



LUND UNIVERSITY

Decision Analysis in Fire Safety Engineering - Analysing Investments in Fire Safety

Tehler, Henrik

2003

[Link to publication](#)

Citation for published version (APA):

Tehler, H. (2003). *Decision Analysis in Fire Safety Engineering - Analysing Investments in Fire Safety*. [Doctoral Thesis (compilation), Division of Fire Safety Engineering]. Department of Fire Safety Engineering and Systems Safety, Lund University.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

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

Decision Analysis in Fire Safety Engineering

Analysing Investments in Fire Safety

Henrik Johansson

Akademisk avhandling

som för avläggande av teknologie doktorsexamen vid tekniska fakulteten vid Lunds universitet offentligen försvaras vid avdelningen för Brandteknik, John Ericssons väg 1, sal V:A, torsdagen den 25 september 2003, kl 13.15.

Decision Analysis in Fire Safety Engineering – Analysing Investments in Fire Safety

Henrik Johansson

Report 1027

ISSN: 1402-3504

ISRN: LUTVDG/TVBB-1027-SE

Number of pages: 330

Cover photograph: Bernt Claesson/Pressens Bild.

Illustrations: Henrik Johansson.

Keywords: Decision analysis, risk analysis, fire safety, epistemic uncertainties, investment appraisal.

Abstract

The application of decision analytical methods to the evaluation of investments in fire safety was investigated, particularly with the aim of being able to suggest a method for analysing a *specific* investment in fire safety for a *specific* factory. Attention was directed above all at the handling of cases of large epistemic uncertainty regarding both probabilities and utilities, Bayesian decision theory serving as a basis for the development of the method. Two extensions of the decision rule used in Bayesian decision theory (the principle of maximising expected utility) were suggested for use in the present context. Together with a model for calculating the expected utility of a specific investment, they provide an evaluatory framework for the analysis of investments in fire safety. The major contributions of the thesis to the area of decision analysis within fire safety engineering are that it provides a better understanding of the use of different decision analytical approaches in a context such as the present one, that it highlights problems of evaluation when large epistemic uncertainties are present, that it suggests a solution for use in such a case, and that it suggests a way in which the reduction in risk can be evaluated in terms of monetary value.

© Copyright: Henrik Johansson and Brandteknik, Lunds tekniska högskola, Lunds universitet, Lund 2003.

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>

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

Summary

The thesis examines the use of decision analytical methods in evaluating investments in fire safety in a specific factory. Particular attention is directed at evaluating the risk reduction that an investment involves in terms of a monetary value and at suggesting how large epistemic uncertainties concerning probabilities and consequences can best be dealt with.

The major contributions of the thesis to the area of decision analysis within fire safety engineering are that it provides a better understanding of the use of different decision analytical approaches in a context such as the present, that it highlights the problems involved in evaluation when large epistemic uncertainties are present, that it suggests a solution to this problem, and that it suggests a way in which the risk reduction achieved can be evaluated in terms of a monetary value.

The central aim of the work presented here was to develop a method allowing investments in fire safety in factories to be evaluated in terms both of the certain costs and benefits the investment involves and of the monetary value of the risk reduction achieved. The focus is on *specific* investments in *specific* buildings, which means that situations may be encountered in which the epistemic uncertainties regarding the probabilities and consequences involved are large due to the limited statistical information available concerning fires in the building of interest. It is assumed that all the fire safety alternatives that are being evaluated comply with the applicable building code and that no evaluation of risk to life is included.

In suggesting a suitable method for evaluating fire safety investments in a context of the type described above, the thesis starts by presenting different decision theories that can provide a basis for a decision analytical method applicable here. Each of the methods taken up is evaluated in terms of how well it can be expected to perform in the present context, *Bayesian decision theory* being deemed to be best here.

Bayesian decision theory is used then to create a model, termed the primary model, for analysing investments in fire safety in a specific building, a model based on quantitative fire risk analysis. That model allows the risk reduction of an investment in fire safety to be expressed in terms of its intrinsic monetary value. This provides a basis for

comparing, in terms of monetary values, the investment in question with some other investment or with the alternative of keeping the building in its present state.

Using Bayesian decision theory as a basis for the primary model implies that only the expected utilities are important in comparing two decision alternatives. Using this traditional decision analytic approach provides no information concerning the robustness of a decision. Investigating the robustness of a decision involves analysing how likely it is that the best alternative would change if additional information regarding the decision problem were received. So as to also provide the decision maker information concerning the robustness of a decision, an additional evaluation of this sort to complement the expected utility evaluation is suggested. This approach is termed extended decision analysis.

In situations in which there is a lack of statistical information concerning the reliability of a specific fire safety system, for example, the decision maker may have difficulties in complying with one of the key assumptions of Bayesian decision theory, that of the decision maker's being able to express probabilities (and utilities) as exact values or as probability distributions. In such cases, one may need therefore, to use some other methods for evaluating the alternatives. One such method found useful here is *Supersoft decision theory* (SSD), which employs the basic type of expected utility evaluation used in Bayesian decision theory but it does not require the decision maker to express probabilities and utilities as exact values. It can thus be used in situations of large epistemic uncertainty. Since it is difficult to know in advance how much information decision makers will have in performing analyses of investments in fire safety, the approach the thesis takes is to suggest an evaluatory framework for such investments, one in which three different evaluation principles are included. Depending on the amount of information a specific decision maker has at hand, he/she can choose to use any of the three methods.

The first method is termed the traditional decision analysis method. It involves employing exact values of probabilities and utilities of the different fire scenarios considered in the analysis. Evaluation of alternatives by use of this method is based on examining the expected utilities of the various decision alternatives. The second method is called the extended decision analysis method. It involves expressing the values of probabilities and utilities as probability distributions. Although the

evaluation of decision alternatives by this method involves use of an expected utility evaluation, in addition to that the decision robustness is also evaluated. The third method finally, is one based on Supersoft decision theory. It involves expressing probabilities and utilities as intervals and then using various criteria (based on expected utilities) for evaluating the investment alternative. Each of these methods can be used in conjunction with the primary model. Together they form the evaluatory framework for the analysis of investments in fire safety suggested in the thesis.

Two case studies were performed for illustrating use of the methods for analysing investments in fire safety discussed and for exploring their practical applicability. The case studies were performed on buildings belonging to the companies ABB and Avesta Sheffield. In both of them, the analyses concerned investments in water sprinkler systems. It was concluded that the method that was employed worked well in practice but required a very significant work effort. This makes it appear likely that, for the method to be very useful practically, it needs to be simplified or to be incorporated into a computer program so that an analysis can be carried out more quickly.

The methods for analysing investments in fire safety presented in the thesis allow for the construction of measures of fire risk that have a concrete decision analytical meaning to the decision maker. With use of such measures, it is possible to update the analysis of the building in question continuously, and thus to continuously update the measure of fire risk as well. A sensible way of doing this is by use of Bayesian networks, which are discussed in the thesis. It is shown how information from fires in the building in question or gleaned from expert judgements can be used to continuously update the measure of fire risk employed in a specific building.

In conclusion, a new method (or methods) for analysing *specific* investments in fire safety in *specific* factories is suggested. Compared with previous suggestions of such methods, it provides a new way of estimating the monetary value of the reduction in risk that an investment in fire safety involves. Most importantly, it explicitly addresses the problem of epistemic uncertainties and can be used to evaluate decision alternatives even when the magnitude of these uncertainties is considerable.

Sammanfattning (Summary in Swedish)

I avhandlingen *Decision Analysis in Fire Safety Engineering – Analysing Investments in Fire Safety* undersöks hur beslutsanalytiska modeller kan användas för att utvärdera olika investeringar i brandskydd för en specifik industribyggnad. En stor del av arbetet behandlar hur den riskreduktion som en investering i brandskydd medför kan värderas i form av ett monetärt värde och hur man kan hantera stora kunskapsosäkerheter rörande sannolikheter och konsekvenser.

Syftet med avhandlingen är att ta fram en metod med vilken man skall kunna utvärdera investeringar i brandskydd i industrier, både i termer av de mer eller mindre säkra kostnaderna (och intäkterna) som en investering innebär och i termer av en monetär värdering av den riskreduktion som investeringen är tänkt att åstadkomma. Denna metod skall vara tillämpbar på *specifika* investeringar i *specifika* industribyggnader vilket, på grund av att mängden statistiska data rörande bränder i en specifik byggnad vanligtvis är liten, ofta innebär att kunskapsosäkerheterna rörande sannolikheter och konsekvenser är stora. Vid användning av denna utvärderingsmetod antas att samtliga investeringsalternativ som analyseras uppfyller kraven i den gällande bygglagstiftningen och att beslutsfattaren inte värderar personsäkerheten i byggnaden då han/hon använder metoden.

För att en metod för utvärdering av investeringar i brandskydd ska kunna föreslås inleds avhandlingen med en presentation av olika beslutsteorier som kan användas som bas för utveckling av en sådan metod. Var och en av beslutsteorierna analyseras med avseende på i vilken utsträckning de förväntas kunna passa för den aktuella typen av beslutsproblem, och slutsatsen från denna analys är att det är den *Bayesianska beslutsteorin* som bedöms som mest lämpad i det här sammanhanget.

Bayesiansk beslutsteori används sedan för att skapa en modell, baserad på en kvantitativ riskanalys, som kallas förväntad nytta-modellen. Denna modell är avsedd att användas vid analys av olika investeringar i brandskydd i en specifik industribyggnad. Genom att använda denna modell kan man i form av ett monetärt värde uttrycka den riskreduktion som följer av en investering i brandskydd. Detta ger möjlighet att jämföra olika investeringsalternativ med hjälp av både de mer eller mindre säkra kostnaderna för investeringen, och den monetära värderingen av riskreduktionen.

Att Bayesiansk beslutsteori används som grund för modellen ovan innebär att det enda som är viktigt vid jämförelse av beslutsalternativ är den förväntade nyttan av vart och ett av alternativen. Om man använder detta traditionella sätt att göra en analys får beslutsfattaren ingen information om hur robust beslutet är. Att undersöka robustheten hos ett beslut innebär att man undersöker hur sannolikt det är att det bästa alternativet ändras om beslutsfattaren skulle erhålla mer information om problemet. För att även förse beslutsfattaren med information om robustheten i beslutssituationen föreslås i avhandlingen att en förväntad nytta-utvärdering kompletteras med en utvärdering av beslutssituationens robusthet.

I situationer där man inte har tillgång till statistisk information rörande till exempel tillförlitligheten hos ett specifikt brandteknisk system kan ett av de viktigaste antagandena i den Bayesianska beslutsteorin, att beslutsfattaren kan uttrycka sannolikheter (och nyttovärden) med hjälp av exakta värden, vara felaktigt. I sådana situationer behöver man alltså använda andra metoder än den Bayesianska beslutsteorin för att analysera beslutssituationen. En sådan metod som bedöms vara användbar i det aktuella sammanhanget är *Hypermjuk beslutsteori*. I Hypermjuk beslutsteori används beräkningar av förväntad nytta vid utvärdering av beslutsalternativen, men denna utvärdering kräver inte att beslutsfattaren kan uttrycka sannolikheter med exakta värden utan tillåter att intervall används. Detta innebär att Hypermjuk beslutsteori kan användas i situationer där det finns mycket stora kunskapsosäkerheter och där beslutsfattaren alltså har svårigheter att ange värden för sannolikheter och konsekvenser.

Eftersom det kan vara svårt att i förväg veta hur mycket information en specifik beslutsfattare kommer att ha tillgång till då han/hon skall analysera investeringar i brandskydd föreslås i avhandlingen tre metoder som kan användas för analys. Vilken av dem som bör användas i en specifik beslutssituation beror på tillgänglig tid och information vid tidpunkten för beslutet.

Den första metoden är baserad på traditionell Bayesiansk beslutsteori och innebär att man använder exakta värden för sannolikheter och konsekvenser. Beslutsalternativen utvärderas med hjälp av en analys av den förväntade nytta som de olika alternativen innebär, och det alternativ som har den högsta förväntade nyttan är det bästa alternativet.

Den andra metoden innebär att kunskapsosäkerheterna rörande sannolikhetsvärden och konsekvenser uttrycks genom sannolikhetsfördelningar. Utvärderingen av beslutsalternativ sker genom att beräkna den förväntade nyttan med alternativen, men dessutom görs en analys av robustheten i beslutssituationen, d.v.s. man undersöker vilken effekt kunskapsosäkerheterna rörande sannolikheter och konsekvenser har på resultatet av analysen.

Den sista metoden bygger på Hypermjuk beslutsteori och innebär att sannolikheter och konsekvenser uttrycks som intervall i stället för som exakta värden. Vid utvärdering av beslutsalternativen används sedan flera kriterier (baserade på den förväntade nyttan) för att avgöra vilket alternativ som är bäst.

Var och en av de tre metoder för utvärdering av investeringar i brandskydd som behandlats kan användas tillsammans med förväntad nytta-modellen som också presenteras i avhandlingen. Utvärderingsmetoderna och förväntad nytta-modellen utgör en ”verktygslåda” från vilken man väljer utvärderingsmetod beroende på de förutsättningar som råder i en specifik beslutssituation.

För att illustrera användningen av metoderna som presenteras i avhandlingen och undersöka huruvida de är praktiskt användbara presenteras två studier där analyser av investeringar i brandskydd genomförts för två industribyggnader som tillhör Avesta Sheffield respektive ABB. I båda fallen analyseras en investering i ett sprinklersystem. Slutsatserna från studierna är att metoderna fungerar tillfredsställande i praktiken, men att tidsåtgången för en analys kan vara mycket stor. Därför är det troligt att metoderna måste förenklas och/eller att ett datorprogram som bygger på metoderna utvecklas för att göra analysarbetet snabbare och lättare.

Genom att utnyttja de metoder för beslutsanalys som utarbetats i avhandlingen kan man skapa riskmått som har en konkret beslutsanalytisk mening för beslutsfattaren. Genom att använda ett sådant mått på brandrisken i en specifik byggnad kan man följa utvecklingen av detta mått i en byggnad genom att kontinuerligt uppdatera analysen. Ett fördelaktigt sätt att göra detta är att använda sig av Bayesianska nätverk. Med hjälp av dessa kan information från bränder som uppstått i den aktuella byggnaden eller från experter användas för att kontinuerligt uppdatera riskmåttet i en specifik byggnad.

Den riskförändring som registreras vid den kontinuerliga uppdateringen kan med hjälp av de beslutsanalysmodeller som presenterats i avhandlingen uttryckas i termer av ett monetärt värde. Detta monetära värde ger beslutsfattaren en uppfattning om hur mycket riskändringen är värd för honom/henne givet att han/hon beslutat sig för att fatta beslut enligt principen om maximerad förväntad nytta.

Avhandlingen presenterar således en ny metod (nya metoder) för att analysera en *specifik* investering i brandskydd i en *specifik* byggnad. Vid jämförelse med tidigare existerande metoder för denna typ av analys bidrar denna metod med ett nytt sätt att uppskatta det monetära värdet av den riskreduktion som en specifik investering i brandskydd innebär. Den mest signifikanta skillnaden mellan den metod som presenteras här och andra metoder är att den explicit behandlar kunskapsosäkerheter rörande sannolikheter och konsekvenser och kan användas för utvärdering av alternativ i beslutssituationer där denna typ av osäkerheter är mycket stora.

TABLE OF CONTENTS

Summary	i
Sammanfattning (Summary in Swedish)	v
Table of contents	ix
1 Introduction	1
Overview and background	2
Objective and aims	6
Normative vs descriptive theories	6
The development of an operational model	7
Overview of the thesis	9
2 Decision analysis	13
The development of decision theory	13
Alternative decision theories	20
<i>Prospect theory</i>	21
<i>Bell's theory</i>	21
<i>Ellsberg's theory</i>	22
<i>Hodges and Lehmann's theory</i>	23
<i>Quiggin's theory</i>	23
<i>Yaari's theory</i>	24
<i>Cumulative Prospect Theory</i>	25
<i>Supersoft decision theory</i>	26
<i>The Delta method</i>	27
3 The application of decision analysis in fire safety engineering	29
Evaluation of decision analysis methods	29
<i>Bayesian decision theory</i>	31
<i>Alternatives theories</i>	33
<i>Conclusions from evaluating different decision analysis methods</i>	39
Suggestion of a primary model for decision analysis	41
<i>Calculating the expected utility of a fire exposure</i>	44
<i>A simpler method for estimating the CE of a fire exposure</i>	54
<i>An illustrative example</i>	55
<i>Risk-adjusted net present value</i>	58
<i>Comments on some of the assumptions</i>	61
Dealing with epistemic uncertainty – prescriptive decision rules	63
<i>Extended decision analysis</i>	68
<i>Supersoft decision theory</i>	78
A framework for the evaluation of investments in fire safety	82

4 Bayesian methods in decision analysis concerned with fire safety	87
Frequency of fire	89
Two possible events	91
More than two possible events	92
Continual updating of fire risk measurement	95
5 Case studies	99
The fire risk analysis model	100
Subjective probabilities	106
Screening procedure	109
The ABB analysis	110
The Avesta Sheffield analysis	120
Comparison of case studies to fire statistics	132
Concluding remarks regarding the case studies	137
6 Summary, conclusions and future work	141
Summary	141
Conclusions	144
Future work	145
7 Acknowledgements	147
8 References	149
Papers	157
Paper 1: Decision analysis concerned with investments in fire safety	157
Paper 2: Investment appraisal using quantitative risk analysis	177
Paper 3: Application of Supersoft decision theory in fire risk assessment	201
Paper 4: A Bayesian network model for the continual updating of fire risk measurement	239
Appendix	265
Appendix A: Bayesian updating of a Dirichlet distribution	265
Appendix B: Estimates pertaining to the ABB analysis	269
Appendix C: Estimates pertaining to the Avesta Sheffield analysis	285
Appendix D: Event trees pertaining to the ABB analysis	299
Appendix E: Event trees pertaining to the Avesta Sheffield analysis	309
Appendix F: Computer codes	317

1 Introduction

Determining whether one should make a particular investment in fire safety is not an easy task since in making such a decision one needs to assess whether the benefits of the investment suffice to compensate for the costs associated with it. Judgements regarding the benefits of an investment of this sort are difficult to make due both to uncertainty regarding whether any fire or fires in the building of interest will occur and to uncertainty regarding the extent to which any fire that starts will spread.

Two questions can be seen as related to this problem. One is how much the risk will be reduced if the investment is made. The other is how much this risk reduction is worth. Quantitative risk analysis, which would provide a measure of how much the risk is reduced, could be used in efforts to answer the first question. Various measures of this sort have been presented by Hall and Sekizawa [1]. To obtain an answer to the second question, one needs to consider the values the decision maker assigns to possible outcomes that are considered. Using decision analysis to analyse investments in fire safety represents a formal way of using the decision maker's assessments and values to arrive at a recommendation of the decision alternative the decision maker should select. The thesis deals with the use of decision theory as the basis for a decision analytic method to be employed for evaluating possible investments in fire safety in a specific factory.

Thus, the situation that is dealt with in the thesis is the following: A decision maker, which in this case is a company (or person/group at a company), needs to choose between a set of possible alternatives for the fire protection to be installed in a specific factory (he/she also has the choice of not installing any additional fire protection in the building). He/she wants in so doing, to evaluate the benefits in terms of the risk reduction the different investments would involve so as to be able to compare the alternatives in terms both of the costs the alternatives involve and their benefits.

It is important to point out the differences between the work presented in the thesis and work carried out in the area of decision analysis in the area of fire safety engineering previously. The main differences can be seen as being related to interest being directed here at *specific* investments in

a *specific* building. Since the number of fires in any particular building is not very large, information regarding fires in a specific building is likely to be scarce. This can result in a high degree of epistemic uncertainty regarding the probabilities of different fire scenarios, which may cause problems if one uses traditional methods of decision analysis. The main difference between the work presented here and the previous work as reviewed in the next section is that the method for evaluating investments in fire safety suggested in the thesis is applicable to situations in which there is a high degree of epistemic uncertainty (the probabilities and utilities cannot be assigned precise values), which means that the other methods may not be suitable. Such situations are likely to arise when one is concerned with a specific investment in a specific building.

Overview and background

The thesis is the result of a project carried out at the Department of Fire Safety Engineering at Lund University. The project, started in 1998 and entitled “Economic optimisation of the industrial fire protection” (“Ekonomisk optimering av det industriella brandskyddet”), is one financed by the Swedish Fire Research Board (BRANDFORSK).

Taking the problems involved in evaluating investments in fire safety as its starting point, the project aimed at developing a method that could be used to assist companies in evaluating investments in fire safety. The method was to be one that would be applicable not simply to factories in general but to specific factories and would include a monetary evaluation of the reduction in fire risk that the fire safety investment considered would provide.

Previous research in fire safety engineering, specifically in the analysis of fire safety alternatives in a particular building, was concerned largely with fire risk analysis models and much less with decision analysis. This can be thought to soon change, since many of the situations in which fire risk analyses are performed can also be seen as being decision situations, involving the question for example of whether a particular level of fire risk is acceptable. A number of distinguished authors have also taken up matters of decision analysis in fire safety engineering recently, such as Donegan in the latest edition of the SFPE handbook [2] and Ramachandran [3].

Prior to the work of Ramachandran and of Donegan, representing the application of decision analysis to fire safety engineering generally, Shpilberg and De Neufville [4] presented a study in 1974 in which they analysed what fire protection measures were best for airports. Not long thereafter, in 1978, Cozzolino [5] described a method involving decision analysis intended for use in evaluating different deductible and aggregate retention levels. In 1979 the National Bureau of Standards issued a report on decision analysis concerned with different strategies for reducing upholstered furniture fire losses [6]. Seven years later, in 1986, three papers presented at the first International Symposium on Fire Safety Science [7], [8] and [9], concerned to various degrees with the practical application of decision analysis, appeared. Since then, papers on a variety of methods of fire safety engineering and decision analysis and their application have been presented. Watts [10], [11] and [12] reviews application of various multi attribute methods (index methods) for the evaluation of different fire protection alternatives for a building. Note that, although several index methods for the evaluation of fire risk have been proposed (such as the Gretener method, [13]), these are not taken up further here since they are difficult to use when one wants to determine the monetary worth of a specific risk reduction. Budnick *et al.* [14] showed, in addition, how the Analytic Hierarchy Process (AHP) method can be used to compare the risk involved in different fire safety alternatives employed in telecommunications central office facilities. Takeda and Young [15] considered the basis for the development of a cost-risk model that can be used to select a cost-efficient fire protection alternative for use in a building. Beck presented later a method termed CESARE-RISK [16] and Benichou and Yung a model termed FIRECAM [17], both of which were intended for use in evaluating different alternative approaches to fire protection in a building.

One can classify the methods taken up above according to whether or not they are applicable to some specific building, and also according to the type of decision rule used in evaluating the decision alternatives. A decision rule provides a way of determining which of a set of alternatives is best. The methods taken up here can be divided into three classes of decision rules: index methods, expected-cost methods, and expected-utility methods. Index methods use some type of index value to determine the alternative which is best. The index value is usually obtained by a weighting procedure based on various *characteristics* of the alternatives. Evaluation of alternatives using the expected-cost methods, in turn, involves considering the possible outcomes of choosing

a particular alternative (instead of the characteristics of the alternatives as such), evaluation being based on the expected costs associated with choosing a particular alternative. The expected-utility methods, finally, consider the outcome of choosing a specific decision alternative. An important difference between the expected-cost methods and the expected-utility methods is that the latter methods allow one to take account of the decision maker's risk attitude, whereas this is not possible with use of an expected-cost method. Also, in the present context the logical basis for use of the expected-cost method's decision rule is somewhat weaker than that for use of the expected-utility method's decision rule (see the beginning of the next chapter and [18] (Paper 4)). The classification of the methods is shown in Table 1. Note that some of the references referred to above, since they discuss the use of decision analysis in fire safety engineering in only general terms, cannot be considered to provide any method for practical application here.

Table 1 Classification of various decision analysis methods employed earlier in fire safety engineering.

	Index method	Expected-cost method	Expected-utility method
Buildings in general		Policy analysis (see [6], for example)	Spilberg and De Neuville [4]
Specific buildings	Watts [10], [11], and [12] Budnick <i>et al.</i> [14]	CESARE-RISK [16] FiRECAM [17]	Ramachandran [3] Cozzolino [5] Van Anne [8], [19]

In the present context, multiattribute evaluation methods, also termed index-methods, are difficult to employ since they fail to address the uncertainty inherent in the problem, i.e. the uncertainty of whether a fire will occur and the extent to which a fire that occurred would spread. Since such methods are unable to do this, it is difficult to use them to evaluate the possibility of fires occurring in a specific building in terms of monetary value, which in the present context would be desirable.

Although the expected-cost methods do address the uncertainty regarding the occurrence and development of fire, in using them one assumes that the decision maker is risk-neutral, which is an assumption that can be questioned for decisions involving the possibility of large

losses¹ for a firm. For public investment decisions in which risk-neutrality is assumed, these methods are usually considered to be well motivated since the losses to the individual taxpayer which a single public investment project could cause are small and since people are assumed to be risk-neutral towards possible losses that are small [20]. Thus, the expected cost methods are more appropriate to use for public investment decisions than for decisions involving investments for a specific firm.

Of the methods taken up above, those described by Ramachandran [3], by Cozzolino [5] and by Van Anne ([8] and [19]) are the ones most suitable in the present context since they focus on specific buildings and employ the rule of maximising expected utility, which allows consideration of risk attitude to be taken. These methods have certain drawbacks, however, concerned primarily with their treatment (or lack of treatment) of ambiguity regarding estimates of the probabilities and of the consequences. Ambiguity concerning probabilities and consequences is likely to be very common when investments in specific factories are to be analysed. This is due to the lack of specific statistical information regarding many of the probabilities of interest, such as the probability that the employees of a factory will succeed in extinguishing a fire. Information regarding such probabilities may not be available in any other form than through the consultation of experts. Thus, the thesis is concerned primarily with the development of a method able to deal with a high degree of epistemic uncertainty regarding the probabilities and consequences involved.

A further difference between the work presented in the thesis and that of Ramachandran's in developing his methods is that he assumed the decision maker's utility function for the outcome of a fire to have a specific form and that he did not explicitly address the (negative) utility of more than one fire occurring in the building of interest during a specific period of time. Calculating the expected utility of the possibility of suffering the losses of fires in a specific building during a specific period of time requires some rather strong assumptions regarding the

¹ Laughunn *et al.* [21] investigated the risk attitude of 224 managers towards below-target returns, finding a majority of the managers to not be risk-neutral. Spetzler [22] found that for high-risk investment projects the majority of the participants were risk averse.

decision maker's utility function, which are not considered in detail by Ramachandran. Cozzolino assumes that the decision maker's utility function for losses has an exponential form, which allows him/her to calculate the expected utility of a series of fires. Assumptions of this type are dealt with in the present thesis, a set of assumptions that lead to a fairly simple way of calculating the expected utility of an investment in fire safety being presented.

Objective and aims

The major objective of the thesis is to present and discuss a method for the evaluation of fire safety investments in a *specific* factory building or complex, a method that includes evaluation of investments in terms of the risk they involve and can be used to evaluate the risk in monetary terms. The method needs to also be constructed in such a way that it is practically applicable to situations in which the information available regarding the probabilities relating to fire and the consequences a fire would have is both limited and uncertain.

The thesis also aims at providing a basis for further research in the area of prescriptive decision analysis concerned with fire safety.

Normative vs descriptive theories

A fundamental distinction highly important in the present context is that between normative and descriptive decision theories. A descriptive decision theory seeks to *describe* the world *as it is*. For example, a model of how people tend to make decisions regarding fire safety investments is a descriptive model. The quality of such a model is determined by the extent to which it accurately predicts the behaviour of people, given a particular context.

A normative decision theory, in contrast, sets up rules for *how things should be*, i.e. how people should make decisions. Also, a normative theory is concerned with completely analysed alternatives, meaning that all the outcomes are known, such as those found in gambles and lotteries.

Prescriptive models are another type of decision analytical models. These are models concerned with *helping* people make well informed,

and thus hopefully better, decisions *in practice* [23]. Such models might be ones, for example, to help decision makers avoid common mistakes in decision making, such as the inclusion of sunk costs, or failing to recognise dominance [24]. The thesis deals primarily with the use of prescriptive models in decision analysis concerned with investments in fire safety. Normative models are used as a point of departure here for creating such models.

One problem that can be encountered in connection with normative and prescriptive models concerns the validation of them. Using the same approach as applied to descriptive models, that of measuring the extent to which a model actually describes the world, would not be feasible here since what one is attempting to do is not to model the world. How then can one ensure the quality of a model one develops?

One way of judging the quality of a prescriptive decision model would be to determine the extent to which it possesses certain desirable properties. Avoiding the inclusion of sunk costs and not recommending decision alternatives that are stochastically dominated are two such desirable properties. Thus, showing that a specific method may lead to such undesirable behaviour is a strong argument against use of that method. In addition to noting desirable properties concerned with evaluation of the different decision alternatives with use of a given model, one can formulate a set of objectives for a prescriptive decision analysis model as a whole. Keeney [23] provides a list of four objectives concerned with selecting a decision analysis model: (1) address problem complexities explicitly, (2) provide a logically sound foundation for analysis, (3) provide for a practical analysis and (4) be open for evaluation and appraisal. The objectives formulated in the present context differ somewhat from these (see chapter 3), although in principle they resemble rather much the objectives that Keeney refers to. Thus, the quality of the method presented in the thesis can be judged about equally well in terms of criteria of the type that Keeney presents and of the criteria provided in chapter 3.

The development of an operational model

The model suggested here for practical use in decision making concerned with investments in fire safety will be termed the operational model. In presenting such a model, the approach taken is to start with a

normative decision theory, i.e. one that shows how people should make decisions in order to be consistent with a set of rules, seen as logically correct, and then to develop from it a model adapted to practical use in the present context.

In employing a normative theory in practice, it can be sensible to introduce certain approximations appropriate to the problem at hand. Since employing a particular normative model might possibly require a degree of time and effort, for example, not regarded as being justified in terms of the end result, one might well seek an approximation of that model. Another reason for employing an approximation of this sort could be that the results the normative theory generated failed to provide all the information the decision maker needed. Such is the case here, as will be discussed later in the thesis. In short, this has to do with the question of how much information the decision maker receives regarding the epistemic uncertainties inherent in the decision problem. Since the information made use of in applying the basic normative model suggested here may not provide the decision maker with all the information he/she requires, it may be seen as necessary to complement the normative decision rule involved by using a set of additional rules as well. The normative decision rule, when combined with a complementary method for evaluating the decision alternatives, result in what is termed here the prescriptive model (an illustration of this is provided in Figure 1). The prescriptive model should be viewed as a *help* for the decision maker when faced with particularly difficult decisions regarding investments in fire safety. In contrast, a normative theory can be seen as a model that dictates how one *should* make a decision, any choice of an alternative except the one dictated by the normative decision rule being considered “irrational”. Even when certain assumptions regarding the normative model are made in efforts to create as adequate a prescriptive model as possible, the model may still not be appropriate for practical application due for example to the time and effort application of it requires. Thus, one may need to make additional assumptions, particularly for making the model useful for dealing with a *specific type* of decision problem, in the present case in choosing between different investments in fire safety. The resulting model is termed the operational model. At the same time, since in applying the operational model to a *specific* decision problem there may be a need for still further approximations, the model used in practice may not be identical with the operational model. This line of thinking, which is adopted in the thesis, is illustrated in Figure 1. The first box there, which

deals with the choice of a reasonable normative theory to base the decision analysis model on, is taken up in the second chapter of the thesis and the beginning of the third. The second box, concerned with the prescriptive model, i.e. the model suggested to be used as a *help* to the decision maker in arriving at concrete decisions regarding investments in fire safety, is dealt with in the third chapter. Since the prescriptive model there applies only to the final evaluation of the decision alternatives, an operational model is needed to indicate how one can analyse the problem in such a way that the prescriptive decision rule can be adequately applied. Chapter 3 is concerned with this model as well. Chapter 4 takes up further aspects of the operational model. The case studies dealt with in chapter 5 illustrate finally what the last box in Figure 1 is concerned with, two examples of how the decision analytical framework suggested can be applied in practice being provided.




Figure 1 *Scheme showing the influence of the different models/theories taken up here.*

Since the thesis deals with all the steps shown in Figure 1, extending from the choice of a suitable normative model that a prescriptive model can be based on to the model's practical application, it takes up matters of relevance both to those concerned with the theoretical foundations of the model and with its application in practice.

Note that the thesis deals with the evaluation of fire risk in monetary terms and that risk to life is not included here. It would be theoretically straightforward to extend the method so as to also include an evaluation of risk of this sort. This is outside the scope of the thesis, however, the method suggested only being intended for use in evaluating the direct and consequential losses due to fire.

Overview of the thesis

A central part of any decision theory is the decision rule employed, indicating how decision alternatives are to be evaluated. The second chapter of the thesis provides a survey of various decision theories and

of the decision rules associated with them, theories and rules which it would appear could serve as a possible basis for a decision analysis model useful in the present context. A further aim of the chapter is to present an overview of the foundations of Bayesian decision theory, one of the decision theories most commonly employed today. This is a very important part of the thesis since it provides a basis for the decision analysis method which is suggested here.

The third chapter starts by comparing carefully the different decision theories seen as potentially useful in the present context and selecting one of them, in terms of the characteristics it possesses, as the basis for a decision model. A set of additional evaluation procedures to be used in conjunction with this model are also suggested. The part of the chapter that follows then concerns what is termed an operational decision model, or the decision analysis model intended for use in practice, also designated as the primary model. The operational assumptions the model involves are discussed. The last part of the chapter deals with decision analysis carried out under conditions of only “vague” information (large epistemic uncertainties) being available, a theoretical and practical framework for evaluating decision problems of this sort concerning investments in fire safety being presented. It consists both of the primary model, which although it can be used in isolation for the evaluation of investments does not in itself indicate in any way how the epistemic uncertainties contained in the model affect the end result. For this reason, two other methods, the extended decision analysis method and the Supersoft decision theory method, are also included in the practical framework employed. The idea is that if a decision maker cannot express the probabilities and consequences of different fire scenarios as exact values, which the primary model alone would require, either of the other two methods can be employed alongside it. The extended decision analysis method requires that the decision maker be able to express his/her uncertainty regarding the probabilities and consequences in terms of probability distributions. If the information available regarding the decision problem at hand is so vague that the decision maker cannot express his/her assessments using probability distributions, use can be made of Supersoft decision theory, evaluations based on it not requiring that probabilities or consequences be expressed precisely or in terms of probability distributions.

The most common situation in analysing investments in fire safety is probably one involving access to only limited statistical information. For

example, there may be only few fires one has observed in the building in question. Provided the building has not been changed since the fires occurred, information regarding these fires is likely to be more relevant than general information regarding the group or category of buildings to which this building belongs. It is thus important to use that information. In chapter four, various probabilities and frequencies of interest in the present context are discussed, together with Bayesian methods for utilising evidence regarding them. In addition, a discussion of how Bayesian networks can be used to measure changes in fire risk continuously is included in the chapter.

In chapter five, two case studies performed using the decision analysis method suggested in the thesis are presented. These studies were performed for the companies ABB and Avesta Sheffield, in both cases the possible investment in a water sprinkler system being analysed. In both case studies, a simple risk analysis model representing the development of a fire in a factory had to be created. The first part of the chapter deals with this model and the last part deal with the two case studies themselves.

Chapter six provides a summary of the work presented in the thesis and presents various conclusions regarding the practical applicability of the method employed. A discussion of possible future research within the present context is also included in this chapter.

Four papers highlighting various aspects of the methods presented are also included in the thesis. Although the overlap between the papers and the material in the chapters referred to above is substantial, the papers are included since they focus on special aspects of the method and provide deeper insight into certain areas.

The first paper, “Decision analysis concerned with investments in fire safety” discusses the background of the decision analysis method taken up, dealing in detail with the extended decision analysis method, discussing such important concepts as decision robustness and uncompensated losses. Also, the paper summarises some of the more important criticisms that have been directed at the maximisation of expected utility decision rule, the rule that forms the basis for the method suggested in the thesis.

The second paper “Investment appraisal using quantitative risk analysis” focuses on the primary model for the evaluation of fire exposures, concerned with the possibility of having an unknown number of fires, each of unknown outcome, during a particular period of time. The information contained in this paper can also be found in chapter 3.

The third paper “Application of Supersoft decision theory in fire risk assessment” deals with the evaluation of decision situations involving a high degree of epistemic uncertainty. In this paper, the application of Supersoft decision theory in a fire safety context is discussed, its use being compared with other methods of decision analysis. Supersoft decision theory is important in the present context since it provides a tool for analysing investments in fire safety even when a high degree of uncertainty exists, i.e. in situations in which the decision maker cannot express probabilities and consequences by use of exact values or distributions.

The fourth paper, “A Bayesian network model for the continual updating of fire risk measurement”, deals with how the decision analysis method suggested can be used together with Bayesian networks for providing a means of continually updating a measure of fire risk in a building.

2 Decision analysis

This chapter provides a brief background of decision analysis, starting with the development of decision theory. It includes accounts of the most commonly used decision rule, the maximisation of expected utility, as well as of various other decision rules. The aim is not to provide a complete account of decision analysis or decision theory but to introduce the reader to the area generally and to point out some of the more important aspects of decision analysis within the present context. Note that since the thesis deals with normative/prescriptive models for decision analysis, this section focuses on models of these types.

The development of decision theory

Decision theory concerns how decisions are made or ought to be made [25]. An early analysis of decisions under risk was conceived within the context of fair gambles [26], pertaining to how much one should pay in order to participate in a particular game. This question is not trivial since there is uncertainty regarding the outcome of any game, some kind of rule being needed for determining whether one should participate in a game for a specific price.

At the beginning of the 18th century, the most natural rule for determining this was the rule of the maximisation of expected value (*MEV*), which considers the expected monetary return of a game to be the price one should be willing to pay in order to participate in the game. Denote the monetary outcome of a game as x , the probability of a particular outcome x_i , as p_i . The game (L) can be described then as $L = (p_1x_1, p_2x_2, \dots, p_nx_n)$. The expected monetary value (*EMV*) of this game would be the sum of each of the products of the probability of a particular outcome and the value of the outcome (equation (2.1)).

$$EMV = p_1 \cdot x_1 + p_2 \cdot x_2 + \dots + p_n \cdot x_n \quad (2.1)$$

The MEV rule can well seem plausible, since in playing a game a large number of times one can expect to end up, *on average*, receiving the expected monetary value per game. However, one can question whether it is reasonable to use the rule when one only plans to play the game a small number of times. Furthermore, there is the so-called St. Petersburg paradox which shows that the MEV rule can lead to unreasonable

results. The St Petersburg paradox concerns a game in which a “fair” coin is thrown until a head appears. The gambler receives $\$2^n$ if the first head appears on the n th throw. The probability of this occurring is $(\frac{1}{2})^n$. Thus, the expected monetary value of the game is $2(\frac{1}{2}) + 4(\frac{1}{4}) + 8(\frac{1}{8}) + \dots = 1 + 1 + 1 + \dots$, which is an infinite number. According to the MEV rule, a person should be willing to give up his entire fortune for participating in this game, which does not seem reasonable.

In 1738 Bernoulli [27] presented the idea that instead of maximising the expected monetary value one should maximise the expected *intrinsic* monetary value, or what today is termed utility. Bernoulli postulated that the utility of one’s money as a whole increases as one acquires more of it, but that it does so at a decreasing rate. Thus, logarithms, for example, can be used to describe the utility of money, the utility of n dollars then being $\log n$. In these terms, the expected utility of the game reported in the St. Petersburg paradox is $E(U) = (\frac{1}{2}) \log 2 + (\frac{1}{4}) \log 4 + (\frac{1}{8}) \log 8 + \dots$, which, at its limits, approaches a finite value. The price that a player should pay in order to participate in the game is m dollars, where $\log m = E(U)$. This exemplifies the maximisation of expected utility rule (MEU), one of the most popular decision rules.

Although Bernoulli’s expected utility rule seemed *reasonable* at the time, it was not until the 20th Century that it received firmer support. Several authors, such as von Neumann and Morgenstern [28] and Savage [29], showed that if the decision maker is willing to accept a number of specific rules (axioms) for his/her preferences between uncertain situations, he/she will *act as if* maximising the expected utility. This was a milestone in the history of decision theory since a logical basis for the intuitively reasonable MEU rule was thus provided².

Note that the utility concept used by Bernoulli (so-called Bernoullian utility [30]) differs from that used by von Neumann and Morgenstern. Bernoullian utility is based on riskless preferences, whereas utility as conceived by von Neumann and Morgenstern is based on comparisons of risky prospects.

² One should note that the axiomatic systems just referred to have been criticised for not providing the logical support just mentioned [61]. This is taken up in a later part of the thesis.

The axioms of expected utility are so constructed that they can be seen as rules guiding a rational person in his/her decision making. Besides the authors mentioned above, there are various other authors who have constructed similar axiomatic systems, such as Herstein and Milnor [31], Oddie and Milne [32] and Krantz *et al.* [33]. The following is an example of an axiomatic system for expected utility consisting of six axioms (the axioms listed are similar to those found in [26]).

1. Ordering

A set of outcomes can be ordered using a “preference or indifference” ordering. For example, a decision maker’s preferences for two outcomes o_1 and o_2 can be described as being either o_1 is preferred to o_2 , that o_2 is preferred to o_1 , or that the decision maker is indifferent between them.

2. Reduction of compound lotteries

A decision maker is indifferent between a complicated compound game and an equivalent game involving only a simple uncertain event, the equivalence of which is determined on the basis of standard probability manipulations. This axiom is illustrated in Figure 2. There, the lottery on the left is somewhat more complicated than the one on the right. According to this axiom, a decision maker should be indifferent between the two lotteries (indifference is represented by “ \sim ”) since the probabilities of the outcomes are the same for both.




Figure 2 Illustration of the second axiom, "Reduction of compound lotteries".

3. Continuity

A decision maker is indifferent between the outcome o_i of a particular gamble and the outcome of an equivalent gamble involving only the best and the worst outcome of the first gamble. Let $L_I = (p_1 o_1, \dots, p_i o_i, \dots, p_n o_n)$ denote a gamble in which o_1 is better than o_i , which in turn is better than o_n ($o_1 > \dots > o_i > \dots > o_n$). The continuity axiom implies that there is a probability u_i such that for a particular probability value the decision maker is indifferent in choosing between o_i and $(u_i o_1, (1-u_i) o_n)$ for some value u_i .

4. Substitutability

A decision maker is indifferent between any lottery involving outcome o_i and a lottery in which o_i is replaced by a lottery that is judged to be equivalent to o_i . For example, if $L = (p_1 o_1, \dots, p_i o_i, \dots, p_n o_n)$, $L_i = (u_i o_1, (1-u_i) o_n)$, and $L_i \sim o_i$, then $L \sim (p_1 o_1, \dots, p_i L_i, \dots, p_n o_n)$.

5. Transitivity

Preference and indifference between gambles are transitive relations. This means that if a decision maker prefers alternative L_1 to alternative L_2 , and alternative L_2 to alternative L_3 , then he/she must prefer alternative L_1 to alternative L_3 .

6. Monotonicity

A decision maker prefers the game $(p_1 o_1, (1-p_1) o_n)$ to game $(p_2 o_1, (1-p_2) o_n)$ or is indifferent between the two if and only if $p_1 \geq p_2$. This axiom seems very reasonable, since if one has to choose between two lotteries each of them involving the same two outcomes, one would choose the lottery in which the probability of winning the better prize is higher.

Assume one would like to compare the lottery $L_I = (p_1 o_1, p_2 o_2, p_3 o_3)$ with another lottery. The first thing to do would be to create a lottery the same as L_I except that, instead of having all the o_i s as prizes, one has gambles involving simply the best outcome (o_1) and the worst (o_3). The continuity axiom states that for each of the prizes there is a gamble that is equivalent to it. The substitutability axiom, in turn, states there to be indifference between the original lottery and a lottery in which the prizes have been exchanged for such equivalent gambles (see Figure 3).




Figure 3 *Diagrammatic representation of the original lottery and of a lottery that is constructed in a manner such that a decision maker should be indifferent between the two lotteries (according to axioms 3 and 4).*

Applying axiom 2 makes it possible to create a lottery (see Figure 4) such that the decision maker should be indifferent between that lottery and the lottery on the right side in Figure 3.




Figure 4 *Illustration of a lottery.*

If one would like to compare the original lottery (the one on the left in Figure 3) with any other lottery, one needs to find a lottery of the structure shown in Figure 4 such that the decision maker is indifferent to it and the lottery that one wishes to compare the original lottery with. This can be achieved using the same technique as demonstrated above. Comparing the two lotteries (of the form shown in Figure 4) would result in one of them being deemed the best or in the decision maker being indifferent between them (according to axiom 6). On the basis of the transitivity axiom, this result can then easily be applied to the two lotteries that one wanted originally to compare.

In Figure 4 one can see that if the u_i 's are called utilities the rule of maximising expected utility (MEU) is in agreement with the axioms. Equation (2.2) provides the formula for calculating the expected utility ($E(U)$) of an uncertain situation, where p_i is the probability of outcome o_i and $U(o_i)$ is the utility associated with that outcome. $V(p_1 o_1, \dots, p_n o_n)$ is the value of an uncertain situation.

$$V(p_1 o_1, \dots, p_n o_n) = E(U) = \sum_{i=1}^n p_i \cdot U(o_i) \quad (2.2)$$

Note that the probabilities von Neumann and Morgenstern [28] employed were connected with games and can be viewed as being “objective”, whereas the probabilities Savage [29] employs are subjective. Savage is regarded as a principal founder of modern decision theory [34], also termed Bayesian decision theory. He provides axioms of a type similar to those presented earlier, using them to derive the decision rule of maximising expected utility (MEU). His axioms, however, as has already been pointed out, lead to probabilities being treated as subjective.

The subjective interpretation of probability employed in Bayesian decision theory involves the probability of a particular event being seen as reflecting the decision maker's choices between uncertain gambles involving the event in question. For example, if a decision maker is indifferent between receiving prize a for certain and a lottery involving his/her receiving prize b if some given event E occurs and prize c if E does not occur, then the probability of event E , or p_E , is defined according to equation (2.3) [35], where $U(a)$ is the utility associated with consequence a , $U(b)$ is the utility associated with consequence b and $U(c)$ is the utility associated with consequence c . Note that $U(b) > U(a) > U(c)$.

$$p_E = [U(a) - U(c)] / [U(b) - U(c)] \quad (2.3)$$

Thus, the probability of an event E is defined in terms of the decision maker's choices among uncertain situations that involve the event E . In practice, obtaining estimates of the probability of event E can be performed by simply having the decision maker state a probability value, or if the decision maker feels uncomfortable in doing this, asking him/her questions concerning choices between various uncertain

situations (see for example, [36]). Assume, for example, that the decision maker is uncertain about whether a particular sprinkler system would extinguish a fire that occurred at a specific place in the building in question. The analyst could ask the decision maker then which of two lotteries (uncertain situations) of the following type he/she would prefer. In the one situation, the decision maker would receive a prize a if the event of interest occurred, i.e. if the sprinkler succeed in extinguishing the fire, and he/she would otherwise receive nothing. In the other situation, the decision maker would draw a ball from an urn containing a known proportion of black and of white balls, for example 90 white and 10 black. If the decision maker drew a white ball, he/she would receive prize a , and would otherwise receive nothing. If the decision maker were indifferent between these two uncertain situations, his/her subjective probability of the event occurring that the sprinkler system succeeded in extinguishing the fire would be 0.9. If the decision maker preferred the uncertain situation involving the urn, the analyst could continue asking questions regarding the uncertain situation while reducing the number and thus proportion of white balls in the urn until the decision maker was indifferent between the two situations. If instead the decision maker preferred the uncertain situation involving the event of interest in the first place (i.e. the event of the sprinkler system succeeding in extinguishing the fire) the number of white balls could be increased until the decision maker was indifferent between the uncertain situations.

There are several methods for eliciting probabilistic judgement (see, for example, Edwards and von Winterfeldt, [37]). In the present thesis, no final assessment of what method of estimating probabilities is most suitable will be made. In the case studies presented, a very simple procedure involving the direct assessment of probabilities by the analyst together with some members of the personnel of the factory in question is employed.

Note that in Bayesian decision theory it is assumed that the decision maker can find a specific probability value for which he/she is indifferent between alternatives of the type described above. This is an assumption that a decision maker may have difficulties to comply with. Assume, for example, that the probability that the price of one US dollar will be 10 Swedish crowns or more on January 1, 2050 is to be estimated. It is highly doubtful that a decision maker would be able to assign a precise value to such a probability and, if the decision maker is unable to do this, one of the key assumptions of Bayesian decision

theory is not valid here. The Bayesian solution to the problem, however, is to create a probability distribution over the possible values of the probability in question to represent the uncertainty regarding that value. Creating such a distribution may require considerable resources, however, in terms both of time and of information, resources the decision maker may not have. The thesis deals with this problem and aims primarily at suggesting methods for the decision analysis of investments in fire safety in which the decision maker may have only vague or imprecise information concerning the probabilities (and utilities) of interest.

Alternative decision theories

Bayesian decision theory is one of the theories most commonly employed in decision analysis. It has been substantially criticised, however, from both a normative and a descriptive standpoint. Some of the most important criticisms involved are discussed in [38] (Paper 1) and their implication for the use of the Bayesian decision theory in the present fire safety engineering context will be discussed further here.

Although the axiomatic system of the basic type referred to in the previous section can possibly be regarded as being the oldest and most influential one within decision theory, other theories are also available. A number of theories have been developed on the basis of criticisms directed at Bayesian decision theory. The criticisms stems from empirical investigations in which people have been found to not behave in accordance with Bayesian decision theory, violating some of its axioms (see [34], [39] and [40]). These criticisms have led to the development of theories that have relaxed some of the assumptions of Bayesian decision theory so as to be able to better explain people's behaviour. Examples of theories of this type are Kahneman and Tversky's Prospect theory [40] and the theory presented by Bell [41]. Some of the better known alternative decision theories will be presented in the present section, providing an opportunity to determine whether they possess some properties desirable in the present context. In doing this, emphasis will be placed on the *decision criterion* employed in each, i.e. the way in which different decision alternatives are compared.

Prospect theory

In evaluating different decision alternatives in terms of Prospect theory [40], use is made of a structure similar to that employed in evaluating alternatives by use of the MEU criterion. In Prospect theory, however, a weighting-function for the probabilities found in the decision problem is introduced. An uncertain alternative that can result in either a positive or a negative outcome is evaluated according to equation (2.4). There, V_I is the value of the decision alternative, the p_i s are the probabilities of the outcomes, denoted as o_i , π is the weighting function for the probabilities and v is the value of the respective outcomes.

$$V_I(p_1 o_1, p_2 o_2) = \pi(p_1) \cdot v(o_1) + \pi(p_2) \cdot v(o_2) \quad (2.4)$$

The evaluation criterion for decisions involving only positive or only negative outcomes is given in equation (2.5). Note that $p_1 + p_2 = 1$ and that either $o_2 > o_1 > 0$ or $o_2 < o_1 < 0$.

$$V_2(p_1 o_1, p_2 o_2) = v(o_1) + \pi(p_2) \cdot (v(o_2) - v(o_1)) \quad (2.5)$$

Note that Prospect theory deals with so called objective or standard probabilities. The authors state, however, that the theory could be extended to encompass situations in which the probabilities are not given ([40], p.288). In Prospect theory each outcome is assigned a value relative to a reference point which has the value of 0.

Bell's theory

Bell [41] uses the von Neumann and Morgenstern theory [28], but instead of basing the utility of an outcome on the decision maker's final asset position, Bell suggests that the utility of an outcome be based, not simply on the final asset position, but also on the regret one may feel due to the potential asset one has given up by making the decision.

Assume that a decision maker has long been betting on a particular number in a lottery and needs to decide whether to continue participating in the lottery and betting on the same number, or to not bet at all. According to Bell, the outcome of the decision maker's not betting but the number he/she has been betting on earlier winning should be evaluated as being worse than his/her final asset position alone.

Bell denotes one's final assets as X and the *foregone assets* as Y . The evaluation of one of two decision alternatives involving two uncertain outcomes is described in equation (2.6). If the decision maker chooses the alternative that is analysed in the equation, he/she will receive outcome o_1 with probability p_1 and outcome o_2 with probability p_2 , whereas if the decision maker chooses the other alternative, he/she will receive outcome o_3 with probability p_1 and outcome o_4 with probability p_2 . There, $U(o_i, o_j)$ is the utility of a particular outcome o_i , given that the decision maker would have received outcome o_j if he/she had chosen the other alternative.

$$V(p_1 o_1, p_2 o_2) = p_1 \cdot U(o_1, o_3) + p_2 \cdot U(o_2, o_4) \quad (2.6)$$

Note that Loomes and Sugden [42] present a similar theory in which decision alternatives are evaluated on the basis of their "expected modified utility". The modification referred to is that of the utility of a particular outcome, a utility which is modified in accordance with the regret the decision maker would be expected to feel due to not being able to receive any of the consequences associated with the other alternative(s) if the alternative in question is chosen.

Ellsberg's theory

Ellsberg [39] suggests a decision theory for choices between uncertain alternatives, for which the outcomes can be assigned "von Neumann-Morgenstern utilities" [39]. Ellsberg relaxes the assumption that a decision maker can assign a precise probability distribution defined on the potential outcomes of the decision. Instead, Ellsberg assumes that a decision maker can estimate a set of distributions Y^0 seen as being "reasonable", and that the decision maker can also estimate a specific probability distribution termed y^0 , representing his/her "estimate". This is the distribution he/she would use if having to choose some specific distribution. The decision maker is assumed to also be able to assign a value ρ representing his/her degree of confidence in the probability distribution selected (y^0). The degree of confidence is expressed here as a value between 0 and 1.

Ellsberg suggests that in evaluating decision alternatives one use a combination of the expected utility of choosing a particular decision alternative, the distributions of which corresponds to the y^0 distribution,

and of the lowest expected utility of choosing the same alternative, provided the distribution employed in calculating the expected utility is within the set of probability distributions Y^0 seen as reasonable. Let $E(U_i)$ denote the expected utility of choosing alternative i , calculated using the distribution corresponding to the estimated distribution y^0 . Let $E(U_i)_{\min}$, in turn, denote the minimum expected utility of choosing alternative i , calculated using a distribution located within Y^0 . An alternative should be evaluated then according to equation (2.7), the decision maker's choosing the alternative that maximises V .

$$V = \rho \cdot E(U_i) + (1 - \rho) \cdot E(U_i)_{\min} \quad (2.7)$$

Hodges and Lehmann's theory

Note that Ellsberg's theory draws heavily upon the work of Hodges and Lehmann [43], who introduce the concept of a "restricted Bayes solution". A restricted Bayes solution is a decision alternative that minimises the expected negative value of an uncertain situation, given that the value of the least favourable outcome is higher than a certain threshold-value. They suggest that a decision alternative should be evaluated according to equation (2.8), where $R_\delta(\theta)$ is the risk function, i.e. distribution of losses, of decision alternative δ , ρ_θ is a number between 0 and 1 indicating the confidence the decision maker has in the probability distribution $\lambda(\theta)$, and $\sup_\theta R_\delta(\theta)$ is the maximum possible loss, given that alternative δ is chosen. Thus, Hodges and Lehmann's evaluation is based on a combination of the expected value and the maximum possible loss of a decision alternative.

$$V = \rho_0 \int R_\delta(\theta) d\lambda(\theta) + (1 - \rho_0) \cdot \sup_\theta R_\delta(\theta) \quad (2.8)$$

Quiggin's theory

Quiggin [44] presents a theory of anticipated utility in which weaker axioms are employed than in Bayesian decision theory. The decision rule used in Quiggin's theory is shown in equation (2.9), where $h_i(p)$ is a weighting function for the probabilities of the different outcomes (o_i) and $U(o_i)$ is a utility function for those outcomes.

$$V = \sum_i h_i(p) \cdot U(o_i) \quad (2.9)$$

Note that Quiggin's evaluation of uncertain decision alternatives resembles that of Prospect Theory in that it employs a utility function which is weighted by a function of the probabilities of the different outcomes. An important difference between Quiggin's weighting function $h_i(p)$ and Prospect Theory's weighting function $\pi(p_i)$, however, is that $\pi(p_i)$ is determined by the probability of the outcome in question, whereas the value of $h_i(p)$ is determined by the probabilities of each of the different outcomes. More precisely, $h_i(p)$ is determined in accordance with equation (2.10).

$$h_i(p) = f\left(\sum_{j=1}^i p_j\right) - f\left(\sum_{j=1}^{i-1} p_j\right) \quad (2.10)$$

The weighting function $h_i(p)$ is thus determined by a function (f) of the cumulative probability of the different outcomes of the decision alternative.

Yaari's theory

Yaari [45] presents a theory of choices between risky alternatives, called the Dual theory of risk, which resembles Bayesian decision theory in that evaluation of the alternatives is preformed using a product measure. In Bayesian decision theory the product of the probability of a particular outcome is multiplied by the utility associated with it in evaluating the alternative in question. In the Dual theory of risk, the value of an outcome is multiplied instead by a function of the probability of the outcome, a function derived by having the decision maker answer questions regarding his/her preferences among various alternatives involving uncertain or certain outcomes. Denote the value of a particular outcome as v_i and the probability of that outcome as p_i . Assign the best outcome a value of 1 and the worst outcome a value of 0. One then derives the function referred to above, termed $f(p)$, using the preference equation presented as equation (2.11). There, $(p; 1)$ is a lottery involving the decision maker's receiving the best outcome with probability p and otherwise nothing. The value of the outcome of the certain alternative

that the decision maker considers equal in value to the lottery, such that he/she is indifferent between the two, is denoted then as $f(p)$.

$$(p; 1) \sim (1; f(p)) \quad (2.11)$$

The process of determining the $f(p)$ function resembles that of determining the utility function with use of Bayesian decision theory. In either case, the decision maker needs to state his/her preferences between alternatives involving simple lotteries and those involving certain outcomes. The difference between the two approaches, however, lies in the fact that in Bayesian decision theory the function (in this case the utility function) is defined over different possible values of the uncertain outcomes, whereas in the Dual theory the function $f(p)$ is defined over the probabilities involved.

In evaluating different alternatives using the Dual theory of risk, one determines the value (V) of a particular decision alternative using equation (2.12). The best decision alternative is that with the highest value. $G(t)$ in equation (2.12) is the probability that the value of the outcome of the alternative is higher than t .

$$V = \int_0^1 f(G(t))dt \quad (2.12)$$

Cumulative Prospect Theory

Cumulative Prospect Theory (CPT) [46] is a further development of Prospect Theory (PT). The difference between the two lies in how decision weights for the probabilities are determined. In CPT a weighting function (w^+) is defined for the probabilities associated with positive outcomes, and another weighting function for the outcomes involving losses (w^-). A negative index is used to denote a negative outcome and a positive index is used to denote a positive outcome. The outcomes can be ordered such that o_{-m} is the worst consequence and o_n is the best consequence. The probability of outcome o_i is denoted p_i . The decision weights, π^+ and π^- , are defined as follows:

$$\pi_{-m}^- = w^-(p_{-m})$$

$$\pi_i^- = w^-(p_{-m} + \dots + p_i) - w^-(p_{-m} + \dots + p_{i-1}) \quad 1-m \leq i \leq 0$$

$$\pi_n^+ = w^+(p_n)$$

$$\pi_i^+ = w^+(p_i + \dots + p_n) - w^+(p_{i+1} + \dots + p_n) \quad 0 \leq i \leq n-1$$

$$w^+(0) = w^-(0) = 0$$

$$w^+(1) = w^-(1) = 1$$

One employs equation (2.13) for evaluating decision alternatives by use of CPT.

$$V = \sum_{i=-m}^0 \pi_i^- \cdot v(o_i) + \sum_{i=0}^n \pi_i^+ \cdot v(o_i) \quad (2.13)$$

The value of a decision alternative (V) is compared then with that of other alternatives, the alternative with the highest value being considered best.

Supersoft decision theory

Supersoft decision theory (SSD) [47] differs from the evaluation techniques described above in that it does not use any one particular evaluation criterion but rather a combination of qualitative criteria and quantitative criteria. It employs three quantitative criteria. All these criteria involve the evaluation of expected utilities.

In evaluating a decision alternative using SSD, one does not need to assign precise values to probabilities and utilities, these parameters instead being assigned as intervals. Thus, the expected utility of a decision alternative cannot be defined as an exact value, alternatives instead being evaluated on the basis of maximum expected utility, minimum expected utility and a measure termed Average. Supersoft

decision theory is described in greater detail in [48] (Paper 3), which is included in the thesis.

The Delta method

The Delta method [49] is similar to the SSD method in that the decision maker does not need to assign precise values to either probabilities or utilities. In using the Delta method, the parameter δ_{ij} is defined as the difference in expected utility between alternatives i and j , or $E(U_i) - E(U_j)$. Obtaining δ_{ij} for the alternatives under consideration results either in a set of admissible alternatives, in there being only one admissible alternative, or in there being no admissible alternative at all. If one ends up without an admissible alternative, this means that none of the alternatives being analysed can be considered as best. If one finds only one admissible alternative, that alternative represents the best alternative. If there are several admissible alternatives, the analysis which concerns the strength of the alternatives, continue in order to determine which alternative is best. The strength of an alternative i as compared with another alternative j is defined as the maximum value of δ_{ij} ($\max(\delta_{ij})$) together with the relative strength Δ_{ij} , defined as $((\max(\delta_{ij}) - \max(\delta_{ji}))/2)$. After having calculated the relative strength of the different alternatives, one can conduct sensitivity analyses of the decision problem in question. This involves reducing the size of the intervals defining the uncertainty regarding the uncertain parameters.

3 The application of decision analysis in fire safety engineering

Several of the theories taken up in chapter 2, as well as various others, will be examined critically here in efforts to determine what theory or theories could be suitable for decision analysis concerned with different fire protection alternatives for use in a factory. At the end of this chapter, a method for analysing investments in fire safety in terms of the theory or theories considered to be best will be presented in detail. As can be seen in Figure 1, the first part of this chapter can be said to deal with selecting a suitable normative model to base a prescriptive model on. The operational model employed, i.e. the model for how the decision rule suggested should be applied in practice, is then taken up. Note that in developing this model, the assumptions in Bayesian decision theory are assumed to hold, or more precisely, that the probabilities and utilities can be estimated exactly.

The final section of the chapter is concerned with how epistemic uncertainty regarding both probabilities and consequences can be dealt with including situations in which the assumptions of Bayesian decision theory do not hold. The development of the prescriptive decision rules is also considered there.

Evaluation of decision analysis methods

To suggest a method for analysing different alternatives for the designing of a fire protection system for a specific building, one needs first to decide what decision theory and what decision criterion to base one's model on. Several decision theories and decision criteria were presented in chapter 2. Here, several of them are analysed to assess how suitable they are for use in the present context, the aim being to select one of them as a basis for the development of a decision analysis method for use here.

A number of different attributes will be employed for judging to what extent a particular decision theory is suitable. Most of these attributes resemble those taken up by Keeney [23] but have been modified somewhat to fit the present context:

- *Whether the decision theory in question is compatible with the approach of quantitative fire risk analysis.*

It is important that the decision theory chosen be compatible with the form of a quantitative fire risk analysis (as described by Frantzych [50], for example). It is also desirable for the theory to be compatible with the measure of fire risk that Hall and Sekizawa [1] have suggested. Here, “compatible” means that the results of the quantitative risk analysis undertaken can be used in the decision analysis that is carried out. This is advantageous since it reduces the work involved in the overall analysis. The logical connections between a quantitative fire risk analysis and the decision analysis to be carried out in conjunction with it, as well as the similarities between them, can also be seen as making it easier to gain acceptance for results of the decision analysis. Decision methods employing fuzzy arithmetic [51], in contrast, can be considered to have a low degree of compatibility with quantitative fire risk analysis since the probability measures the fire risk analysis would provide would be of only limited use in arriving at fuzzy-logic representations.

- *Whether the decision theory in question is adequately established and has a sound theoretical basis.*

This attribute has to do with how much scrutiny the decision theory has been subjected to. A decision theory’s having a sound theoretical basis helps the decision maker answer questions such as “Why should I choose the decision alternative which is best according to this particular decision theory?”. Having a satisfactory answer to such a question is important, since otherwise the decision alternative recommended on the decision theory in question can be viewed as arbitrary.

- *Whether the practical application of the method in the context of fire safety engineering is easy enough.*

A method intended for use in analysing different alternatives for the fire protection of a building needs to be sufficiently easy to use that the efforts required, partly in terms of time, are not greater than what would seem reasonable in light of the results use of the method provides. If a method is too complicated, its practical usefulness is limited.

- *Whether the results the method provides are easily understood and give the decision maker information relevant to the decision in question.*

The results a method yields need to be easy to present and explain to decision makers and need also to provide information directly relevant to the context in question. They should indicate which decision alternative is best, how much the decision is influenced by uncertainties, how additional information for reducing the uncertainty can best be sought, and the like. Moreover, the analysis needs to be transparent in the sense that the decision maker easily can follow the calculations and the assumptions made.

Note that, although it would be possible to employ the Analytic Hierarchy Process (AHP) method [52] for choosing a theory, for example, the choice process employed here is performed instead in a less formal way, each of the theories being analysed and commented on in isolation.

Bayesian decision theory

Employing Bayesian decision theory and the decision criterion of maximising expected utility (MEU) represents a rather attractive approach here, particularly since MEU finds wide application in a variety of contexts. Bayesian decision theory is highly compatible with the use of quantitative risk analysis (QRA). To illustrate this, consider the general form of a decision problem formulated using the Bayesian decision theory. There, an uncertain situation can be described as $(p_1 o_1, \dots, p_n o_n)$, p_i being the probability of outcome i , o_i . The decision rule employed implies that each o_i is associated with a utility number $U(o_i)$ describing the decision maker's preferences among the uncertain outcomes involved. The MEU rule implies that uncertain alternatives should be evaluated in terms of their expected utilities. In considering the general form that QRA results take, one can note a clear similarity between the results of a QRA and the standard form of an uncertain situation to be evaluated using Bayesian decision theory. The results of a fire QRA can generally be described as a set of probabilities, each associated with a consequence $(p_1 c_1, \dots, p_n c_n)$. If the consequences c_i are regarded as being the outcome of an uncertain situation, the results of a QRA and the standard way of describing an uncertain situation by use of

Bayesian decision theory is very similar, making it thus easy to use a QRA in combination with Bayesian decision theory. The only thing one needs to do in order to use the result of a QRA in a decision analysis is to transform the consequences into utility-values.

Bayesian decision theory appears to be the normative decision theory which is best established (see [53], for example). Although it has been criticised substantially, there is still no theory that is obviously superior to it, at least not superior to the MEU rule, which is discussed in [38] (Paper 1).

Due to its similarities to a QRA, Bayesian decision theory would appear relatively easy to employ in a fire protection context. There could nevertheless be difficulties in estimating the probabilities of different fire scenarios, since Bayesian decision theory requires these estimates to be precise. This could be difficult to achieve in the present context due to there being only limited information regarding the occurrence of various fire scenarios. Although one can express the uncertainty regarding a given probability value by use of a probability distribution, which would make a probability estimate possible even when there is considerable uncertainty regarding it, Bayesian decision theory states that whatever uncertainty there is concerning the different probabilities and consequences does not affect the decision alternative deemed best, the only thing of importance being the expected utilities of the alternatives. This can be regarded as negative in the present context, since the uncertainties regarding the probabilities and the consequences of the different fire scenarios are likely to be substantial. The decision maker would undoubtedly want to have an idea of whether epistemic uncertainty of this sort can affect which decision alternative is best (see the discussion on robust decisions in the section termed “Dealing with epistemic uncertainty – prescriptive decision rules”, and references [38] and [48] (Paper 1 and 3)).

The expected utilities of the different alternatives represent part of the results obtained in the application of Bayesian decision theory. Although the expected utilities facilitate comparison of the alternatives, the real benefit of employing Bayesian decision theory is that this enables the results of a decision analysis to be expressed in monetary terms. The concept of Certainty equivalent (CE) is used in connection with this. The CE of an uncertain decision alternative is the monetary amount the decision maker considers to represent the value of the decision

alternative in question. Thus, the theory not only provides the decision maker a conception of which decision alternative is best, but also provides an evaluation of how good/bad a given decision alternative is in terms of some monetary value. From a pragmatic standpoint, the use of monetary values for the evaluation of decision alternatives is very appealing, since decision makers can be expected to be familiar with the monetary scale, whereas other scales such as the utility scale can be experienced as being more complex.

Due to its compatibility with a quantitative risk analysis, its firm theoretical foundation, and the simple form in which the results of using it can be presented, the decision rule of maximising expected utility is a very strong candidate for use in the present context. However, as has been pointed out, some form of complimentary evaluation may be needed for the MEU rule to be used in the present context.

Alternatives theories

Many of the theories taken up in chapter 2 can be considered to be extensions of Bayesian decision theory. None of them offer any obvious advantages as compared with the MEU criterion that would justify employing that criterion, in the present context, instead of the MEU criterion. However, since some aspects of these theories are appealing, especially those of Prospect theory, Ellsbergs' method, the SSD method and the Delta method, they will be used to provide certain help in designing the prescriptive method for the evaluation of different fire protection alternatives to be employed here. An evaluation of various decision criteria of possible use is presented below.

Prospect theory [40] adds a weighting function for the probabilities, one which the MEU criterion does not employ. In a fire risk model this means that, in addition to estimating the different probabilities of the different fire scenarios, one should estimate a function that can be used to adjust each of the probabilities. The weights involved cannot be interpreted as probabilities since they are not required to sum to unity. Since the weights introduced, although based on probabilities, do not sum to unity, Prospect Theory does not satisfy stochastic dominance (see for example [46], page 299). In comparing two alternatives a_1 and a_2 , alternative a_1 stochastically dominates a_2 if $P(X_1 \geq x_i) \geq P(X_2 \geq x_i)$ for all x_i , where $P(X_i \geq x_i)$ is the probability of a consequence x_i or better, given

that one chooses alternative a_1 . Similarly, $P(X_2 \geq x_i)$ is the probability of a consequence x_i or better, given that one chooses alternative a_2 . Since satisfying stochastic dominance is a desirable feature for a normative/prescriptive decision theory to have, it is regarded as negative in terms of Prospect Theory.

The reason for Prospect theory being developed was to *describe* a certain type of behaviour that could not be explained by the MEU criterion. Thus, the usefulness of Prospect theory in the present context, in which *prescriptive* or *normative* guidance is needed, can be questioned. Prospect theory nevertheless provides an idea that seems reasonable from a practical standpoint. According to this theory, uncertain situations involving losses or gains that are “certain” are evaluated on the basis of the value of the certain part *plus* the value of the uncertain part (see equation (2.5)). In analysing alternative investments in fire safety, this would imply that one first evaluates the costs and benefits that are largely certain (investment costs and maintenance costs, for example) and then add an evaluation of the uncertain consequences due to possible fires. The line of reasoning here is that analysis of a fire protection alternative is clearer and easier if evaluation of the costs and benefits that are certain is separated from the remainder of the analysis and are not mixed in with the evaluation of the fire risk. However, the possible advantages of employing the decision criterion of Prospect theory instead of the MEU criterion does not appear to be sufficient to compensate for the fact that Prospect theory does not satisfy first degree stochastic dominance.

The decision rule suggested by Bell [41] implies, alongside use of the MEU criterion, that in evaluating the utility of a particular outcome one also take account of the regret the decision maker would feel of having precluded the possibility of having achieved some desirable outcome (having missed it) through having chosen an alternative to which this outcome did not belong. The “regret” aspect included in this theory would probably have only a marginal effect in the present context, however, since most of the decision alternatives to be analysed here are of the type for which one cannot *know* for certain that the outcome would have been different if the decision maker had chosen some other alternative. Assume, for example, that a decision maker is to decide between the alternative of investing in a sprinkler system and of keeping the building in its present form. In evaluating the utility of a major fire,

given that one choose not to invest in the sprinkler system, one cannot *know* that the sprinkler investment would have made any difference in the outcome, even if one may *suspect* that it would have. This illustrates the difference between the type of situations analysed by Bell's decision rule, in which one can determine *for sure* that a particular outcome would be precluded by the choice of a given alternative and the quite different situation involved in deciding between different fire protection alternatives. Due to this difference, Bell's decision criterion is thus not regarded as offering any significant advantage compared to the MEU criterion.

Although Ellsbergs' decision rule (equation (2.7)) appears appealing since it involves the uncertainty regarding the probability values in the model being recognised explicitly, in a fire protection context one would probably find it very difficult to determine the ρ -value, which is the probability that the estimated distribution defined on the outcomes is correct. Also, from a prescriptive standpoint, why should the decision rule only employ the minimum value of the expected utility rather than some other plausible value of it? In addition, Ellsberg provides no axiomatisation of his theory, instead using the MEU rule as a point of departure and empirical observations as a basis for changing it, admittedly in a way that seems reasonable, given the observed behaviour of people generally. Since in the present context interest is directed at developing a prescriptive model for fire safety, descriptive considerations such as those underlying the development of Ellsbergs' model can be seen as less important than if one wanted to develop a model to *describe* how people actually make decisions regarding fire safety. However, Ellsberg's idea of considering a class Y^0 of plausible probability distributions is of interest. In making estimates of the probabilities of different fire scenarios, especially catastrophic fire scenarios, one is likely to be very uncertain. This uncertainty can be interpreted in the same way as Ellsberg does, namely by assuming there to be a class of plausible probability distributions defined on the various possible outcomes of a fire. Since each of these plausible distributions is associated with a particular expected utility value, there is a set of plausible expected utility values for each fire protection alternative. In the suggestion of a method for decision analysis of different fire protection alternatives that will follow shortly, use will be made of Ellsbergs' idea of a set of plausible probability distributions.

Hodges and Lehmann's [43] suggested approach to evaluating a decision alternative does not appear to offer any advantages in the present context compared with an ordinary approach to evaluating the expected loss. This is because for many of the decision alternatives likely to be involved in the present context the maximum possible loss is the same, representing complete destruction of the building in question. Since Hodges and Lehmann's evaluation criteria combine an evaluation of the maximum possible loss and of the expected loss, only the expected loss can be used to distinguish the various alternatives. Accordingly, Hodges and Lehmann's evaluation of decision alternatives is not regarded as useful in the present context.

Quiggin's theory [44], Yaari's Dual theory of risk [45] and Cumulative Prospect Theory [46] propose the use of a weighting function for probabilities. Also, the weighting function is defined over the *cumulative* probabilities as opposed to the probabilities of the different outcomes, which is the case in Prospect theory. The authors referred to above provide axiomatisations of the decision rule used in their theories, and their theories could probably be used in conjunction with a quantitative risk analysis. Adding a weighting function to the probabilities, however, but not to the consequences would presumably be difficult in practice, especially in view of the difficulties this would create in communicating use of the method to decision makers within the companies involved. This is because introducing a weighting function for the probabilities does not appear as easy to understand as introducing a weighting function for the consequences different outcomes would have. Moreover, there is no measure of fire risk (known to the author, at least) in which probabilities are adjusted by use of a weighting function, whereas "utility-based" measures of consequences have been suggested [1]. Also, in comparing Quiggin's and Yaari's theories by use of common ratio tests³ Malmnäs [54] found the application of each to perform in much the same way as an expected utility evaluation. Although the arguments

³ Common ratio tests are used to compare the performance of different decision rules. The tests involve evaluating two uncertain decisions alternatives, each having two possible outcomes. For each of the two alternatives, the probability of the best outcome is multiplied by a factor r , given the same value in the case of both alternatives, a value which is usually very small. By examining for various r -values which alternative a given decision rule recommends, one has a basis for comparing the performance of that decision rule with another.

presented above against the use of Quiggin's theory, Yaari's theory or CPT are not particularly strong, neither of the methods appears to possess any features making them obviously superior to Bayesian decision theory or the use of the MEU principle. Thus, the practical problems that can be anticipated in using a weighting function for the probabilities as well as the fact that they behave in much the same way as an expected utility evaluation in common ration tests appear sufficient to justify choosing the expected utility criterion rather than Quiggin's theory, the Dual theory of risk or Cumulative prospect theory as a basis for the prescriptive rule to be employed here.

Both the methods employed in Supersoft decision theory (SSD) [47] and the Delta method [49] use a somewhat different approach than the other methods taken up here. Instead of employing exact probabilities, both methods utilise an imprecise representation of both probabilities and utilities, i.e. probabilities and utilities are not assigned as precise values (or probability distributions) but as intervals. This complicates the use of these methods in conjunction with a traditional fire risk analysis since the values used there are generally assigned as being precise or as being probability distributions. However, the problem does not appear insurmountable, since sensitivity analyses in which uncertain parameters are adjusted from their lowest and highest values are often included in quantitative risk analysis. The lowest and highest values used in a sensitivity analysis could be used in an analysis employing either the SSD method or the Delta method.

The degree of usefulness of these two methods in the present context can be seen as high. In using methods of these two types, however, complicated calculations may be called for (depending on the problem at hand), reducing the methods practical usefulness. It thus remains to be seen whether the computer programs that have been developed and are being developed [49] and [55] can facilitate use of these methods. Nevertheless, some of the ideas the methods involve appear very interesting and are taken into account in designing the method recommended here for analysing different fire protection alternatives. Note that the usefulness of these methods in the present context is due to their not requiring exact representation of probabilities and utilities and that they can thus be used in contexts in which the assumptions employed in Bayesian decision theory (and in many of the other theories referred to above) do not hold.

Although there are simpler methods not taken up in the thesis earlier, which could be used in the present context, many of them are not suitable here. These include the so-called Laplace paradigm [56], the Wald paradigm [57] (also called the minimax or maximin decision rule), the Savage paradigm [58] and the Hurwicz paradigm [59]. In showing why these methods are not suitable, use will be made of a hypothetical decision situation. Assume that a decision maker has to decide between two fire protection alternatives, one involving an investment in some kind of fire safety system and the other involving the building being kept in its original state. Denote the cost of the investment as C_{inv} and represent the possible outcomes of a fire (expressed in monetary terms) by the uncertain situation $L = (p_1x_1, p_2x_2, \dots, p_nx_n)$, such that consequence x_i occurs with probability p_i , etc. Assume also that $x_1 > x_2 > \dots > x_n$. It is assumed that in using these methods one is completely ignorant concerning the values of the probabilities p_i .

Employing the Laplace paradigm involves one's choosing the decision alternative that minimises the expected costs, the probabilities of the different outcomes being treated as being equal. Since an investment in a fire safety system (consisting, for example, of a water sprinkler system, smoke alarm, smoke ventilation, etc.) does not generally reduce the maximum *possible* consequence but reduces instead the *probability* of such a consequence, all outcomes involved, ranging from those of total destruction of the building to those of no damage at all occurring are possible even if the investment in question should be made. Accordingly, the Laplace paradigm would always result in the best alternative being to keep the building in its original state. To see this, denote the uncertain outcome of not having invested in the fire safety system as $L_1 = (p_{1,1}x_1, p_{2,1}x_2, \dots, p_{n,1}x_n)$ and the uncertain outcome involved in having made the investment as $L_2 = (p_{1,2}(x_1 - C_{inv}), p_{2,2}(x_2 - C_{inv}), \dots, p_{n,2}(x_n - C_{inv}))$. These two uncertain situations differ in the probabilities of the outcomes and in the investment cost (C_{inv}) added to each of the possible outcomes in the case of the alternative of investing in fire safety. The Laplace evaluation ($V_{La}(L)$) of these two uncertain situations will always yield a higher value for the alternative that involves no investment being made. This is because:

$$V_{La}(L_1) = \sum_{i=1}^n \frac{1}{n} \cdot x_i > \sum_{i=1}^n \frac{1}{n} \cdot (x_i - C_{inv}) = V_{La}(L_2)$$

This demonstrates clearly that the Laplace paradigm is unsuitable for use in the present context. The same type of argument can be directed against use of the Wald paradigm, according to which the decision maker should choose the alternative for which the costs for the worst outcome are lowest. Recalling the example given above, it is clear that for both alternatives the worst consequence is complete destruction of the building. When the investment alternative is chosen the cost of this consequence will be the cost of consequence x_n *plus* the cost of the investment, C_{inv} , but when no investment is made it is only the cost of consequence x_n . Thus, according to the Wald paradigm the best alternative would always be to recommend not investing in a fire safety system. Similar arguments can be directed against use of the Savage paradigm and the Hurwicz paradigm. The problem of using these methods is that they ignore the probabilities of the different outcomes and thus ignores as well the positive effects of the most common investments in fire safety, which generally aim at reducing the probability of a serious fire.

Conclusions from evaluating different decision analysis methods

A conclusion that can be drawn from the evaluation of the different decision analysis methods just presented is that there apparently is no “ultimate” decision analysis method. Since all the methods considered can be criticised in one way or another, one needs to evaluate the methods here on the basis of how well they would function in the present context. Thus, Prospect theory, for example, fails to satisfy stochastic dominance and can therefore be seen as having a weaker normative foundation than the MEU principle. Bell’s theory, in turn, appears to add nothing of significance to what the MEU criterion provides, since the usefulness of “regret” in the present context can be regarded as very limited. Ellsberg’s decision criterion appears reasonable from a practical perspective but the normative basis of it appears weaker than that for the MEU criterion. Hodge’s and Lehmann’s approach appear to be very similar to that of the MEU criterion when employed in the present context and therefore it does not provide any obvious advantage as compared to using the MEU criterion. Quiggin’s theory, the dual theory of risk (Yaari’s theory) and Cumulative prospect theory seem reasonable and appear about as useful as the MEU rule. However, since these methods introduce weighting functions for the probabilities involved, they appear to be more complicated to use in conjunction with fire risk

analysis than the MEU principle is. Accordingly, they are not employed here. Both Supersoft decision theory and the Delta method appear to possess advantages in evaluating decision situations of the type encountered here. Since the basic ideas the methods involve can easily be used in combination with methods developed on the basis of Bayesian decision theory, Supersoft decision theory will be incorporated into the theoretical and practical framework to be suggested. In contrast, Laplace's method and Wald's method, which are of a simpler type, can lead to rather strange recommendations in the present context. Thus, they will not be considered here further. The principle of maximising expected utility (MEU) thus appears to be a strong candidate for use as a decision rule here. The theory has many advocates, such as Winkler [60], for example, who states:

“Is there a theory that is more appealing [, than utility theory,] from a *normative* viewpoint? ...Recent efforts to develop new axiomatic theories with normative orientations are exciting and stimulating, but are any of them convincing enough to cause us to shift loyalties? That's a matter of personal opinion, of course: for me, the answer at the moment is no.” ([60] p. 247)

Malmnäs [61] has criticised Bayesian decision theory, showing that a decision maker who has accepted the axioms of the theory does not necessarily make decisions in accordance with the MEU principle. Malmnäs also concludes, however, that in comparing different evaluation criteria it is difficult to find a better criterion than the MEU criterion⁴:

“...the prospects for finding an evaluation [decision rule] that is much better than $E(A,f)$ [MEU] are not particularly bright.” [47]

Examining the possible extensions of the MEU rule that have been presented here one can conclude that no single criterion appears to be substantially better than the MEU criterion for use in the present context. Accordingly, the MEU criterion will be used as a point of departure for

⁴ For practical decision situations, Malmnäs [54] has suggested combining qualitative evaluations with use of expected utilities.

the method to be suggested for use in the present context. The criterion will be extended, however, in efforts to deal with some of the (possible) shortcomings it has when applied to problems of fire safety. What will be suggested here is also not *a* method for the evaluation of different fire protection alternatives, but rather a prescriptive framework within which decision analysis can be performed. The difference between this framework and a specific decision criterion is that it includes several possible ways of evaluating alternatives, its being up to the individual decision maker to choose some one of these for use with a particular problem.

Suggestion of a primary model for decision analysis

Since Bayesian decision theory was found to be the most suitable theory to base a decision model on in the present context, that theory will be used as a point of departure in constructing a model for the analysis of different investments in fire safety. Note that two of the particularly important assumptions in Bayesian decision theory is: (1) that each of the various outcomes possible in the decision situation can be assigned a unique utility value, and (2) that the decision maker can assign a specific probability value to each of these outcomes. The model created for calculating the expected utility of investing in a specific fire safety alternative (using assumptions (1) and (2)) is called the primary model. That model can be used in combination with various prescriptive decision rules. In the thesis, three such rules are provided, its being up to the decision maker to determine which rule is most applicable to his/her individual case. The decision rules considered are the principle of maximising expected utility, made use of by Bayesian decision theory; an extension of this rule, termed extended decision analysis; and a decision rule based on Supersoft decision theory (SSD). Although the three rules differ in their evaluation of decision alternatives, they all employ some kind of evaluation of expected utilities, allowing the primary model to be used in conjunction with each.

Note that it is assumed that all fire safety alternatives being evaluated by the methods suggested here comply with the building codes that apply and that thus no evaluation of occupant safety is included here.

So as not to be overwhelmed by the task of constructing the primary model, one can start by evaluating smaller segments of the problem and

then put these together to form a complete model. It is reasonable to start by endeavouring to calculate the expected utility of *a single* fire.

Assume that in a hypothetical decision situation a decision maker has to choose between suffering the consequences of a fire in the one or in the other of two separate buildings. Assume for simplicity too that the decision maker owns both buildings. How could the decision maker calculate the expected utility of each of the two alternatives here? First, one would need to know what the outcomes of the decision are. These can be described in terms of the fire scenario that occurs if a fire should occur in either of the two buildings. Note that the term “scenario” is used to refer to an element of the fire situation that involves “...a complete physical description of the fire; the environment in which it began, developed, and ended; and the consequences of its occurrence.” [1]. Assuming that one has identified a finite number of possible fire scenarios, one needs to be able to estimate for each of them the probability of its occurring. A quantitative fire risk analysis can be used for doing this. In addition to having estimated the probabilities of the different outcomes of a decision, one needs to be able to assign a utility value to each of the outcomes such that the utility values represent how good or bad the different uncertain outcomes are in relation to each other. If one can do this, one can describe the occurrence of a fire in a specific building using the “lottery”-form introduced earlier: $L = (p_1 o_1, p_2 o_2, \dots, p_n o_n)$, where n is the number of fire scenarios considered in the analysis. In comparing the decision alternatives, one should calculate the expected utility, according to equation (3.1), for each of the alternatives and then choose the alternative having the highest expected utility.

$$E(U) = \sum_{i=1}^n p_i \cdot U(o_i) \quad (3.1)$$

In determining the utility of a specific outcome the decision maker should select a value $U(o_i)$ such that he/she is indifferent between the outcome o_i and a hypothetical lottery in which he/she will receive the best outcome, o_1 , with probability $p = U(o_i)$ and the worst outcome, o_n , with probability $(1-p)$, its being assumed that $o_1 > \dots > o_i > \dots > o_n$ (see chapter 2). Assigning utilities in this way is clearly impractical in the present context, however, since a fire risk analysis may easily consist of hundreds of fire scenarios, each requiring the decision maker to make judgements regarding hypothetical lotteries. Instead, it is suggested here

that all fire scenarios (outcomes) be evaluated in terms of their *intrinsic* monetary value. This intrinsic monetary value can be translated into a utility value using a utility function that can be constructed requiring of the decision maker only a limited number of judgements regarding uncertain situations.

Note that the intrinsic monetary value of a particular fire scenario is the monetary value that the decision maker regards as being equal to the value of the fire scenario in question, his/her being indifferent between suffering the consequences of the fire scenario and paying the monetary amount. In finding the intrinsic value of a particular fire scenario, the decision maker needs to evaluate the *uncompensated losses* the fire scenario involves. The uncompensated losses can, for example, be lost market shares due to the business interruption, fines, bad reputation, and the like (see [38] (Paper 1), for a discussion of uncompensated losses). In evaluating fire scenarios in terms of their intrinsic monetary value, one assumes that all aspects relevant to the decision maker can be captured in terms of their monetary value. In the present context, involving choices between different fire protection alternatives for factories, it is likely that such is the case. On the other hand, if a decision maker feels that some aspects of the consequences cannot be expressed in terms of monetary value, some other method of evaluating them can be employed. The method referred to is that of Multi-objective utility theory, which deals with the evaluation of consequences involving several non-commensurable objectives (see [62], for example). The uncompensated losses associated with a particular outcome of a fire o_i , is termed x_i .

Having expressed each fire scenario in terms of its intrinsic monetary value allows the decision maker to continue then with the utility evaluation of the scenarios. The whole point of the monetary evaluation of the fire scenarios is to make it easier to assign each of them a utility value, fewer judgements regarding uncertain hypothetical situations thus being required. The idea is to construct a utility function on the basis of a relatively limited set of questions that are posed to the decision maker. These can be used to “translate” each of the uncertain monetary values involved into a utility value. The construction of a utility function is taken up in most introductory books on decision analysis (see [36], for example). A common way of constructing one is to ask the decision maker a series of question regarding his/her preferences towards different lotteries involving monetary outcomes. The decision maker’s answers can be used to deduce the utility function. Note that there are

many methods for assessing a utility function. A comprehensive survey of many such methods is provided in [63].

The form of the utility function indicates whether the decision maker is risk averse, risk neutral or a risk seeker. A convex utility function (a utility function having a positive second derivative) is called a risk-seeking utility function, a concave utility function (a utility function having a negative second derivative) a risk-averse utility function, and a utility function that is linear is called a risk neutral utility function. In the thesis, not much attention is directed at the elicitation of utility functions since this is covered thoroughly elsewhere (see the reference referred to above, for example, the methods presented there being applicable here).

The process of calculating the expected utility of a single fire can be summarised as follows: (1) create a model for the development of a fire such that the probability of each of the different fire scenarios can be estimated (a suggestion for such a model being given later in the thesis), (2) determine the uncompensated losses for each fire scenario and evaluate them in terms of their intrinsic monetary value, (3) assess a utility function defined on monetary outcomes, (4) use the utility function to calculate a utility value for each of the fire scenarios, and (5) calculate the expected utility of a single fire using equation (3.1).

In practice, being able to calculate the expected utility of *one single* fire is not very useful. Instead, what is of interest is the expected utility of choosing a specific fire protection design in a building. Thus, one needs to evaluate the possibility of having more than one fire in the building of interest as well as evaluating the more or less certain costs and benefits associated with a particular fire safety design, such as investments costs, maintenance costs, and the like.

Calculating the expected utility of a fire exposure

Assume one is interested in evaluating a particular fire safety design during a particular time period, one of five years for example. The possibility of having an unknown number of fires, each with an unknown outcome during a particular period of time, is called a *fire exposure*. To evaluate one fire exposure and compare it with another, one can calculate the expected utility of each of the two exposures and compare their values.

Calculating the expected utility of a fire exposure is more complicated than calculating the expected utility of a single fire. In the case of a fire exposure one has to consider the decision maker's preferences both for the event of more than one fire occurring and for fires that occur at different times. Expressing the expected utility of a fire exposure in as simple a way as one can express the expected utility of a single fire (equation (3.1)) requires that a number of rather strict assumptions be made concerning the decision maker's utility function for more than one fire. Since one cannot know in advance what assumptions are valid for a given decision maker, the approach here is to start by deriving a utility function using a set of assumptions that may or may not be applicable to the decision maker in question (the applicability of which needs to be determined on a case-by-case basis, i.e. from one decision maker to the next). The aim of using such assumptions generally is to find a form of the utility function that facilitates insofar as possible a practical analysis of the problem. The utility function should both seem reasonable to the decision maker and involve a work effort no greater than seems motivated by the results it provides. If the decision maker cannot accept the assumptions one has made in deriving the utility function here, one needs to investigate the decision maker's preferences in detail and derive some other form of utility function the decision maker can accept.

Emphasis is thus placed on the method suggested to the decision maker for calculating the expected utilities being practical, applicable and acceptable to the decision maker. At the end of this section, however, a cruder method for calculating the expected utility of a fire exposure is also considered. This method is based, not on detailed assumptions regarding the decision maker's preferences for fires but on more intuitive ideas concerning the utility of a fire exposure generally. The advantage of this method is that it may be easier to use in practice, at the same time as its drawback is that determining whether it is appropriate for the decision maker at hand can be more difficult.

As indicated above, evaluating the expected utility of a fire exposure is more complicated than evaluating the expected utility of a single fire, its requiring a number of assumptions regarding the decision maker's preferences for more than one fire. Here, various assumptions of this sort will be investigated and it will be shown how they can be used to calculate the utility of a set of fires. It is up to the individual decision maker to determine which assumptions he/she feels is justified for the problem at hand. Nevertheless, there are considerable practical problems

in assessing a utility function for more than one fire if one makes use of only the most relaxed assumptions. The major aim here is to consider a number of assumptions that in practice allow one to estimate a utility function more easily.

It is assumed that the decision maker can evaluate all possible fire scenarios (outcomes) in terms of their intrinsic monetary value, which is the value making him/her indifferent between occurrence of the fire scenario in question, o_i , and losing the monetary amount x_i . Thus, it is assumed that $U(o_i) = U(x_i)$. Hereafter, the outcome of a fire will be expressed in terms of its intrinsic monetary value x .

Thus far in the thesis, little has been said about multi-attribute utility theory. This is because in practice it is probably easier to express the losses a fire has caused in terms of a single monetary attribute and to translate this into a utility value than to assess a separate utility function for each attribute of interest, lost market shares due to business interruption, bad reputation, and the like, and to then use some way of weighting these utilities in the overall evaluation. Nevertheless, when considering fires occurring at different times one does need to take different attributes into account separately. One can view the outcome (fire scenario) of the first fire as being one attribute, the outcome of the second fire as being another, and so on. Let $U(x_1, x_2)$ be the utility of 2 fires occurring in a specific building, where x_1 and x_2 are the intrinsic monetary values of the outcomes of fires 1 and 2, respectively. Denote the worst outcome of the first fire as x_1^0 , the worst outcome of the second fire as x_2^0 , the best outcome of the first fire as x_1^* and the best outcome of the second fire as x_2^* . Assume that x_1' is that outcome of the first fire which would make the decision maker indifferent between the two alternatives shown in Figure 5. The alternative to the left involves outcome x_1' occurring in the case of the first fire and x_2^0 in the case of the second. The alternative to the right involves a probability of 0.5 of the outcome x_1^i occurring in case of the first fire and the outcome x_2^0 in the case of the second, and a probability of 0.5 of the outcome x_1^j occurring in the case of the first fire and the outcome x_2^0 in the case of the second fire.




Figure 5 Two decision alternatives that the decision maker is indifferent between.

If the outcome x_2^0 for the second fire can be changed to some other outcome (which is the same for both alternatives) and this does not affect that outcome of the first fire which would make the decision maker indifferent between the two alternatives (the outcome x'_1), then the outcome of the first fire and the outcome of the second fire are utility independent [62]. In the present context, utility independence between fires would appear to be a reasonable assumption.

If two attributes, x_1 and x_2 , are utility independent, $U(x_1, x_2)$ can be expressed by the multilinear expression contained in equation (3.2) [62].

$$U(x_1, x_2) = U(x_1, x_2^0) + U(x_1^0, x_2) + k \cdot U(x_1, x_2^0) \cdot U(x_1^0, x_2) \quad (3.2)$$

$$U(x_1^0, x_2^0) = 0$$

$$U(x_1^*, x_2^*) = 1$$

$$k = \frac{1 - U(x_1^*, x_2^0) - U(x_1^0, x_2^*)}{U(x_1^*, x_2^0) \cdot U(x_1^0, x_2^*)}$$

The utility function in equation (3.2) can be extended to n attributes [62], making it possible to estimate the utility of n fires. In practice, however, it is considered very difficult to use such a utility function due to its complex form when n is large (see [62] page 293). Instead, one can suggest that additional assumptions be employed.

If the decision maker is indifferent between the two uncertain decision alternatives shown in Figure 6, for all (x_1, x_2) and for an arbitrary x_1^i and

x_2^j , then the utility function $U(x_1, x_2)$ can be expressed as the additive utility function [62] shown in equation (3.3).




Figure 6 Two uncertain decision alternatives.

$$U(x_l, x_2) = U(x_l, x_2^0) + U(x_1^0, x_2) \quad (3.3)$$

$$U(x_1^0, x_2^0) = 0$$

$$U(x_1^*, x_2^*) = 1$$

Although the assumption of additive utility has not been investigated in the context of fire protection engineering, it appears less likely to be applicable than the mutual utility independence assumption. Assume that x_l in the figure is the best outcome of the first fire, that x_2 is the best outcome of the second fire, that x_l^i is the worst outcome of the first fire, and that x_2^j is the worst outcome of the second fire. If the decision maker chooses the decision alternative to the left in Figure 6, he/she will either suffer two fires for each of which the outcome occurring is the most serious one or two fires for each of which the outcome occurring is the least serious one, whereas if he/she chooses the decision alternative to the right in the figure, he/she knows that the result of the two fires will in both cases be such that in the one fire it is the most serious outcome and in the other fire it is the least serious outcome that will occur. The additive utility function in equation (3.3) can be expressed as equation (3.4) [62], where $k_1 = U(x_1^*, x_2^0)$ and $k_2 = U(x_1^0, x_2^*)$. Note that $k_1 + k_2 = 1$, that $U_1(x_l)$ is the conditional utility function for the outcome of the first fire, and that $U_2(x_2)$ is the conditional utility function for the outcome of the second fire.

$$U(x_1, x_2) = k_1 U_1(x_1) + k_2 U_2(x_2) \quad (3.4)$$

$$\begin{aligned} U_1(x_1^*) &= 1 \\ U_1(x_1^0) &= 0 \\ U_2(x_2^*) &= 1 \\ U_2(x_2^0) &= 0 \end{aligned}$$

Note that equations (3.3) and (3.4) can be extended to include n attributes [62] and can thus be used to calculate the utility of several fires. A utility function for n fires using the assumptions referred to above is shown in equation (3.5), where $U_i(x_i)$ is the conditional utility function for the outcome of the i th fire. Note that $\sum_{i=1}^n k_i = 1$.

$$U(x_1, x_2, \dots, x_n) = \sum_{i=1}^n k_i \cdot U_i(x_i) \quad (3.5)$$

Calculating the *expected* utility of a series of fires while using the independence assumption and the additive utility assumption, as well as assuming the fires to be probabilistically independent, involves expressing the expected utility of the fires as the sum of the expected utilities of the individual fires ($E(U_i(x_i))$) (see [62] page 242) multiplied by their respective scaling constants k_i (see equation (3.6)).

$$E(U(x_1, x_2, \dots, x_n)) = \sum_{i=1}^n k_i \cdot E(U_i(x_i)) \quad (3.6)$$

If one assumes that the fires occur in a short period of time, during a single year for example, it is reasonable to assume that the k -factors and the conditional utility functions for each of the fires will be the same. This means the utility of the consequences the various fires can bring about not being affected by the order in which the different fires occur.

Since the sum of the constants k_i is 1 and the conditional expected utility of a fire is the same, regardless of whether it is the first fire or some other fire, the expected utility of n fires is $E(U(x))$, which is the expected utility of a single fire.

Thus, determining the expected utility of n fires occurring in a short time interval is not very complicated if the assumptions referred to above are employed. Here, however, the interest is in determining the expected utility of a *fire exposure*, which involves the possibility of virtually any number of fires occurring during a particular time period.

To provide a better understanding of the results of a decision analysis using the methods described here, it is suggested that the attractiveness/unattractiveness of a particular fire exposure be presented in terms not of its expected utility but of its certainty equivalent (CE). The CE is the monetary amount, for certain, that can be seen by the decision maker as representing the utility of a particular uncertain situation. The decision maker should thus be indifferent between the alternative involving the uncertain situation and the monetary amount which the CE of the uncertain situation represents. To calculate the CE, one needs to calculate the expected utility of the fire exposure in question and to then find what present monetary sum the decision maker would be indifferent between having to pay as opposed to being faced with the possible consequences of the fire exposure.

It was concluded that the expected utility of n fires occurring in a short period of time is equal to the expected utility of a single fire and that it is thus not dependent on the number n as such. Therefore, the expected utility of a fire exposure during a short period of time is also equal to the expected utility of a single fire. The utility of the certainty equivalent is calculated by assuming the consequences of the first fire to be equal to the certainty equivalent of the fire exposure in question and the consequences of the rest of the fires that might in principle occur during this period to be equal to the best consequence possible (i.e. that none of these additional fires occurring). If the expected utility of the fire exposure in question and the utility of the certainty equivalent are set to being equal, one can solve the equation (in this case (3.7)) for the value of the certainty equivalent. $E(U_E)$ is the expected utility of the fire exposure, which is equal to the expected utility of a single fire ($E(U(x))$), and k is the same for all the fires. $U(CE)$ is the utility of a consequence equal in monetary value to that of the certainty equivalent, and $(n-1) \cdot k$ is the utility of the best possible outcome of the remainder of the n fires. Note that the right-hand side of equation (3.7) is constructed using equation (3.5).

$$E(U_E) = E(U(x)) = k \cdot U(CE) + (n-1) \cdot k \quad (3.7)$$

Rearranging equation (3.7), noting that $k = 1/n$, results in equation (3.8).

$$U(CE) = 1 - (n \cdot (1 - E(U(x)))) \quad (3.8)$$

This equation provides an expression allowing the certainty equivalent of n fires to be calculated. Substituting λ , the number of fires *expected* to occur during the period of time in question, for the number of fires n provides a way of calculating the certainty equivalent of a brief fire exposure shown in equation (3.9).

$$CE = U^{-1}(1 - (\lambda \cdot (1 - E(U(x)))) \quad (3.9)$$

The equations above can be used for short time intervals for which one can assume that the k -factors in equation (3.6) are all equal and where the conditional utility functions for the fires are the same. A decision maker concerned with a longer time interval, however, is likely to consider the time factor to be important. The reason for this can be seen as being the same as in the area of traditional capital investment analysis, where a cost that occurs ten years from now is better than the same cost occurring today. In a fire-safety context, this would imply that the consequences due to a fire occurring ten years from now is regarded as being less severe than those if the fire should occur today.

To continue the development of a utility function enabling the certainty equivalent of a fire exposure extending over a long period of time to be calculated it will be assumed that the decision maker evaluates monetary outcomes at different times in accordance with the principle of discounting often used in capital investment appraisals. There, costs or benefits occurring in the future are discounted to the present in such a way that the discounted value represents the value at present that is equal to the value in question at some specific time in the future. Since the severity of the outcomes of fires is measured in monetary terms, discounting seems to be a reasonable way of determining a decision maker's preferences for fires that occur at different times. Thus, the decision maker is assumed to be indifferent between a fire that occurs during the j th year involving consequences that are deemed equal to the monetary value x_j and a fire that occurs at present involving consequences that are judged to be equal to the monetary value $x_{0,j}$. One can calculate $x_{0,j}$ using equation (3.10), in which r (a discount rate) determines the decision maker's strength of preferences for monetary

outcomes at different times and j is the year in which the fire in question occurs.

$$x_{0,j} = \frac{x_j}{(1+r)^j} \quad (3.10)$$

The assumption concerning the discounting of future monetary values allows one to calculate the expected utility of n fires, each of them occurring in a different year, by use of a modified version of equation (3.6). In this version, all the fires are assumed to occur at present (an assumption the discounting of future monetary values permits). Accordingly, one only needs to employ a single conditional utility function and a single k -value. The resulting equation (3.11) employs a conditional utility function termed $U_0(x_{0,j})$, since it applies to monetary outcomes occurring at present (at time 0), and a k -value termed k_0 . $E(U_0(x_{0,j}))$ is the expected utility of a fire occurring during year j , the monetary consequences for that year having been discounted to the present by use of equation (3.10).

$$E(U(x_1, x_2, \dots, x_n)) = \sum_{j=1}^n k_0 \cdot E(U_0(x_{0,j})) \quad (3.11)$$

This equation (3.11) can be modified so as to remove the restriction stating that only one fire can occur during a given year. This results in equation (3.12), in which $x_{j,i}$ is the monetary equivalent of the i th fire occurring during the j th year, and n_j is the number of fires occurring during year j . Note that $\sum_{j=0}^{n_y} n_j \cdot k_0 = 1$ and that n_y is the number of years considered in the analysis.

$$E(U(x_{1,1}, \dots, x_{1,n_1}, \dots, x_{n_y,1}, \dots, x_{n_y,n_{n_y}})) = \sum_{j=0}^{n_y} (n_j \cdot k_0 \cdot E(U_0(x_{0,j}))) \quad (3.12)$$

In calculating the certainty equivalent of a specific number of fires occurring during a period of several years one needs to find the present monetary value for which the decision maker is indifferent between paying it (its being assumed to be a negative value) and exposing

himself/herself to the uncertain consequences of the fires. To determine the certainty equivalent, one can use equation (3.5) to find an expression for the utility of the certainty equivalent and set that equal to the quantity given by equation (3.12), as done in equation (3.13). Note that the right side of equation (3.13) contains the conditional utility of the certainty equivalent times k_0 , plus the utility of the consequences of all the fires that are assumed to not occur, the conditional utility of such a fire is 1, times k_0 .

$$\sum_{j=0}^{n_y} (n_j \cdot k_0 \cdot E(U_0(x_{0,j}))) = k_0 \cdot U_0(CE) + \left(\left(\sum_{j=0}^{n_y} n_j \right) - 1 \right) \cdot k_0 \quad (3.13)$$

Using equation (3.13) to calculate $U(CE)$ results in equation (3.14).

$$U_0(CE) = 1 - \sum_{j=0}^{n_y} (n_j \cdot (1 - E(U_0(x_{0,j})))) \quad (3.14)$$

The goal of the analysis is to calculate the certainty equivalent of a fire exposure (CE_E) and not the certainty equivalent of a specific number of fires each occurring at a different time (CE), where the latter can be calculated using equation (3.14). Accordingly, n_j in equation (3.14), which is the number of fires occurring the j th year, needs to be replaced by λ_j , which is the number of fires *expected* to occur that year. This results in equation (3.15), in which the certainty equivalent of the fire exposure (CE_E) is calculated.

$$CE_E = U_0^{-1} \left(1 - \left(\sum_{j=0}^{n_y} \lambda_j \cdot (1 - E(U_0(x_{0,j}))) \right) \right) \quad (3.15)$$

Note that this equation yields the same CE_E as equation (4) in paper 2 [64], provided that in that equation the utilities of the different fire scenarios are given in numbers between 0 and -1, which means the worst consequence being assigned a utility value of -1 and the best a utility value of 0.

A simpler method for estimating the CE of a fire exposure

Although the work effort involved in calculating the expected utility of a fire exposure using the procedure suggested above (equation (3.15)) is probably not very great in practice, a simpler way of calculating the CE can be suggested, one in which an easier approach is taken to calculating the expected utility of a fire exposure or, better expressed, the certainty equivalent of it. This method is not derived from any such assumptions regarding the decision maker's utility function as presented above, but instead is developed in a more intuitive way.

The idea behind this simpler method is to divide the period of interest into segments of one year each. The certainty equivalent of a fire occurring during any one of these one-year periods is assumed to be calculated by first discounting the uncompensated losses (x) to their present value and then calculating the expected utility of such a fire by use of the utility function for a fire occurring at present. Thus, the intrinsic monetary value $x_{i,j}$ of fire scenario i occurring during year j can be calculated by use of equation (3.10).

The certainty equivalent of a fire exposure (CE_E) can be calculated by taking the sum, for all the years considered, of the certainty equivalent of a single fire times the number of fires expected to occur in a single year, λ_j . Denote the certainty equivalent of a fire occurring in year j as CE_j , and the total number of years considered in the analysis as n_y . The certainty equivalent of a fire exposure can then be calculated using equation (3.16).

$$CE_E = \sum_{j=0}^{n_y} \lambda_j \cdot CE_j = \sum_{j=0}^{n_y} \lambda_j \cdot U^{-1}(E(U_0(x_{0,j}))) \quad (3.16)$$

This method for calculating the expected utility of a fire exposure can be easier to use in practice than the method suggested in the previous section. It is more difficult, however, to determine whether a specific decision maker's preferences are adequately captured by use of this method than by use of the previous one. Note that this way of calculating the CE is very similar to the approach suggested by Ramachandran (see [3] page 146 equation (8.21)).

An illustrative example

A limited hypothetical example will be used to show how the methods for evaluating a fire exposure taken up above can be used in practice.

Assume that a risk manager of a firm wants to evaluate the investment in a smoke detection system for a factory, believing this will enhance the chances employees have of extinguishing a fire. To describe the outcome of a fire in the standard form $(p_1o_1, p_2o_2, \dots, p_n o_n)$, the decision maker can create a model for the development of a fire, a model taking account first of whether the fire is detected before it reaches a certain size, above which it is seen as impossible for the employees to extinguish and secondly of whether the employees will succeed in extinguishing it given that it is possible for them to do so. An event-tree model illustrating the different fire scenarios is shown in Figure 7. In one of them the fire is not of sufficient potential for it to cause any significant damage. This could be a fire occurring in a metal wastepaper basket, for example, with no possibility of spreading further. The second fire scenario is of a fire with the potential of destroying the building but which is extinguished by the employees. The last scenario is of a fire with the potential of destroying the building and that the employees are unable to extinguish. The consequence assigned to each of the fire scenarios is a monetary value that the decision maker considers to represent the consequences involved, i.e. the uncompensated losses. In estimating the utility values associated with the consequences in each case, one needs to create a utility function that can be used to translate each of the monetary consequences into a utility value. The function represents the decision maker's attitude towards risk. It can be found by having the decision maker answer a set of questions regarding his/her preferences concerning various uncertain situations. To start with, the best and the worst consequence can be assigned a utility value of 0 and 1, respectively. The utility value of the remaining fire scenario can be found by letting the decision maker decide for which probability value p he/she is indifferent between the two decision alternatives, as shown in Figure 8. Alternative 1 involves an uncertain situation in which the decision maker will lose \$1,000,000 with a probability of $1-p$, alternative 2 involves the decision maker losing \$10,000 with certainty. The utility of the second fire scenario from the top of the event tree in Figure 7 (scenario o_2), that of the employees extinguishing a fire which has the potential of destroying the building, is equal to the probability p ($U(o_2) = p$).




Figure 7 Illustration of three possible fire scenarios in a factory.




Figure 8 Illustration of the decision maker's being indifferent between two decision alternatives.

Assume that if p is 0.996 the decision maker is indifferent between the two decision alternatives shown in the figure. One can plot the decision maker's utility curve using the three points (0,1), (-10000, 0.996) and (-1000000, 0). That plot, however, would provide only a very rough description of the decision maker's risk attitude. In order to have a greater number of points on the utility curve, one should thus ask the decision maker further questions regarding various uncertain situations similar to those presented above. Assume, for example, that in investigating the utility of the monetary outcomes -\$100,000, -\$200,000, etc. one finds that the decision maker's preferences towards monetary losses can be represented by a utility function of the form given in equation (3.17), which is shown in Figure 9. In that equation, $U(x)$ is the utility of the monetary outcome x .

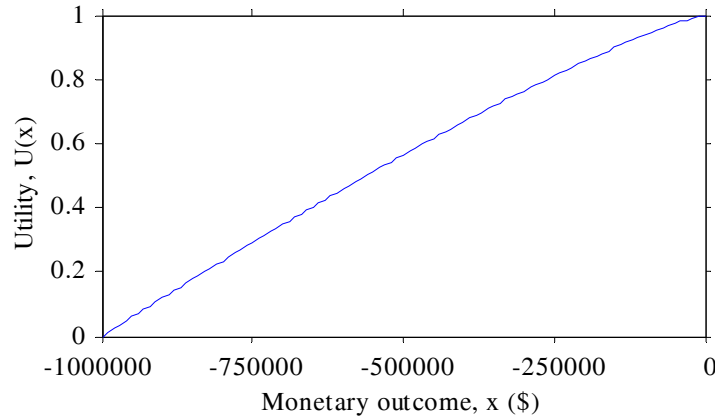


Figure 9 Illustration of a utility function.

Having determined the utilities of the different fire scenarios enables one to calculate the expected utility of a single fire (see equation (3.1)). Since, as was pointed out above, what is usually of interest is the expected utility of a fire exposure, however, one needs to take account of the possibility of more than one fire occurring during a specified period of time.

$$U(x) = 1 - \left(\frac{x}{-1000000} \right)^{1.2} \quad -1000000 \leq x \leq 0 \quad (3.17)$$

In determining the expected utility of a fire exposure, one needs to assess the decision maker's utility function for more than one fire. Assume that the risk manager accepts the assumptions regarding the decision maker's preferences for fires occurring at different times discussed above (the assumptions leading to equation (3.15)). Assume that the decision maker decides in the case just described to use an r -value of 0.1 (in equation (3.10)) and wants to evaluate the exposure over a period of five years. Assume also that the probability of a fire having only a small potential (p_{pot}) is 0.95, the probability of the employees extinguishing the fire (p_{emp}) is 0.7, and that the frequency of fires (the expected numbers of fires) in the building is estimated to be 0.5 per year. Using equation (3.10), one can calculate the intrinsic monetary value (at present) of the different fire scenarios occurring at different times. Using equation (3.17) one can also calculate the corresponding utility values for each of

these scenarios (see Table 2). Using the probabilities of the different fire scenarios, these utility values can be used to calculate the expected utility of a fire occurring during a particular year, which in turn can be used to calculate the certainty equivalent of the fire exposure (see equation (3.15)). The result is -\$49,600. Thus, the decision maker should be willing to pay \$49,600 in order to avoid suffering the consequences of the fire exposure of interest. The expected *monetary loss* due to this fire exposure is \$29,100 (using a discount rate of 0.1 per year). The difference between the expected monetary outcome and the CE, in this case \$20,500, is called the *risk premium* [36].

Table 2 *The intrinsic monetary value of uncertain consequences occurring during different periods of time and the corresponding utility values.*

Present		Year 1		Year 2	
Consequence	Utility	Consequence	Utility	Consequence	Utility
0	1	0	1	0	1
-\$10,000	0.9960	-\$9,091	0.9964	-\$8,264	0.9968
-\$1,000,000	0.0000	-\$909,091	0.1081	-\$826,446	0.2045

Year 3		Year 4		Year 5	
Consequence	Utility	Consequence	Utility	Consequence	Utility
0	1	0	1	0	1
-\$7,513	0.9972	-\$6,830	0.9975	-\$6,209	0.9978
-\$751,315	0.2904	-\$683,013	0.3671	-\$620,921	0.4355

Risk-adjusted net present value

Although the certainty equivalent of a particular fire exposure is a measure that could be useful, the real value of calculating certainty equivalent is that it allows one to *compare* on the basis of certainty equivalent the risk of different fire exposures. Assume, for example, that a decision maker would like to have a measure of how much money it is reasonable to pay for a particular fire safety investment for the building considered in the example just presented. In calculating this value, it is reasonable to have an estimate of the certainty equivalent of the *present* fire exposure, as was just calculated above, and to compare it with the certainty equivalent of the fire exposure resulting from the investment in question. The difference between these certainty equivalents is a

measure of how much it is reasonable to pay in order to change from one level of fire exposure to another.

Since the certainty equivalent for the fire exposure of the factory of interest in its original form was calculated above, it is sufficient here to carry out similar calculations for the fire exposure of the factory given that a particular fire safety investment has been made. Assume that the investment under consideration is a smoke detection system intended to decrease the time from the start of fire until the employees can initiate extinguishing operations. Assume also that this investment is estimated to increase the probability that the employees will extinguish a fire (p_{emp}) to 0.85.

Performing the same calculations as above, except for p_{emp} being assumed to be 0.85, yields a certainty equivalent of -\$28,200. Thus, there is an increase in certainty equivalent of \$21,400 if the investment is chosen. In other words, the decision maker should be willing to pay \$21,400 for the investment in question. Note that if the decision maker had had a risk-neutral attitude to risk, the investment in question would have been worth \$14,100.

In analysing different investments in fire safety, one needs to consider the more or less certain costs and benefits of the different decision alternatives as well as the evaluation of risk. Therefore, in estimating how good or bad a specific decision alternative is in comparison with other alternatives, use should be made of the sum of the certainty equivalent of the fire exposure (CE_E) and the value of the costs (c_{Inv}) and benefits (b_{Inv}) that are certain. This sum will be termed the certainty equivalent of the investment in question (CE_{Inv}), the calculation of which is shown in equation (3.18).

$$CE_{Inv} = CE_E + b_{Inv} - c_{Inv} \quad (3.18)$$

Assume that the investment in the smoke detection system costs \$15,000. This would imply, if no other costs (or benefits) are taken into account, that the *difference* between the alternatives in terms of CE_{Inv} would be \$6,400. This value will be termed the *risk-adjusted net present value* of the investment in question. The reason for using a term very similar to that employed in traditional investment appraisal (the net present value method) is that the two values are calculated in a similar way. The difference is that in the model presented here use is made of

the *intrinsic* monetary value of different outcomes, whereas in investment appraisal the *actual* monetary value is employed. Also, in traditional investment appraisal no consideration is taken of risk attitude, whereas in the analysis method presented here risk attitude can have a very strong effect on the end result. The classical net present value and the risk-adjusted net present value are used in a similar fashion. If the net present value is positive, an investment is considered to be good and should be made. The same can be said of the risk-adjusted net present value, in that if that value is positive the benefit in terms of risk reduction is greater than the costs, meaning that the investment should be made. Note that in determining a risk-adjusted net present value for an investment one needs to compare the investment under consideration with an alternative. That alternative is usually that of keeping the building in its present state.

The model just presented for calculating the expected utility, the certainty equivalent and the risk-adjusted net present value of an investment in fire safety will be referred to as the primary model or primary method. Later in the thesis this model will be complemented with various ways of dealing with epistemic uncertainty, the entire process of using the primary model, together with the different ways of dealing with epistemic uncertainty, being termed the *framework* for analysing investments in fire safety. The different concepts contained in the primary model are shown in Figure 10. The texts in italics positioned between the boxes are the conditions necessary in order to go from the one box to the other.




Figure 10 Illustration of the process of determining the Risk-adjusted net present value.

Comments on some of the assumptions

Note that in the process of developing a way of calculating the expected utility of a fire exposure, also termed the primary model, a number of assumptions were made. Thus far, nothing has been said about the epistemic uncertainty regarding probabilities and consequences of different fire scenarios. Instead, it has been assumed that the decision maker can express these parameters using specific values. Suggestions concerning how this type of uncertainty can be dealt with, which is a very important part of the thesis, will be considered in a later section.

The first of the assumptions referred to that will be taken up is the assumption that the monetary consequences (uncompensated losses) due to a fire occurring in the future are seen as being equivalent to the monetary consequences of the fire occurring at present, given that its consequences are determined by the formula for discounting of future losses (equation (3.10)). This assumption does not appear very controversial, since the principle of discounting is widely used in capital investment appraisal. If the decision maker feels, however, that the assumption does not apply to his/her preferences, he/she can use equation (3.6) instead as a point of departure for calculating the expected utility of a fire exposure (provided that the other assumptions holds). The work load in doing this will be greater than if the assumption referred to were accepted, since the decision maker would need to estimate a separate utility function for each year ($U_i(x_{ij})$) and would also need to estimate a k -value (see equation (3.6)) for each of the years.

The second assumption that will be taken up is that of additive utility. This assumption is probably the one that decision makers would have greatest difficulties in accepting of the assumptions used in deriving the equation for the certainty equivalent of a fire exposure. However, if the decision maker does not accept this assumption (but only accepts the independence assumption), the calculation of the certainty equivalent of a fire exposure becomes very difficult, especially if a large number of years are considered in the analysis. This is due to the k -term in equation (3.2), which will lead to the calculation of the certainty equivalent being more difficult than when the assumption of additive utility is employed. Although it would probably be possible to find some way of calculating the certainty equivalent in this manner, doing so is not aimed at in the thesis.

The third assumption made in the primary model that will be taken up is that the probabilities of different fire scenarios are constant. Thus, in calculating the CE of the fire exposure it is assumed that the probability of a major fire, for example, is the same the first year, the second year, etc. Note that it is possible to make use of probabilities that are dynamic in character, i.e. that change from year to year. Since no investigation has been made (as far as the author is aware) of how much the probabilities associated with different fire protection systems change over time, or if they do change at all, the model described in the present thesis assumes the probabilities to be constant. This area would be interesting to investigate further, however, in future research.

Another assumption made in the primary model is that fires can occur at any time during the period of interest. This means that even though a serious fire may have just occurred, a new fire can take place at any time. Although a more realistic approach might be to assume that during a particular period of time following the first fire a new fire would not take place. This would call for a more complicated analysis since the time required for building up the production capacity again after a fire depends upon which fire scenario is involved. Such an analysis is outside the scope of the present thesis.

Dealing with epistemic uncertainty – prescriptive decision rules

The only type of uncertainty dealt with in the primary model is that of aleatory uncertainty, which is irreducible uncertainty regarding the outcome of an uncertain situation. There is another type of uncertainty, however, which is highly important in the present context, that of epistemic uncertainty.

In modelling the development of a fire in a specific building, various nondeterministic features of the building and of the environment need to be taken into account. By nondeterministic is meant "...that the response of the system is not precisely predictable because of the existence of uncertainty in the system or the environment, or human interaction with the system." [65]. Epistemic and aleatory uncertainty are distinguished on the basis of whether a source of nondeterminism is irreducible or is reducible, as it is in the former and the latter case, respectively [65]. Epistemic uncertainty and how to deal with it in analysing investments in fire safety is one of the key topics in the thesis and is taken up in all the papers included here [18], [38], [48] and [64] (Papers 4,1,3 and 2).

Note that epistemic uncertainty has been variously called type B uncertainty [66], reducible uncertainty, subjective uncertainty and cognitive uncertainty [65], and has also been conceptualised in terms of "second-order probabilities" (see, [67], for example). Here, the term epistemic uncertainty will be employed. Sometimes, if the meaning is clear from the context, only the term uncertainty will be used.

To clarify how the terms epistemic uncertainty and aleatory uncertainty are used here, consider the following example. Assume that whether the

employees in a specific factory of interest succeed in extinguishing a fire plays an important role in determining the outcome of the fire. Since before a fire has occurred one is unable to know whether the employees will succeed in extinguishing a fire, what is involved here is aleatory uncertainty. Although the value of the probability that the employees will succeed in extinguishing a fire may likewise be uncertain, that uncertainty can be reduced if further information is received, such as concerning the number of fires the employees succeeded in extinguishing earlier. Thus, the uncertainty concerning the value of the probability in question represents epistemic uncertainty.

One might argue that the uncertainty concerning whether the employees will succeed in extinguishing a fire is also an epistemic uncertainty since if one knew the exact location of the fire, as well as the location of the employees, the heat release rate of the fire, etc., one could predict whether the employees would succeed in extinguishing the fire. If one pursued such a line of argumentation in a consequent way, this would lead to there being no aleatory uncertainties but only epistemic ones. In the present context, however, the occurrence and development of fire in a specific building is seen as a nondeterministic process in which some uncertainties are *practically* possible to reduce whether others are not, the first type being called epistemic uncertainties and the second type aleatory. In the case studies included in the thesis, the aleatory uncertainties are presented in the form of the event trees shown in appendix D and E and the epistemic uncertainties, which in this case are uncertainties concerning damage costs and probability values, are presented in appendix B and C.

In Bayesian decision theory, no distinction between aleatory and epistemic uncertainty is made. This has been criticised since there is difference between a decision situation in which a decision maker is very certain regarding the probability values, such as in many games, and one in which he/she is uncertain about these values (see for example Ellsberg's criticism in [39]). It is possible, however, to describe the epistemic uncertainty regarding a particular probability value in using Bayesian decision theory, but this uncertainty makes no difference there in the evaluation of the decision alternatives. The reason for this is that, although the epistemic uncertainty regarding a probability value can be described in Bayesian decision theory by use of a probability distribution defined over the different probability values of interest and this probability distribution can be taken into account in calculating the

expected utility of a decision alternative, only one expected utility value for a given decision alternative is calculated. That expected utility value provides no information regarding the epistemic uncertainty the decision problem contains. In Bayesian decision theory the expected utility of an alternative is the only thing that should matter to a decision maker in choosing between alternatives, all information relevant to the decision being incorporated into that value. Although in the present context the expected utility is the value employed in choosing between alternatives, this value can be complemented by an evaluation of the effect of epistemic uncertainty on the expected utilities. Thus, in addition to calculating the expected utility of a decision alternative, a decision maker can also estimate the effect which the epistemic uncertainty regarding the probabilities and utilities involved can have on the expected utilities of the alternatives. First, however, one needs to consider the question of why a decision maker would want to analyse the epistemic uncertainties in addition to calculating the expected utilities of decision alternatives.

First of all, one can note that people in general seems to think of decision situations involving epistemic uncertainty as being less desirable than those not involving it. Ellsberg's account of situations in which people were asked to place bets on the event that a red or that a black ball will be drawn from an urn confirms this [39]. The situation he described was as follows: A person sitting in front of two urns labelled 1 and 2 knows that urn 1 contains 100 balls but does not know in what proportion there are red or black. He/she also knows that urn 2, in turn, contains exactly 50 red and 50 black balls. Most people are found to be indifferent between placing a bet of \$100 on the event of drawing a black ball from urn 1 and placing a bet of the same size on the event of drawing a red ball from urn 1. The same type of preference holds for urn 2. Many people feel, nevertheless, that they would rather bet on red in urn 2 than on red in urn 1 and that they would rather bet on black in urn 2 than on black in urn 1. These preferences violate the axioms of expected utility and show that people tend to view a situation in which the probabilities of the outcomes are known as being basically different from a situation in which probabilities are not known exactly. The question, however, is whether this should matter in efforts to determine which decision alternative is best. Although in terms of normative Bayesian decision theory the answer is no, in the present context, in which a prescriptive method for analysing investments in fire safety is of interest, the question is not easy to answer. In analysing investments in fire safety in

a specific building, statistical information regarding some of the probabilities of interest, such as the probability that the employees will succeed in extinguishing the fire, is likely to be scarce. This can result in large epistemic uncertainties. Since decision situations that differ in terms of epistemic uncertainty are viewed as being basically different by people generally, it would seem not unreasonable to assume that a decision maker would want to have some notion of the epistemic uncertainties when analysing investments in fire safety.

Another reason for considering epistemic uncertainties separate from aleatory uncertainties is that the alternative which is best may change as new information arrives, depending on the degree of epistemic uncertainty present. To illustrate what is meant, consider the following decision situation. Assume that if a decision maker chooses decision alternative 1, event A will occur with probability p_A and event C with probability $1-p_A$, but that if the decision maker instead chooses alternative 2, event B will occur for certain (see Figure 11).




Figure 11 Illustration of two decision alternatives.

Assume furthermore that the utility of event A occurring is 1, of event B occurring is 0.45 and of event C occurring is 0. Assume in addition that the decision maker is uncertain (epistemic uncertainty) about the value of probability p_A and that this uncertainty can be represented by conditioning the value of p_A on the parameter θ_A , the value of which can be represented by a probability distribution. The value of $p_A(\theta_A)$, or the conditional probability that event A will occur, is equal to θ_A , thus $p_A(\theta_A) = \theta_A$. This is an example of how epistemic uncertainty could be dealt with in Bayesian decision theory. Note that $p_A(\theta_A)$ could have been expressed instead as $P(A|\theta_A)$.

Assume that the probability distribution in question is a beta distribution with the parameters $\alpha=1$ and $\beta=1$ (see equation (4.5)), $f(\theta_A)$ thus being a

beta distribution with the parameters just referred to. Since the expected value of such a distribution is 0.5, the expected utility of alternative 1 is 0.5. This implies that Bayesian decision theory would recommend the decision maker to choose alternative 1 since that alternative has the highest expected utility. Assume that the probability $p_A(\theta_A)$ is connected to some event of interest to the decision maker. It could be the probability, for example, that the employees would succeed in extinguishing a fire in a specific building. Assume that the decision maker learns that the employees failed to extinguish a fire that occurred in that building recently. This information, together with the prior distribution $f(\theta_A)$, would result in a posterior distribution for the probability in question. The posterior distribution, in this case would be a beta distribution with the parameters $\alpha=1$ and $\beta=2$. The expected value of such a distribution is $1/3$. Thus, after the new information was received, the expected utility of alternative 1 would be $1/3$, which is lower than the expected utility of alternative 2. Accordingly, Bayesian decision theory would recommend the decision maker to choose alternative 2. Which alternative was best changed, therefore, simply because of the decision maker receiving a seemingly small piece of information regarding the probability of $p_A(\theta_A)$. If, on the other hand, the prior distribution of $p_A(\theta_A)$, $f(\theta_A)$, had not have been as “vague” or “flat” as the one used in this example, which alternative was best might not have changed when the new information was received. For example, if the prior distribution had been a beta distribution with the parameters $\alpha=10$ and $\beta=10$, arrival of the new information would not have changed the alternative adjudged to be best. Note that when the term “epistemic uncertainty regarding a probability value” is employed it refers to a distribution of the type discussed above ($f(\theta_A)$).

Due to the potential presence of large epistemic uncertainties (“vague” prior distributions) in decision analyses concerned with fire safety and to their effect on whether the best alternative will change as a result of more information concerning a decision problem being obtained, it is in the decision maker’s interest to take account of epistemic uncertainties in making decisions concerning fire safety. Doing so enables him/her to identify situations in which a small amount of information can markedly change the conclusions to be drawn.

Two ways of taking account of epistemic uncertainties when analysing a decision are discussed in [38], [48] and [64] (Papers 1, 3 and 2), which

are included in the thesis. These two methods differ in terms of how much information regarding epistemic uncertainties the decision maker needs to provide in order to perform a decision analysis. The first method is termed here extended decision analysis, and the second method Supersoft decision theory. Extended decision analysis involves the epistemic uncertainty regarding the probability and utility values being expressed as probability distributions. This uncertainty is then related to the expected utility of the decision alternatives in question, allowing the expected utility to be expressed as a probability distribution. Note that if one uses Bayesian decision theory as described in chapter 2, an approach that will be termed traditional decision theory, the only value of interest in such a distribution is the expected value of the distribution. It can be argued, however, that the form and position of the distribution which results from relating the epistemic uncertainty to the expected utility of a decision alternative provides the decision maker valuable information and should thus be included as a basis for the decision. This method is discussed in [38], [48], [64] (Paper 1, 3 and 2), and in the next section.

The second method, Supersoft decision theory, does not require that the decision maker provide as much information as extended decision analysis requires. Instead of using distributions to represent epistemic uncertainty regarding probability and utility values, intervals are employed there. Consequently, it is not possible to utilise the maximisation of expected utility criterion, a set of criteria being used instead. Supersoft decision theory as applied in the present context is discussed in greater detail in [48] (Paper 3).

Extended decision analysis

Assume that epistemic and aleatory uncertainties are separated, such that the occurrence of specific events during a fire, such as the sprinkler system's extinguishing or not extinguishing a fire, are treated as representing aleatory uncertainty, and that the uncertainty regarding the value of the probability of such events is treated as representing epistemic uncertainty. This implies that, irrespective of the amount of information concerning the sprinkler system and its environment that has been received, one cannot determine for certain whether the sprinkler system would extinguish a fire. The information could be used, however,

for reducing the uncertainty concerning the probability that it would extinguish a fire.

Using this approach to separating epistemic and aleatory uncertainty, one can express the probability of a specific fire scenario i as a conditional probability $p_i(\theta_i)$, the value of which depends on the parameter θ_i . In this case, θ_i is the value that the probability of scenario i occurring has. Since there can be epistemic uncertainty regarding this value, it is represented by a probability distribution $f_i(\theta_i)$ showing how likely the different values of the probability in question are. Consequently, uncertainty regarding a utility value $U_i(\gamma_i)$ is represented by the probability distribution $g_i(\gamma_i)$. If one wishes to analyse the effect of the epistemic uncertainties on the expected utility of an alternative, one can calculate the *conditional* expected utility of a fire, conditional on the values θ_i and γ_i (see equation (3.19)). Since the values are uncertain due to epistemic uncertainty and are thus represented by probability distributions, one can express the conditional expected utility $E(U|\theta, \gamma)$ as a probability distribution. That distribution shows how likely different values for the expected utility would be if no epistemic uncertainties were present. Although such a distribution, if seen in isolation, does not provide the decision maker very much information, when several decision alternatives are analysed and one can start comparing them in terms of conditional expected utilities these distributions can provide valuable insight into the decision problem.

$$E(U|\theta, \gamma) = \sum_{i=1}^n p_i(\theta_i) \cdot U_i(\gamma_i) \quad (3.19)$$

When two decision alternatives are compared, one is interested in which alternative has the highest expected utility. The calculation of expected utility when epistemic uncertainty is present is not basically different from the calculation of expected utility when only aleatory uncertainty is present. For calculating the expected utility using continuous distributions to represent epistemic uncertainty use can be made of equation (3.20).

$$\begin{aligned}
 E(U) &= \int \dots \int_D \left(\sum_{i=1}^n p_i(\theta_i) \cdot U_i(\gamma_i) \right) \cdot \\
 &\cdot f_1(\theta_1) \cdots f_n(\theta_n) \cdot g_1(\gamma_1) \cdots g_n(\gamma_n) d\theta_1 \dots d\theta_n d\gamma_1 \dots d\gamma_n \\
 D &= \left\{ (\theta_1, \dots, \theta_n, \gamma_1, \dots, \gamma_n); \sum_{i=1}^n p_i(\theta_i) = 1 \right\}
 \end{aligned} \tag{3.20}$$

In using equation (3.20) one can label the different decision alternatives such that the one with the highest expected utility becomes alternative 1, the one with the second highest expected utility alternative 2, and so on.

According to Bayesian decision theory, this is all that is necessary to determine which alternative is best. As has been pointed out in the present study, however, the decision maker may wish to also include an analysis of the effect of the epistemic uncertainties on the results obtained (see the previous section). In doing this, one can analyse the *robustness* of the best decision alternative. The term robustness is used here to denote how likely it is that the best alternative would change if all epistemic uncertainties were eliminated through the decision maker's receiving sufficient resources in terms of time and information to allow this to be done. A suitable way of presenting the robustness of a decision alternative to a decision maker, where there are only two decision alternatives, is to provide him/her with a distribution showing the difference in terms of conditional expected utilities between the best alternative, $E(U_1|\theta, \gamma)$, and the other alternative, $E(U_2|\theta, \gamma)$. Note that one could also examine the difference in Certainty equivalents between two alternatives, since if the Certainty equivalent of one alternative is higher than that of another, the expected utility of that alternative is also higher.

If one does not wish to present the results in the form of distributions, one can instead calculate a *robustness index*, R . This is the probability that the conditional utility of the alternative seen as best ($E(U_1|\theta, \gamma)$) being higher than that of the other alternative ($E(U_2|\theta, \gamma)$) (see equation (3.21)). One can also consider this value as being the probability that if the epistemic uncertainties are eliminated the alternative seen as best will be the same.

$$R = P(E(U_1|\theta, \gamma) > E(U_2|\theta, \gamma)) \tag{3.21}$$

If there are more than two decision alternatives, the decision maker can compare the alternative seen as best with each of the others and then calculate various robustness indexes.

In order to exemplify the extended decision analysis approach, use will be made of the brief example discussed in the previous section. Assume that the decision maker cannot assign precise probability values to the probabilities used in the example, instead using probability distributions to represent the epistemic uncertainty regarding these values. Assume that the decision maker believes the minimum plausible value for the probability that the potential of the fire is small (p_{pot} in Figure 7) is 0.9, that the most likely value is 0.95 and that the maximum value is 0.99. The uncertainty involved here can be represented by a triangular probability distribution having the values just mentioned as its minimum, most likely and maximum values, respectively. Assume that a triangular distribution having the minimum value of 0.5, a most likely value of 0.7 and a maximum value of 0.9 is being used to represent the uncertainty regarding the probability that the employees will succeed in extinguishing a fire (p_{emp} in Figure 7). Assume in addition that the uncertainty regarding the frequency of fires in the building is represented by a triangular distribution having the minimum value of 0.3, the most likely value of 0.5 and the maximum value of 0.7.

Since the probability values in the model shown in Figure 7 are represented by probability distributions, it is possible to express the Certainty equivalent of a decision alternative as a probability distribution. A histogram representing this epistemic uncertainty regarding the Certainty equivalent of the fire exposure discussed above is provided in Figure 12.

In arriving at this histogram, 10000 Monte Carlo simulations were performed. One can perform this same type of analysis for the decision alternative of investing in a detection system, the difference being that the probability value for the employees succeeding in extinguishing the fire is increased. The triangular distribution representing this probability value has a minimum value of 0.7, a most likely value of 0.85 and a maximum value of 0.95. The histogram of the CE of the exposure (not counting the investment costs) resulting from the investment in a detection system is shown in Figure 13.




Figure 12 *Histogram illustrating the epistemic uncertainty regarding the certainty equivalent of the exposure when keeping the building in its present state.*




Figure 13 *Histogram illustrating the epistemic uncertainty regarding the certainty equivalent of the exposure for the alternative of investing in a detection system.*

It is clear that the epistemic uncertainty regarding the probability values in the example above has a strong effect on the CE of the exposures. It is

not clear, however, what effect this should have on the decision to be recommended. In the analysis of the two decision alternatives using the primary model, without epistemic uncertainty being considered, the alternative of investing in the detection system was recommended since that alternative has a higher CE than the alternative of keeping the building in its present state. The risk-adjusted net present value of the detection investment was \$6,400. Since the CE:s of the two exposures are expressed as probability distributions, it is also possible to express the risk-adjusted net present value of the investment as a probability distribution. A histogram representing the risk-adjusted net present value, one which is the result of 10000 Monte Carlo iterations, is shown in Figure 14.




Figure 14 Histogram showing the risk-adjusted net present value of the investment in a smoke detection system.

In using the primary method, the value of the risk-adjusted net present value was estimated to be \$6,400. Looking at Figure 14, one can see that there is considerable epistemic uncertainty associated with the estimate. There is a good chance, in fact, that the risk-adjusted net present value is negative if the epistemic uncertainties could be reduced, which means that the investment should not be made. A purely intuitively judgment based on looking at the histogram in Figure 14 does not allow one to determine which alternative is best.

Although the example discussed above is hypothetical, it serves to illustrate the difference between using only the primary model and using the primary model in combination with specific consideration of the epistemic uncertainties. Having performed an estimate of the effect of epistemic uncertainties on the risk-adjusted net present value, the question arises of how best to interpret the kind of distribution shown in Figure 14 in terms of deciding which decision alternative is best.

The decision rule that is used here is the maximisation of expected utility. That rule, however, is complemented here by an additional evaluation, that of the robustness of the decision which was discussed earlier. Decision robustness is also discussed in [38] and [48] (Paper 1 and 3).

As noted above, one way of measuring this is to look at the distribution of the difference in conditional expected utility between two alternatives or the risk-adjusted net present value and to note whether the distribution overlaps the 0-value. If no overlap is found, one can conclude that the decision is robust, since the decision alternative deemed best is the best for all plausible values of the probabilities and utilities involved. This indicates that one should choose the alternative having the highest expected utility. This is the first case presented in Table 3. If the distribution overlaps the 0-value, however, as illustrated in Figure 14, the situation is quite different. Here, there are plausible combinations of probability and utility values that lead to the one decision alternative being best and other combinations which lead to the other alternative being best. It can be suggested that the decision rule of maximising expected utility be employed, but that instead of only providing the decision maker with information regarding which alternative is best and using an exact estimate of the risk-adjusted net present value, one should also provide the decision maker information concerning the robustness of the decision. This can be achieved by presenting the decision maker histograms, such as the one in Figure 14, and/or the *robustness index*. The index is defined in equation (3.21) and it is equal to the amount of the resulting distribution (for the difference in conditional expected utility between two alternatives) that is positioned on the side which indicates the best alternative to be the one with the highest expected utility. In the example discussed above, the robustness index is 58%. Thus, for 58% of the values illustrated by the histogram in Figure 14 are positive. This index can serve as an indicator of when it is not possible to draw a clear conclusion regarding which decision alternative is best.

How low the index needs to be in order for the decision situation to appear unclear, so that no definite conclusion can be drawn regarding which alternative is best, is up to the individual decision maker to decide. It seems reasonable, however, to consider a decision situation to be robust if the index is somewhere around 90% or higher. Accordingly, the decision situation shown in Figure 14 would not be considered to be robust.

Different types of decision situations can occur in this respect. The distribution of the risk-adjusted net present value can be positioned completely on the positive part of the scale. This is the situation illustrated by case 1 in Figure 15. Another type of decision situation is that of the distribution of the risk-adjusted net present value overlapping the 0-value but its not exceeding the limit for considering the decision to be robust. This is illustrated by case 2 in Figure 15. Still another situation is that of the distribution of the risk-adjusted net present value overlapping the 0-value and the overlap being greater than the decision maker can accept in order to consider the decision robust. In such a case it can be recommended that the decision maker search for further information regarding the problem at hand so as to reduce the epistemic uncertainty. Another possibility, if the major part of the distribution of the investments risk-adjusted net present value is positioned on the negative side of the scale, is to not consider the investment at all. This decision situation is illustrated by case 3 in Figure 15. The last decision situation to be considered is that of the entire distribution of the risk-adjusted net present value being on the negative part of the scale. In such a case, the investment should not be made. This is illustrated by case 4 in Figure 15. The four decision situations just taken up are summarised in Table 3.

Table 3 *Description of different decision situations and suggested actions.*

Decision situation (the case numbers refer to Figure 15)	Suggested action
Case 1: All the plausible values of the risk-adjusted net present values indicate the investment of interest to have the highest expected utility.	The decision situation is robust. The investment should be made.
Case 2: <i>Almost</i> all of the plausible values of the risk-adjusted net present values indicate the investment of interest to having the highest expected utility. The robustness index is above the limit for deeming the decision situation to be unclear.	The decision situation is not 100% robust but the investment should nevertheless be made.
Case 3: The robustness index lies below the limit for deeming the decision situation to be clear.	The decision situation is not robust. It is not possible to determine which alternative is best. The decision maker should look for more information regarding the parameters of interest so that the epistemic uncertainties can possibly be reduced and a robust decision situation be achieved.
Case 4: All of the plausible values of the risk-adjusted net present values indicate the investment of interest to <i>not</i> have the highest expected utility.	The investment should not be made.

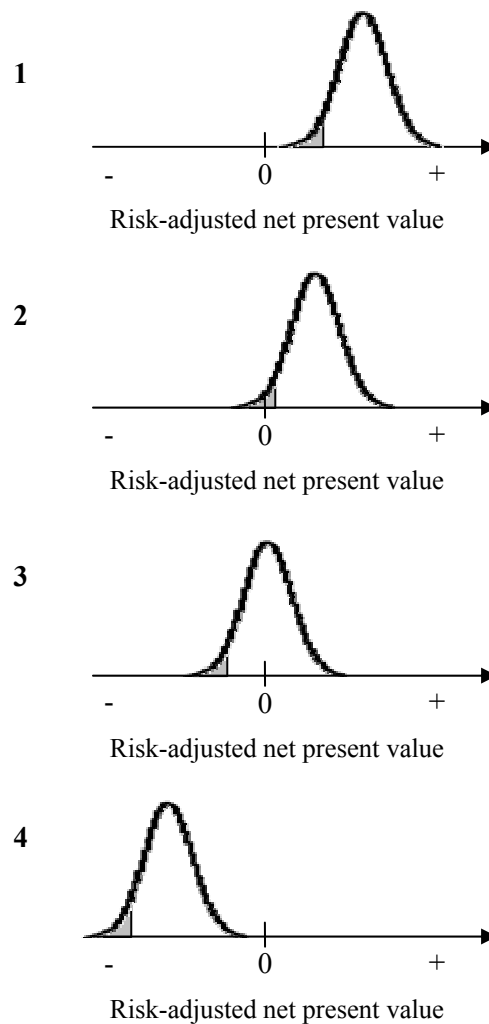


Figure 15 Illustration of different decision situations. The shaded area of the distributions represents the part of the distributions that is allowed to be on the negative side on the scale, such that the decision is still considered to be a robust one.

Note that the illustrations in Figure 15 is concerned with the comparison of the decision alternative of making an investment in fire safety and the primary decision alternative, usually that of keeping the building in its present state. If more than two decision alternatives are available, the

decision maker should start by evaluating the expected utility of each of the decision alternatives so as to be able to identify the alternative with the highest expected utility. This alternative could then be compared with all the other alternatives in terms of the difference in conditional expected utility. This comparison could be performed using extended decision analysis, in which the differences in conditional expected utility are expressed as probability distributions. The concept of robustness could then be used for this type of comparison so as to determine whether the best decision alternative is robust.

Having performed an analysis using the form of extended decision analysis described here provides the decision maker with (1) an evaluation of the different decision alternatives in terms of a monetary value, based in part on an evaluation of the risk reduction that the investment in question achieves, (2) a recommendation of which decision alternative to choose, and (3) an evaluation of the robustness of the decision situation. This should give the decision maker a sound basis for the decision to be made.

Supersoft decision theory

The form of extended decision analysis just considered provides the decision maker with ways of expressing his/her epistemic uncertainty regarding probabilities and utilities. Sometimes, however, it can be difficult to assign distributions to represent one's uncertainty regarding the probability and utility values of different fire scenarios, especially when one is interested in scenarios that are catastrophic or extreme in other ways. In such cases, it is desirable for the decision analysis method employed to not force the decision maker to be more precise than he/she wants to be. A method allowing this is Supersoft decision theory (SSD) [47]. It will be shown here how SSD can be combined with the primary model, presented earlier in the present chapter, so as to create a framework within which it is possible to evaluate decisions regarding fire safety involving events for which both the probabilities and the utilities are extremely uncertain. Since a thorough presentation of the application of SSD to problems of fire safety is to be found in [48] (Paper 3) the presentation here will be brief.

To illustrate why someone may wish to use Supersoft decision theory rather than a more traditional approach, consider the following hypothetical example:

Assume the decision maker wants to calculate the expected utility of a fire occurring in a particular factory. Figure 16 presents a simple event tree model indicating how a fire might develop in the building. If a fire occurs, it can either be small and run out of fuel or grow larger. The probability p_{small} in the event tree represents the probability that the fire will be small. The consequences if a fire remains small are denoted as A in Figure 16. If the fire grows, it can either be the case that it is extinguished by the employees of the factory, the consequences under such conditions being termed B , or that the employees do not succeed in extinguishing it, the consequences this results in being termed C . The probability that the employees will succeed in extinguishing the fire is termed p_{emp} . Suppose that the decision maker concludes that the utilities of the different consequences are $U(A) = 1$, $U(B) = 0.8$ and $U(C) = 0$. Assume in addition that the probability p_{small} that a fire will remain small, is estimated to be 0.9.




Figure 16 Event tree showing the possible fire scenarios in a hypothetical building.

Assume that in estimating the probability that the employees in the building will succeed in extinguishing a fire, given that the fire is not small, the decision maker runs into trouble, he/she being uncertain regarding the value of p_{emp} , only being able to conclude that it is between 0.2 and 0.8 in value. According to Bayesian decision theory, a decision maker must assign probabilities exactly or use a probability distribution to represent the uncertainty concerning the probability value (this is termed extended decision analysis here). In assigning a probability distribution the decision maker should use whatever information regarding the event in question which he/she has available in arriving at a distribution for representing the epistemic uncertainty regarding the probability in question (see previous section). Some decision makers

may argue that expressing epistemic uncertainties in terms of a probability distribution requires time and resources he/she may not have available. If unable to do so, they cannot make use of Bayesian decision theory here. The question is whatever one can use the (inexact) information the decision maker has provided, despite Bayesian decision theory not being fully applicable. For example, using the information contained in the example presented above, one can conclude that if the decision maker could assign an exact value or a probability distribution to p_{emp} , the expected utility of a fire would lie somewhere between 0.916 and 0.964. This information might be used, for example, to determine that if the expected utility of a fire in another factory is lower than 0.916, the expected utility of a fire in the factory considered is higher than that in the other factory. As mentioned above, an inexact representation of p_{emp} (in the form of an interval) is not allowed in Bayesian decision theory. This means that using intervals for the expected utility in comparing alternatives cannot be done in using that theory.

There are decision theories that do not require that probabilities be assigned exact values. Theories which can use the information contained in the example above to provide the decision maker with a recommendation of which alternative is best. Supersoft decision theory (SSD) is one such method. The reason it is taken up here is that it involves performing expected utility evaluations of alternatives, which means the primary model presented earlier for calculating the expected utility of a fire exposure can be used in combination with SSD. The main difference between Bayesian decision theory and SSD lies in the decision rule that is employed. In Bayesian decision theory decision alternatives are evaluated on the basis of their expected utilities only, whereas the evaluation of alternatives using SSD is based on qualitative evaluations in combination with the evaluation of expected utilities. However, since the probabilities and utilities need not be assigned as precise values or as probability distributions, it is not possible to use the rule of the maximisation of expected utility when evaluating the alternatives. Instead, three criteria are used simultaneously in evaluating a given decision alternative. The key point in using SSD is that by using vague statements regarding probabilities and utilities in the primary model it is possible to identify an interval within which all the plausible values for the expected utility lie. Despite one's not knowing how likely the different values within the interval are, one can still evaluate decision alternatives on the basis of the position and size of the intervals. Vague statements regarding probabilities and utilities are thus viewed as

boundary conditions for the expected utility of an alternative, those boundary conditions being termed here the *decision frame*.

Given a decision frame, it is possible to evaluate the three criteria involved: *Min*, *Max* and *Average*. The *Min* and *Max* criteria are the minimum and maximum values, respectively, for the expected utility of the decision alternative, given the decision frame. The *Average* criteria can be seen as the expected utility, given that the probabilities and utilities are treated as being represented by uniform distributions between their minimum and maximum value (see [48] (Paper 3) for a more thorough account of these criteria).

Assume that the analysis of the investment in the detection system presented earlier is performed using SSD. The decision frame can then be summarised as being:

$$\begin{aligned} p_{\text{pot}} &\in [0.9, 0.99] \\ p_{\text{emp}} &\in [0.5, 0.9] \text{ (For the primary alternative)} \\ p_{\text{emp}} &\in [0.7, 0.95] \text{ (For the alternative to invest in a detection system)} \\ \lambda &\in [0.3, 0.7] \end{aligned}$$

The *Min*, *Max* and *Average* values for the expected utility of the decision alternatives, given the decision frame, can then be calculated. Since the original decision problem was analysed in terms of certainty equivalent (CE), the SSD evaluations will also be performed in terms of CE. This does not matter for the analysis of which alternative is best, since if one alternative has a higher CE than another, it also has a higher expected utility, and vice versa. The results of the evaluation can be found in Table 4.

Table 4 Results of an SSD evaluation of decision alternatives. The monetary sums in the table are the CEs of the respective decision alternatives.

Evaluation	Decision alternatives	
	Investment in detection system	No investment
Min	-\$132,000	-\$178,000
Average	-\$50,000	-\$54,000
Max	-\$17,000	-\$4,000

In evaluating one decision alternative as being better than another, all three criteria need to indicate the alternative deemed best to be higher in expected utility (CE). One can see, in Table 4, that only two criteria indicate the alternative to invest in a smoke detection system as being better. Thus, no clear conclusion regarding which alternative is best can be reached, given the decision frame at hand. This should come as no surprise since in analysing the same decision problem by use of extended decision analysis resulted in Figure 14, a figure which indicates very definitely the decision situation being far from clear.

The application of Supersoft decision theory in the present context is described in greater detail in [48] (Paper 3), which is included in the thesis.

A framework for the evaluation of investments in fire safety

Instead of suggesting the use of a specific prescriptive rule for evaluating investments in fire safety, a framework that can be seen as a kind of “toolbox” and consists of three of the methods taken up in the present chapter is provided, its being up to the decision maker to choose one of the three prescriptive methods that the framework includes.

Before the decision maker applies any of the methods included in the framework, it is important that he/she removes from further consideration any alternatives viewed as unrealistic due to their being too expensive, involving risks to individual persons considered to be too high, and the like. Employing such a “screening” of alternatives before applying the framework presented here is important since the expected utility approach suggested cannot be used to distinguish which alternatives cost more than the budget allows or involve some person or group being exposed to unreasonable risks.

A primary model for estimating the expected utility of different fire safety investments has been presented in the thesis (see equation 3.15). That method can be utilised in conjunction with either extended decision analysis, Supersoft decision analysis or even traditional decision analysis. The evaluation criterion to be applied to the decision alternatives depends on which of these methods the decision maker chooses to employ. There are considerable differences between them, as discussed in detail in [48] (Paper 3), the choice of which model to

employ not being an easy one. It is probably not a very good idea to try to provide rules for when a particular method should be employed. Instead, that choice should be left up to the decision maker.

In making this choice, the decision maker can start by evaluating the three methods in terms of how precise it appears that the estimates of the different parameters in the primary model need to be. The traditional decision analysis method is the one requiring the highest degree of precision, followed by extended decision analysis and Supersoft decision analysis (SSD) in that order. One can always start by analysing a problem by use of SSD and then make use of extended decision analysis if one discovers that the information available justifies the uncertainties being expressed as probability distributions. This is a prudent approach, and if one can conclude by use of the SSD method that one particular decision alternative is clearly the best (see case 1 in figure 9 in paper 3), there is no need to continue using any of the other approaches. One could end up instead, however, with a slightly more complicated situation in which all the SSD criteria results in the recommendation of one particular alternative, but there are also plausible combinations of probabilities and utilities that lead to some other alternative having the highest expected utility (see case 2 in figure 9 in paper 3). Although in such a case one may nevertheless be fairly certain that the alternative to which the SSD method points is best, one might want to have a measure of how robust the decision involved would be. Thus, one might want to carry out an analysis using extended decision analysis and in so doing calculate a robustness index. Note that this is only possible if the decision maker feels that the available information justifies expressing probability values and utility values as probability distributions. If, in contrast, the initial SSD analysis gives no indication of which alternative should be chosen (case 3 in figure 9 in paper 3), the decision maker could continue looking for information so as to attempt to reduce the uncertainties involved and might (if this is possible) perform an extended decision analysis so as to be able to assess the robustness of the decision situation.

The theoretical and practical framework employed is shown in Figure 17, which indicates the type of decision rule to be used, depending on which method is chosen. In using traditional decision analysis in combination with the primary model, the result of evaluating a decision alternative is to obtain a single value for the expected utility. Employing extended decision analysis instead results in a distribution of conditional

expected utility, enabling one to determine the robustness of the decision alternative in comparing it with other alternatives. Use of Supersoft decision analysis, finally, involves the result of evaluating an alternative being expressed as an interval, the interval between the minimum value for the expected utility and the maximum value for it. The average-value is also calculated for each alternative. In the present context, a traditional decision analysis involves using the primary model without considering the effects of knowledge uncertainties, using exact values for the probabilities and consequences instead. This approach is dealt with in paper 2 of the thesis [64] as well as in the present chapter between pages 55-63. The extended decision analysis method is taken up in paper 1 [38] and in the present chapter between pages 68-78, and the Supersoft decision analysis method in paper 3 [48] as well as in the present chapter between pages 78-82.

Note that all three methods employ the evaluation of expected utility, which is the normative decision rule they are based on. In practical terms, however, the three methods differ in terms of how the expected utility criterion is used to evaluate alternatives, the prescriptive decision rules they employ differing.

The framework presented here does not contain any method for the practical elicitation of probabilities or utilities from the decision maker. This matter has been dealt with in various publications (see the summary of methods in [63] and [68], for example), the work carried out here not being aimed at developing any new methods of this sort. Regardless of whether one uses the methods suggested in the references taken up above or one employs some other method, one should be aware of the fact that people are sometimes poor probability assessors and that we are subject to a set of biases that can influence our estimates [69]. An interesting area for future research would be to examine different methods for the elicitation of probabilities and utilities in a context such as the present one. Since such a work can probably be more suitably carried out, however, in a psychological setting than in a technical one, no attempt is made in the thesis to develop such methods.




Figure 17 The different parts of the theoretical and practical framework suggested here for use in evaluating investments in fire safety.

4 Bayesian methods in decision analysis concerned with fire safety

In using the extended decision analysis described earlier in the thesis it is vital to be able to describe uncertainty regarding a probability value using all available information regarding that particular probability. In doing this one needs to be able to combine information from different sources in a logical way. Of particular interest in the present context is the combination of subjective judgement with statistics from the building of interest. The area of Bayesian statistics can be of much use in doing this and it fits very well into the framework of the extended decision analysis model. Therefore, this section is devoted to illustrating how the use of Bayes' theorem can be used in combination with extended decision analysis model described earlier.

The use of Bayesian methods within engineering are by no means new but have previously been widely used (see [70] and [71], for example). In fire safety engineering they have been used in connection with fire safety in nuclear power plants [72], [73] and [74].

Using Bayesian decision theory as a basis for the models developed in the present thesis implies that a subjectivistic view of probability [75] is adopted. The term "probability" will refer to a numerical measure of the state of confidence one has regarding the state of some uncertain quantity. Such a state may well be some particular probability value. Here, θ will be used to denote the uncertain parameter, which may be the probability of some event or a frequency of fire, for example. Thus, $P(\theta = 0.7) = 0.95$ represents the decision maker's being very confident (0.95) of the θ - parameter having a value of 0.7.

In the Bayesian approach one should think of an uncertain parameter such as the frequency of fire as having the possibility of being in one of several states, and in the context at hand as having some specific value. Since one is uncertain of what the value of the parameter is, one's degree of confidence that the parameter has some particular value is expressed by use of a probability measure. If this is done for every value that the parameter can have, this results in a probability distribution defined over all possible parameter values. Such a probability distribution is very useful when one wants to visualise one's belief regarding a specific

parameter. It is this type of probability distribution which is employed in extended decision analysis, discussed in the previous chapter.

If one starts with a particular belief (expressed in the form of a probability distribution) regarding some parameter and receives additional information that one wants to incorporate into the previous body of knowledge, one can make use of Bayes' theorem to do so (see equation (4.1)).

$$P(\theta_i|E) = \frac{P(E|\theta_i) \cdot P(\theta_i)}{\sum_i P(E|\theta_i) \cdot P(\theta_i)} \quad (4.1)$$

$P(\theta_i)$ is the probability that one assigned – prior to obtaining the new information – to θ_i being the correct state. $P(E|\theta_i)$ is the probability that the new information E would have been observed given that the state of the parameter of interest was in fact θ_i . $P(\theta_i|E)$ is the probability one assigns, after obtaining the new information E , to θ_i being the correct state. Since the application of Bayes' theorem in the present context is usually easy when the discrete form of Bayes' theorem presented in equation (4.1) is employed, this will not be taken up here. Instead, the focus will be on various continuous distributions that can be highly useful in the present context. These continuous distributions are the Gamma distribution, the Beta distribution and the Dirichlet distribution. These are conjugate distributions that can be highly applicable to in particular situations in which decision analysis concerned with investments in fire safety is involved. A discussion of how they can be employed in the present context follows.

A common way of performing a quantitative risk analysis (QRA) that can be the basis for a decision analysis is to first attempt to assess the frequency of fire in the building and to then attempt to model the probability of the different fire scenarios that seem possible, using conditional probabilities in an event tree. This approach is used in the case studies presented later in thesis. The event tree approach has also been used for a QRA carried out in a hospitals ward [50] and [76], in a hotel [77] and in an office building [78]. In employing a QRA framework of this sort, one can readily encounter uncertain situations of the following types: (i) that one is uncertain about the frequency of fire in the building of interest, (ii) that one is uncertain about the probability

of an event when there are only two possible alternative events that can occur, and (iii) that one is uncertain about the probability of an event when there are more than two possible alternative events that can occur. These three types of situations will be dealt with in the present section.

In employing a Bayesian approach to analysing fire protection measures by use of the methods presented earlier in the thesis one should begin by collecting and documenting all the evidence or facts available that could be of relevance. This collection of facts would then be called the “evidence base”. The evidence can consist of precise evidence such as that “The sprinkler system in this building has extinguished 9 out of 10 fires”, of imprecise evidence such as that “The sprinkle system in this building has extinguished somewhere around 7 to 9 fires out of 10”, or of expert judgements. This collection of evidence can then be used to change the belief regarding the different uncertain parameters of interest. Previous examples of how to treat precise evidence have been given in [74] and [79]. Examples of Bayesian methods using imprecise evidence have been given in [72], [73] and [74], and examples of Bayesian methods using expert judgements in [72], [80], [81], [82] and [83].

The main benefit of such an approach in dealing with evidence is that the stakeholders can agree upon the question of what evidence that should be used in the analysis before the analysis itself is performed. To document all the applicable evidence is also beneficial since it leaves no questions regarding what the estimations in the analysis is based upon. When all the stakeholders have agreed upon the relevance of the evidence it is possible to use Bayesian methods to combine all the evidence available, the results obtained being used in the final decision analysis of the problem.

Frequency of fire

If one is interested in the frequency of fire in a specific building (λ), the Poisson distribution can be used to calculate the probability that some particular number of fires will occur there during a given period of time. Using the Poisson distribution involves assuming that whatever fires take place occur independently of one another and with a constant tendency to occur. The Poisson distribution can be employed for calculating $P(E|\lambda)$, used in updating one’s belief regarding the frequency of fire on the basis of new evidence (see above). Examples of the use of Bayesian

methods when the process of interest is a Poisson process with respect to fire safety have been given in [74].

When the Poisson distribution is used to calculate $P(E|\lambda)$, the Gamma distribution is a conjugate distribution (see for example [71]). Thus, if one describes one's uncertainty regarding the frequency of fire using a Gamma distribution, the resulting posterior distribution will also be a Gamma distribution but with other parameter values. In using a distribution from the Gamma family here, one must of course agree to its representing one's belief regarding the frequency in question. However, since the Gamma family is flexible, it should be possible to find a distribution of this type suitable to the needs of the decision maker. The Gamma distribution is shown in equation (4.2), where α and β are the parameters of the distribution and λ is the frequency of fire.

$$f(\lambda) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} \exp(-\beta\lambda) \quad (4.2)$$

Since the Gamma distribution is conjugate to the Poisson distribution, updating it by use of Bayes' theorem is very simple. Assume that α' and β' denote the prior parameters of the Gamma distribution and that α'' and β'' denote the respective posterior parameters. The posterior parameters can then be calculated from the prior parameters by use of equations (4.3) and (4.4), in which r is the number of fires that have been observed during the time t [71].

$$\alpha'' = \alpha' + r \quad (4.3)$$

$$\beta'' = \beta' + t \quad (4.4)$$

As an example of a situation where a Gamma distribution can be updated using fire statistics from a specific building a problem from a decision analysis performed in a factory belonging to Avesta Sheffield (one of the case studies included in the present thesis) can be used. In this example, the prior knowledge regarding the frequency of fire in the building, which was a cold-rolling mill, was scarce. Because of this, a vague prior distribution was chosen to represent the belief regarding the frequency of fire in the building before any fire statistics from the building had been observed. A prior Gamma distribution with parameters $\alpha' = 0$ and

$\beta' = 0$ was used to represent a situation with vague knowledge regarding the frequency of fire (recommended in [86]).

Information regarding the number of fires that had actually occurred in the Cold-rolling during the last six years was then obtained. During that period, sixty fires had occurred. This might sound a lot but many of these fires were very small and a large part of them occurred within the rollers (where fires are expected to occur) where they were extinguished by automatic fire extinguishment systems. Using this information regarding the total number of fires in the building during the six-year period it was possible to update the vague prior distribution. The result was a Gamma distribution with the parameters $\alpha'' = 60$ and $\beta'' = 6$, which is shown in Figure 18.




Figure 18 The posterior Gamma distribution representing the belief regarding the frequency of fire λ in the cold-rolling mill belonging to Avesta Sheffield.

Two possible events

In analysing fire scenarios, one frequently encounters uncertain situations that involve only two possible, mutually exclusive events, such as the sprinkler system's extinguishing the fire or not doing so, or the building staff succeeding in extinguishing the fire or not. In such cases one can treat the number of successful trials, such as the number of times a fire was extinguished by the sprinkler system, as being binomially distributed. The parameters of the binomial distribution are p , which in

the present context could be the probability that the sprinkler system would extinguish a fire, and n , which would be the number of fires observed. A situation of this sort has been studied by Apostolakis [72], who used Bayes' theorem to combine indirect evidence with direct evidence for the demand availability of sprinkler system.

The Beta distribution is the conjugate distribution that applies when the probability of receiving the evidence, $P(E|p)$, can be calculated using a Binomial distribution [71]. Thus, by using a prior distribution which is in the form of a Beta distribution one can simplify the use of Bayes' theorem here considerably.

The Beta distribution can be written as equation (4.5), where a and b are the parameters of the distribution itself and p is the probability of successful trials.

$$f(p) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1} (1-p)^{b-1} \quad (4.5)$$

When a prior distribution of the Beta family is updated by use of Bayes' theorem and the Binomial distribution is used to calculate $P(E|p)$, the parameters of the posterior distribution can be calculated according to equations (4.6) and (4.7), where a' and b' are the parameters of the prior distribution and a'' and b'' the respective parameters of the posterior distribution [71]. The number of successful trials is denoted as r and the total number of trials that were observed as n .

$$a'' = a' + r \quad (4.6)$$

$$b'' = b' + n - r \quad (4.7)$$

More than two possible events

Another situation that one can be called upon to deal with is one in which there are more than two possible, mutually exclusive events. One might wish, for example, to describe a situation in which a fire might occur in any one of several different areas in a building, given that a fire has occurred in the building. This can be described by a probability distribution defined over all of the possible events of the type "the fire

will occur in area 1”, “the fire will occur in area 2”, etc. To calculate the probability that, say three fires occurred in area 1, one fire in area 2 and two fires in a third area, one can make use of the multinomial distribution. In such a case, the multinomial distribution can be used to calculate the probability that the evidence in question would have been observed given a specific probability distribution over the different areas. The multinomial distribution can be written as equation (4.8), where r_i is the number of events of type i (for example “the fire occurred in area i ”), p_i is the probability of an event of type i , n is the total number of events ($n = r_1 + \dots + r_{k+1}$) and $k+1$ is the number of different types of events.

$$p(r_1, \dots, r_{k+1}) = \frac{n!}{r_1! \dots r_k! (n - r_1 - \dots - r_k)!} p_1^{r_1} \dots p_k^{r_k} (1 - p_1 - \dots - p_k)^{(n - r_1 - \dots - r_k)} \quad (4.8)$$

When a multinomial distribution is employed to calculate the probability of observing the evidence, the Dirichlet distribution is a conjugate distribution (see Appendix A). The Dirichlet distribution can be written as equation (4.9), where $\nu_1, \nu_2, \dots, \nu_{k+1}$ are the parameters of the distribution and $p_1, p_2, \dots, p_k, (1 - p_1 - \dots - p_k)$ are the different probabilities of event 1, event 2, etc.

$$f(p_1, \dots, p_k) = \frac{\Gamma(\nu_1 + \dots + \nu_{k+1})}{\Gamma(\nu_1) \dots \Gamma(\nu_{k+1})} p_1^{\nu_1 - 1} \dots p_k^{\nu_k - 1} (1 - p_1 - \dots - p_k)^{\nu_{k+1} - 1} \quad (4.9)$$

If the decision maker’s belief regarding one of these probabilities is of interest one can study the marginal distribution of that particular probability. For p_i , the marginal distribution is a Dirichlet distribution with two parameters, ν_i and ν_A , where ν_A is equal to the sum of all the other ν s except for the ν_i [84]. A Dirichlet distribution with only two parameters is the same as a Beta distribution and thus the marginal distribution for a specific p_i is a Beta distribution with the parameters $a = \nu_i$ and $b = \nu_1 + \dots + \nu_{k+1} - \nu_i$ (see equation (4.5)).

Updating the Dirichlet distribution follows the same simple pattern as for the Beta distribution, the posterior parameters being given by equation (4.10), where $\nu'_1, \nu'_2, \dots, \nu'_{k+1}$ are the prior parameters of the Dirichlet distribution, $\nu''_1, \nu''_2, \dots, \nu''_{k+1}$ are the posterior parameters and r_1, r_2, \dots, r_k ,

r_{k+1} , are the number of events of the different types that occurred (see Appendix A).

$$\begin{aligned} v_1'' &= v_1' + r_1 \\ v_2'' &= v_2' + r_2 \\ &\vdots \\ v_{k+1}'' &= v_{k+1}' + r_{k+1} \end{aligned} \quad (4.10)$$

Assume that the decision maker's prior belief regarding the probability that a fire will occur in one of three different areas in a particular building can be represented by a Dirichlet distribution with the parameters $v_1 = 2$, $v_2 = 5$ and $v_3 = 10$. The parameters p_1 , p_2 and p_3 correspond to the probability with which a fire is assumed to occur in areas 1, 2 and 3, respectively. The marginal distribution of the different frequencies is shown in Figure 19.




Figure 19 Marginal distributions for the parameters p_1 , p_2 and p_3 of the prior Dirichlet distribution.

Assume that the decision maker learns that there have been 4 fires in area 1, which is the area where a fire was least expected to occur. One can update the decision maker's belief regarding the probability of fires in the different areas then by using the rather simple equation (4.10). This results in a posterior Dirichlet distribution which has the parameters $v_1 = 6$, $v_2 = 5$ and $v_3 = 10$. The marginal distributions of the probabilities p_1 , p_2 and p_3 are shown in Figure 20.




Figure 20 Marginal distributions for the parameters p_1 , p_2 and p_3 of the posterior Dirichlet distribution.

As expected, the four fires that occurred in area 1 led to the posterior marginal distribution for the probability of a fire in that area being adjusted in the direction toward a value of 1. In the posterior distribution, the mean value of the marginal distribution of p_1 is even higher than the mean value of the marginal distribution of p_2 .

Continual updating of fire risk measurement

Now that an overview of a number of important Bayesian methods useful in the present context has been presented it is important to consider briefly how they can be used in practical terms in the present context.

All the Bayesian methods taken up here can be used to combine subjective judgement with measurements conducted in or applying to the building of interest. For example, one could use a general investigation of the fire frequencies in buildings of different types (for example [85]) as a point of departure for creating a prior distribution of the frequency of fires in a specific building, after which that prior distribution can be updated using information concerning how many fires have actually occurred in the building in question. Bayesian updating of this type is easy to perform by use of either conjugate distributions such as discussed

in previous sections or discrete distributions to represent the uncertainty regarding the parameters of interest.

Of greater concern here, however, is the measurement of changes in fire risk by use of the Bayesian methods that have been considered. In [18] (Paper 4) a method is presented involving Bayesian networks being combined with the decision analysis framework presented in the thesis. The key ideas in that paper are to use decision analysis to provide the basis for a measure of fire risk that can be useful in factories and to show how the measure can be calculated by use of Bayesian networks. Since the measure of fire risk can be updated through the use of Bayesian networks when new information is obtained, it constitutes a way of measuring fire risk continually.

Since in many factories fires seldom occur, the information of actual fires used to update the Bayesian network may need to be complemented by information from experts, such as representatives of the fire department, and the like. A model for how expert judgement can be incorporated into the use of Bayesian methods is suggested in the paper. The method is based on earlier ideas of Apostolakis and Mosleh [82], who suggested that expert estimates should be treated as if the expert actually had observed the phenomenon in question. Thus, if an expert estimates the probability that a sprinkler system will succeed in extinguishing a fire to be 0.9, this is interpreted as if he/she actually had observed a number of fires, 90% of which were extinguished by the sprinkler system.

A technique called fractional updating is used to update the Bayesian network. In using that method, uncertainty regarding a probability estimate is expressed by use of a fictitious sample size (s). The higher the fictitious sample size is the more certain the decision maker is concerning what the estimate of the probability in question should be. If the fictitious sample size is set to 10, for example, and the estimate of the probability that a sprinkler system would succeed in extinguishing a fire is 0.9, this is interpreted as if the sprinkler system had succeeded in extinguishing 9 out of 10 fires.

One issue not dealt with in the paper referred to above is how to show the uncertainty regarding a particular probability when the fractional updating method is employed. Consider a situation in which one is uncertain of whether an event will occur or not, such as whether a

sprinkler system will succeed in extinguishing a fire, for example. Assume that the uncertainty regarding the probability of the event in question can be represented by a non-informative prior distribution from the Beta-family. A Beta prior distribution with the parameters $a = 0$ and $b = 0$ is such a distribution [86]. When new evidence is obtained, the parameters of the Beta-distribution are updated in accordance with equations (4.6) and (4.7). Thus, if one observed a sample of size s and in r cases the event in question occurred, the posterior distribution would be a Beta-distribution with the parameters $a = r$ and $b = s - r$. Since the mean value of a Beta distribution is $a/(a + b)$, the mean value of the posterior distribution is r/s , which is the value used to represent the probability in question in the fractional updating method [87]. Thus, one can view the fractional updating method as representing the Bayesian updating of a non-informative Beta prior distribution in which the mean value of the posterior distribution is used as a Bayesian estimator. If there are more than two possible events that can occur, one can use a Dirichlet distribution as the prior distribution and update it when new information is received.

Thus, the conjugate distributions considered in this chapter can be seen as being connected with the Fractional updating method employed in [18] (Paper 4). Whereas in using the Fractional updating method one does not need to indicate the uncertainty one has regarding the probability in question using a probability distribution, it is important in using the extended decision analysis method employed in the thesis, that the decision maker assess or be provided knowledge concerning uncertainty of this sort. The connection between the Fractional updating method and the conjugate distributions discussed here is thus important when that method is employed. For further information concerning this method for the continuous measurement of fire risk, see [18] (Paper 4), which is included in the thesis.

5 Case studies

Two case studies using the methods presented in the thesis are included here. The first case concerns the company Asea Brown Boveri (ABB) and the second concerns Avesta Sheffield⁵. Only analyses employing extended decision analysis are presented in this chapter. In Paper 3 [48], a small part of the case studies are analysed by use of Supersoft decision theory and of traditional decision theory, so as to compare the use of these methods. Both case studies, performed in 1998 when the analyst (the author) spent 2 weeks in each factory making interviews, visual inspections, and the like, involved analysing investments in water sprinkler systems. Help in conducting the analyses was provided above all by Ingemar Grahn⁶ and Olle Österholm⁷ from Avesta Sheffield and Bo Sidmar⁸ and Mikael Zeeck⁹ from ABB. In referring to “people from ABB” and “people from Avesta Sheffield”, these are the persons referred to.

In performing a decision analysis of this sort for a factory, one needs to be able to estimate the probabilities of the fire scenarios taken into account. Since the focus in this thesis is on decision analysis, no attempt is made to develop advanced risk analysis methods, a simple risk model based on event-trees being employed in both case studies. Note that the decision analysis methods described in the thesis could be utilised in combination with any fire risk analysis model used for estimating the probabilities of different fire scenarios in the event of fire.

The results of the case studies presented in the chapter differ somewhat from those presented in [38] and [64] (Paper 1 and 2). This is due to improvements in the risk model involved having been made. More precisely, a different procedure for screening the uncertain parameters was used in both the ABB and the Avesta Sheffield case as presented here than employed initially [38] and [64] (Paper 1 and 2). Also, whereas

⁵At the present time Avesta Sheffield has merged with the stainless steel division within the company Outokumpu and is called Avesta Polarit.

⁶Ingemar Grahn is the risk manager of Avesta Sheffield.

⁷Olle Österholm is an engineer working at the cold-rolling mill in Nyby.

⁸Bo Sidmar is the risk manager of ABB Sweden.

⁹Mikael Zeeck is responsible for, among other things, the fire safety in ABB:s building 358.

in the initial Avesta Sheffield analysis the consequences when the sprinkler system extinguishes a fire were assigned a value of 0, this was changed in the present analysis to values varying (depending on the area in which the fire occurs) between 10 thousand SEK and 150 thousand SEK. The screening process employed in the original analysis involved investigating how much the uncertainty regarding a specific parameter could change the CE of a *specific* alternative. Instead of investigating the change in the CE of a specific alternative, the change in the *difference* in CE between the decision alternative of investing in a sprinkler system and that of not investing in it is employed in the screening process here.

The chapter starts with a description of the risk analysis model employed. The case studies are then taken up. The chapter concludes with a brief discussion of the practical applicability of the decision analysis methods presented in the thesis.

The fire risk analysis model

The goal of using the fire risk model is to estimate the probability of different fire scenarios, given that a fire has occurred. In doing this, a number of events which, depending on whether they occur or not, can affect the fire scenario that occurs will be investigated. Event trees are used to visualise the model. It is assumed that whether the events in question will occur during a fire cannot be determined beforehand. The probabilities of these events are thus treated as representing aleatory uncertainty.

The events involved concern a set of “systems” intended to limit or extinguish any fire that occurs in a building. The following systems of this sort are considered here:

Active systems. These are systems designed to actively extinguish a fire, such as water sprinkler systems, CO₂-systems, or light-water systems. In estimating the conditional probability that an active system will extinguish a fire (i.e. conditional on all the preceding events in the event tree), one can use as a point of departure any investigations that may be available concerning how reliable the system is. One should remember, however, that the numbers an investigation provides are estimates for a whole group of systems and that the reliability of the system in question can differ from this. One should best use any value obtained in an

investigation of this sort as a starting point in attempting to estimate the reliability of the *specific* system at hand. One could use the results of investigations generally to create a prior distribution pertaining to the reliability of the system and then use statistics obtained for the specific system considered to perform a Bayesian updating of the system's reliability (see the previous chapter).

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

Fire department. Since a fire department can affect the outcome of a fire, its usefulness in this respect can be represented by the conditional probability that it will succeed in extinguishing a fire. This probability could be estimated in collaboration with representatives of the fire department in question. It would probably be estimated in terms of a rather large probability interval or broad probability distribution (representing epistemic uncertainty). Särndqvist [88] and Tillander and Keski-Rahkonen [89] has provided information on the performance of fire departments in manual fire fighting operations, information that could be useful in making such an estimate.

Fire growth potential. After ignition in the first fuel package involved, a fire may continue to grow, so that further fuel packages are involved, remain steady and consume all the fuel available, or go out directly due to the conditions no longer being sufficient to sustain combustion. In the risk model employed the growth potential of a fire is modelled by the probability either that the fire will be small, i.e. have only a negligible effect on the company, that it will have the potential to spread to a moderate extent, or that it will have the potential to spread to a large extent. Note that what is of interest is a fire's *potential*. If a fire had the potential for achieving a large spread, it could sooner or later, if left unattended, destroy or seriously damage the entire fire compartment

where it started, although the spread of fire might nevertheless be very limited if extinguishing operations were successful.

Employees. If employees detect a fire and have the appropriate equipment, they may succeed in extinguishing a fire before it grows to any significant size. The probability of employees extinguishing a fire would be expected to depend on such factors as their training, the amount of fire fighting equipment they have access to, and the like.

The events that are of interest in the risk model and pertain to the different areas described above are modelled in an event tree. An illustration of the event trees used in the case studies is shown in Figure 21, a full description of those used in the two case studies being found in Appendix D and E. In the event tree, shown in the figure, the initiating event is “Fire has occurred”. The first uncertain event then concerns the question of which fire compartment the fire has occurred in. In estimating the probability of fire for each of the fire compartments, given that a fire has occurred somewhere in the building, information regarding the different fire compartments needs to be taken into account. For example, both the size of a fire compartment and the activities performed there can be expected to affect the probability of a fire occurring in the fire compartment. The next uncertainty in the event tree concerns the growth potential of the fire. This is conditional upon the particular area in which the fire has occurred. One could conceive of the fire growth potential as being in some one of several “states”, such as large, medium or small. A large growth potential of the fire could be taken to mean that if the fire is unattended it will at least destroy the fire compartment in which it began. An example of a fire having a large fire growth potential is that of a fire starting in a storage rack containing a large amount of combustibles. A fire with a medium fire growth potential might be said to be one that has the potential to destroy large parts of the fire compartment of concern, such as a particular machine located there, or its causing substantial damage to equipment sensitive to smoke, for example. A fire with a small growth potential can be conceived as one that would not cause any significant damage to the company’s facilities.

The next uncertain event in the model concerns whether the automatic detection system, if one is present, will detect the fire. The states of the detection system are assumed to be “Working” and “Not working”. “Working” means that the system has detected the fire and has notified the appropriate personnel.

It is assumed that the employees can possibly extinguish a fire using the manual fire equipment present in the building. The possible events considered are “The employees will succeed in extinguishing the fire” and “The employees will not succeed in extinguishing the fire”, the terms “Extinguish” and “Not extinguish”, respectively, being employed here. Whether the employees succeed in extinguishing the fire is dependent upon the fire compartment the fire has occurred and whether or not the fire detection system is working.

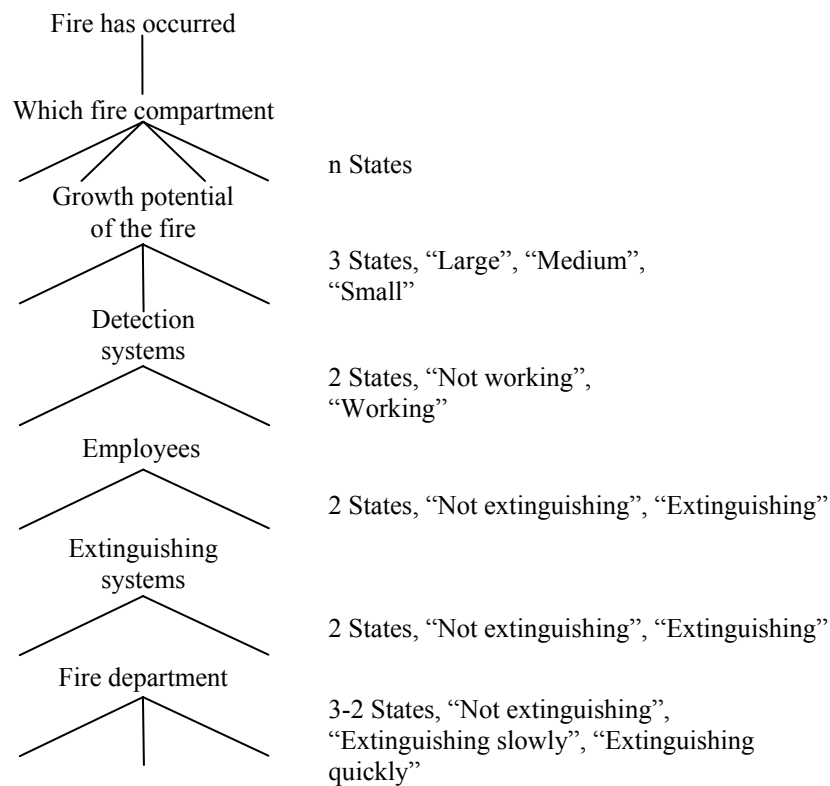


Figure 21 Illustration of an event tree model used to calculate the probability of different fire scenarios.

If the building is equipped with a fire extinguishing system, for example a sprinkler system, the possibility this has of extinguishing the fire needs to be taken account of in the event tree. The uncertainty of whether an

extinguishing system will succeed in extinguishing a fire is represented in the figure by the states “Extinguishing” and “Not extinguishing”.

The last uncertain event contained in the event tree concerns whether the fire department will succeed in extinguishing the fire. In the ABB analysis, the fire department was represented by three states: “Not extinguishing”, “Extinguishing slowly”, and “Extinguishing quickly”. The difference between “Extinguishing slowly” and “Extinguishing quickly” could be viewed as being that of whether successful extinguishing operations can be launched by the fire department personnel that arrive at the factory first or whether they have to wait for additional forces before they can successfully extinguish the fire. In the Avesta Sheffield analysis, the state of the fire department’s operations can be either “Extinguishing” or “Not extinguishing”. The reason for the states in the two case studies differing is that in the ABB case whether the fire was extinguished quickly or slowly was judged to have a strong effect on the damage costs, whereas in the Avesta Sheffield case the difference between a slowly extinguished fire and a quickly extinguished fire was not judged to be critical for damage costs. This is due to the difference between the two factories in the equipment present. In the ABB building there was a lot of electronic equipment, which is sensitive to smoke. Quick extinguishing operations could thus prevent the equipment from being exposed to large quantities of it and thus limit the damage considerably. In the Avesta Sheffield case, on the other hand, the equipment was not very sensitive either to smoke or to heat, as indicated by earlier fires.

The event tree shown in Figure 21 can be used to calculate the probability of a fire growing so as to involve the entire fire compartment. It is desirable to also consider, however, what happens after a fire has grown to that extent. It is important in doing this to take the effectiveness of the fire compartment into account. In estimating the probabilities of different fire scenarios, given that a fire has spread so as to involve the entire fire compartment, use is made of a simple probabilistic model. Each barrier between adjacent fire compartments is assigned a probability value, representing how likely it is that the barrier in question would prevent a fully developed fire from spreading from one side of the barrier to the other. Since not much information is available regarding the probability that a fire compartment barrier will succeed in limiting the spread of fire beyond it, estimating such a probability involves large uncertainties. In estimating the probabilities of different extents of fire

spread, given that a fire has spread so as to fully involve a particular fire compartment, use is made of a simple computer program, one which calculates the probability of having a particular combination of working and not working fire compartment barriers. The program determines which fire compartments can be expected to be destroyed for each combination of working and non-working barriers. The computer program, written in MATLAB (Version 6.1.0.450 Release 12.1), is included in Appendix F.

As far as the author is aware, not much work has been done to determine dependencies between variables pertinent to fire risk analysis. For example, the probability that a particular water sprinkler system will operate may not be completely independent of the probability that the smoke detection system will work. It may be that both the reliability of the sprinkler system and the reliability of the smoke detection system are dependent upon maintenance and that accordingly they are not independent. The only dependencies considered in the case studies presented here are dependencies between the probability of various events and the event of fire occurring in a specific fire compartment. Also, whether the smoke detection system detects a fire or not affects both the probability that the employees will extinguish the fire and that the fire department will do so. Another type of dependency between variables not dealt with in the case studies considered here concerns the epistemic uncertainty regarding the probability values. Thus no account is taken, for example, of whether the value of the probability that the employees will succeed in extinguishing a fire in a particular area is affected by the epistemic uncertainty regarding the value of the probability that they will succeed in extinguishing a fire in some other area. This assumption is particularly important when extended decision analysis is used, since in performing such an analysis the epistemic uncertainty regarding the different probability values is expressed in terms of probability distributions. If a dependency between the different probabilities exists, the probability distributions that represent, as is the case in the model used, the epistemic uncertainty regarding the probability values in question should not be modelled as being independent. The effect of assuming probabilities to be independent when in fact there is a dependency between them is that the final distribution, which in the present context represents the difference in certainty equivalent (CE) between the two decision alternatives, *may not be conservative with respect to the robustness* of the decision. This means that the dependencies can cause the entire distribution of the CE

to be shifted in any direction. The issue of dependencies between probabilities in an extended decision analysis model is a topic that should be of considerable interest in future research. It is taken up in chapter 6.

Subjective probabilities

Note that according to Bayesian decision theory, probabilities are perceived as subjective and are defined with respect to the decision maker's choices between uncertain situations (see chapter 2, page 18-19). It is recognised that this interpretation may very well lead to persons having difficulties in accepting the results of a decision analysis. This should not pose a problem, however, as long as the *decision maker* accepts a subjective interpretation of probability. After all, the analysis is carried out for the decision maker and for no one else. This means that the results of a decision analysis are not necessarily valid for anyone other than the decision maker. From this, it follows that a decision analysis is never completely objective. Accordingly, even if a specific decision alternative has been found to be the best for one decision maker to select, some other decision alternative may be best for another decision maker.

In practice, a decision maker is unlikely to possess all the knowledge required to perform a decision analysis, his/her having to rely on other persons' estimates and base his/her own estimates on estimates received from others. Probably the easiest way of doing this in practice is to have experts provide their estimates of different probabilities, the decision maker then either accepting these estimates as his/her own or adjusting them in a way he/she considers reasonable. Estimates of (subjective) probabilities and utilities have been studied extensively, a number of different methods for encoding subjective judgement having been suggested (see [68], for example, for a review of different methods for the encoding of probabilities and [63] for a review of different methods for the encoding of utilities). In the thesis, no attempt has been made to develop encoding methods specifically applicable to the present context.

For probability encoding, use has been made of a simple method called "direct response" [68]. This involves probabilities being assessed directly as numbers between 0 and 1 and the uncertainty of the assessments being characterised by a minimum, a maximum and a most

likely value being assessed. In the literature on probability encoding techniques [68], for example) it is often assumed that there is an analyst who elicits probabilities from an expert. In the case studies reported on in the thesis, there was no single expert providing estimates, the analyst (the author) and people from the companies working instead as a group to come up with the estimates needed.

Although it is desirable to have access to “objective” information about the building in question, such as fire statistics pertaining to it or other statistics of relevance, such as the results of investigations of the reliability of water sprinkler systems generally, not much in the way of fire statistics or the like relevant in the present context is usually available. In particular, there is a lack of information regarding the probability either that the employees in a particular building or that the fire department will succeed in extinguishing a fire. Under such conditions, one needs to rely on more “indirect” information, such as the quality of the manual fire extinguishing equipment, the number of people in the building, the amount of combustibles there, and the like. On the basis of this indirect information, one can make estimates of the probabilities in question. However, since one is very likely to feel uncertain regarding them, it is helpful to be able to use either the extended decision analysis method or the Supersoft decision analysis method described above.

Although general information regarding such probabilities as that of the employees succeeding in extinguishing a fire tends to be scarce, there are other probabilities regarding which a considerable amount of information exists, such as concerning the reliability of water sprinkler systems, for example. A summary of the general information regarding different probabilities available in the case studies is presented below.

Automatic water sprinkler system

Most investigations report sprinkler systems to have a high level of reliability, typically one of more than 0.9. The references which provided help in estimating the probability that the sprinkler system in question would succeed in extinguishing a fire are indicated in Table 5. There, one can see, along with the references, the country in which each investigation was performed and the level of reliability found.

Table 5 *References concerning estimates of the reliability of water sprinkler systems.*

Reference	Country	Reliability
Rutstein and Gilbert [90]	Great Britain	0.95
Young [91]	Great Britain	0.985
Stirland [92]	Great Britain	0.95
Marryatt [93]	Australia	0.99
Maybee [94]	USA	0.985

Sui and Apostolakis [72] provide a highly detailed account of an investigation of the reliability of sprinklers in which they relate the reliability of sprinklers to the type of industrial plant in which they are found. The reliability varies between about 0.96 and 0.99, depending on which type of plant is involved. Two industrial areas that appeared relevant to the case studies were those of “Metal products” and “Miscellaneous”, in which the reliability of the sprinkler systems was adjudged to be about 0.99 and 0.98, respectively.

Since there is no reason to believe that the reliability of the sprinkler systems considered in the case studies should differ substantially from that reported in the references cited, estimates of the probability that a sprinkler system will succeed in extinguishing a fire should be high.

Smoke detection systems

Investigations of the reliability of smoke detection systems appear to not be as common as those concerning the reliability of sprinkler systems. The sources of information involved are [95], [96]. The BSI guide [95] cites a value of 0.1 for the probability that a smoke or heat detector will fail to detect a fire. In [96], this probability is reported to be between 0.26 and 0.05, depending upon the detection system.

Fire compartmentation

There appear to not be many investigations pertaining to the probability that a specific fire compartment will succeed in limiting the spread of fire. In Fire engineering guidelines [96], this probability is reported as depending upon whether or not a flashover has occurred in the compartment in question and the fire resistance of the construction involved. Here, the probability that a particular construction will fail to limit the spread of fire, given that a flashover has occurred, is of interest. If a wall has no documented fire rating but has no openings in it (such as

doors), the probability is estimated to be about 0.5 [96]. If the wall is fire-rated and contains no openings, the probability is estimated to be about 0.95 [96]. If the wall has no documented fire rating and has openings without automatic shutters the probability is estimated to be only about 0.3 [96]. If the wall has a fire rating and has openings in which there are automatic shutters, the probability is estimated to be about 0.9 [96].

The information concerning general investigations relating to the probabilities referred to above was used as a point of departure when estimates were made of probabilities pertaining to the specific buildings involved. Note that for some of the probabilities no general information was available, estimates of these being based solely on visual inspection and on judgements by the analyst and by the personnel from the factory in question.

Screening procedure

To perform extended decision analysis for the two case studies, it was necessary to have a screening procedure to determine which parameters, i.e. probabilities and consequences, should be represented as probability distributions in the extended decision analysis. The reason for needing to employ such a screening process was a practical one. Each of the two case studies involved over 100 parameters that can be considered to be of varying uncertainty. Since the work involved in assessing probability distributions for each one of the parameters and performing simulations using these distributions was judged to be too great to make it practical, a screening procedure was adopted so as to be able to identify the parameters having only limited influence on the overall epistemic uncertainty, i.e. the spread of the resulting distribution of the difference between the alternatives in terms of CE. A maximum, a minimum and a most likely value were estimated for each of the parameters. In the screening process, the difference in CE between the decision alternatives is calculated using the values that appear most likely, or what in chapter 3 is termed traditional decision analysis. Each of the parameters is then changed, one at a time, from its most likely to its maximum and minimum value, respectively. The change in CE connected with this is noted. The parameters that result in the least change in the difference between the decision alternatives in terms of CE are those then that are

not dealt with as being uncertain, i.e. are not represented by distributions in extended decision analysis.

In performing the screening analysis of probabilities that could affect each other, such as the probabilities of different fire potentials, a maximum and a minimum value were estimated for each probability. In determining the effect of changing one of the parameters, each parameter was changed from its maximum to its minimum value while the relationship of the remaining parameters to each other was held constant. Thus, if there are three probabilities (p_1 , p_2 and p_3) that are required to sum to 1 ($p_1 + p_2 + p_3 = 1$), the screening procedure would involve setting p_1 for example, at its highest value and adjusting p_2 and p_3 so that the sum of the three would still be 1 and that the ratio of p_2 to p_3 remained unchanged.

The ABB analysis

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 manufacturing plants and electric power plants. The company, which has a turnover of approximately 2.4 billion SEK, has about 1400 employees in the Västerås and the Malmö region in Sweden¹⁰.

The present analysis deals with the investment in a water sprinkler system for a building called building 358. In that building, 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 represent a highly important segment of the ABB group, for reasons such as their providing other companies within the group with circuit cards.

The building is situated in an industrial area in Västerås. It is approximately 55000 m² in size and is divided up into eleven different fire compartments. The nearest fire department, in the city of Västerås, needs 6 to 10 minutes of driving time to reach the building. The building

¹⁰ These numbers were valid in 1998 when the analysis was carried out. Since then, a large part of ABB's activities in the building have been sold to the company Flextronics.

is equipped with a smoke detection system with a communication link to the fire department. There is presently a water sprinkler system for the entire building, although in the middle of the nineties this was not the case. Since the activities currently being carried out in the building are similar to what they were then, the present analysis will be for the building *without* its having a sprinkler system, so as to determine whether an analysis by means of this method, if carried out in the mid-nineties would have shown the sprinkler system to be a good investment. The decision alternatives available are (a_1) investing in a sprinkler system, (a_2) not investing in a sprinkler system.

The costs of the sprinkler system amount to approximately 10 million SEK, maintenance being estimated to cost some 0.1 million SEK per year. These are the only economic matters not related to fire that are considered here. It was decided to use an r -value (see equation (3.10)) of 0.15 to represent the decision maker's (ABB's) preferences for fires occurring at different times.

The uncompensated losses associated with each fire scenario, i.e. the monetary value the decision maker regards as being equal to the fire scenario in question, were estimated by personnel from ABB. No thorough investigation of these losses was carried out, such as investigating market shares lost due to the business interruption following a fire, and the like. Instead, use was made of the monetary amount that the insurance company would have to pay ABB in case of each of the fire scenarios. The idea was to relate the uncompensated losses to this value by assuming them to be equal to the insured losses, as proposed in [97]. Only those losses to ABB pertaining to ABB Automation Products were estimated. No account was taken of the effect that a business interruption in building 358 might have on other companies within the ABB group. The estimated losses in the case of a fire destroying a whole area of the building are presented in Table 6. There, "Direct losses" represents the value of the equipment destroyed in that area and "Consequential losses" the monetary value associated with the business interruption following a fire in that area. Figure 22 shows the relative locations of the different areas within the building. The losses associated with various less severe fire scenarios are presented in Appendix B.

Table 6 *Losses associated with the destruction of a particular area in building 358. The areas referred to are shown in Figure 22. Losses are measured in millions of SEK. The losses presented in the table are abbreviated using a C followed by their respective area number and "ABB", C2ABB, for example, representing the sum of the consequential losses and the direct losses associated with the destruction of area 2.*

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




Figure 22 *The relative location of the different areas within building 358. The areas without numbers are not used by ABB Automation Products.*

In calculating the certainty equivalents of the decision alternatives, a model for the development of a fire is needed. The general structure of such a model was discussed in the previous section. The event tree used to represent the uncertainty regarding which fire scenario occurs if a fire should start in building 358 is too large to be presented with ease here, although it is shown in Appendix D. The tree includes more than 100 different fire scenarios, these leading to frequently varying degrees of fire spread (though sometimes the same) within any given fire compartment. Also, if a fire has spread to involve an entire fire compartment, a number of additional fire scenarios representing the destruction of differing combinations of fire compartments are possible. These fire scenarios are not included among the 100 scenarios mentioned above. Estimates of the probabilities of the differing extents of fire spread between the various areas were obtained by use of a computer program described in Appendix F.

In making estimates of the values of the probabilities used in the model of fire spread in the building, use was made of information available regarding the reliability of fire safety systems generally (presented earlier in this chapter) and regarding past fires in the building, along with estimates made by ABB personnel, by the fire department of Västerås and by the analyst (the author). Since the estimates of the parameters were considered uncertain, a maximum, a most likely and a minimum value were estimated for each parameter in the model. These estimated values are presented in Appendix B. The frequency of fire is the only parameter to which no maximum, most likely and a minimum value were assigned. Instead, the frequency of fire was estimated on the basis of previous investigations. Using the relationship between floor area and frequency of fire given in [85] indicates the frequency of fire in the ABB building to be 0.38, 0.55 or 0.74 fires per year, depending on whether the industrial group “Electrical engineering”, “All manufacturing industry” or “Other manufacturing” was involved. On the basis of the relationship given in [98] (the ignition frequency is suggested to be 10^{-5} per m^2 for buildings above 1000 m^2 floor area), the frequency of fire in the ABB building is 0.55 fires per year. Using these general estimates as a point of departure, the frequency of fire in building 358 was estimated to be somewhere between 0.3 fires and 1.25 fires per year. The uncertainty regarding the frequency of fire was represented by a Gamma distribution with the parameters $\alpha = 15$ and $\beta = 20.8$. This prior distribution was updated using information on how many fires had occurred in the

building. During 1996, 1997 and 1998 there were, in total, 5 fires in the building. On the basis of this information, one can update the prior distribution so as to arrive at a posterior distribution. The process of updating a Gamma distribution is described in chapter 4. The resulting posterior distribution is a Gamma distribution with the parameters $\alpha = 20$ and $\beta = 23.8$. This distribution, shown in Figure 23, is used in the analysis to represent the uncertainty regarding the frequency of fire. There it is referred to as parameter P3ABB.




Figure 23 *Gamma distribution representing the uncertainty regarding the frequency of fire in building 358.*

A screening process was employed so as to reduce the complexity of the analysis. The aim of the screening process was to determine which parameters had only a marginal effect on the epistemic uncertainty regarding the difference between the two alternatives in terms of CE in the extended decision analysis model, parameters which can thus be treated as having exact values. In the screening process, each of the parameters was adjusted, one at a time, from its minimum to its maximum value. The effect of this change on the difference between the decision alternatives in terms of CE was noted, the parameters being ranked according to the size of the change involved. Note that in changing parameters that affect other parameters (for example when changing the probability of a large fire), the remaining parameters (the probability of a medium fire and that of a small fire) are assumed to have values such that their ratio remains constant. In analysing the impact of changing the frequency of fire in the screening process, its maximum

value is assumed to be 1.35 fires per year and its minimum value 0.45 fires per year. Using these two values as the maximum and minimum value for the frequency of fire is an approximation since the frequency of fire is represented by a Gamma-distribution, which has 0 as its lowest value. A fire frequency of between 0.45 and 1.35 fires per year represents approximately a 96% confidence interval for the parameter of interest.

The results of the screening process are presented in terms of a tornado diagram in which the change in the difference in CE for each of the parameters being studied is shown. The parameters leading to the largest change in the difference between the two decision alternatives in terms of CE are located at the top of the diagram and those leading to the smallest change at the bottom. Since many parameters (over 100) are involved in the analysis, it is impractical to present the effects of all of these in a single diagram. Instead, only those able to produce a total change of more than 1% are presented in Figure 24. In the diagram, one can see that it is the uncertainty regarding the frequency of fire that has the strongest influence on the difference in CE between the decision alternatives.

In the extended decision analysis, only the parameters resulting in a total change in the difference in CE of more than 20% (counting the sum of the decrease and the increase in the difference in CE) are treated as being uncertain. The reason for this is practical, in that it would be very cumbersome in an extended decision analysis to treat all the parameters (more than 100) as probability distributions and that the result would hardly be worth the effort. Choosing a limit of 20% for the parameters treated as uncertain in the extended decision analysis results in 9 parameters being treated as uncertain there. These parameters are the following: the frequency of fire (*P3ABB*); the probability of a fire in area 2 being of small potential (*P37ABB*); the probability of a fire in area 2 being of large potential (*P39ABB*); the probability of a fire in area 6 being of large potential (*P83ABB*); the probability that the employees would succeed in extinguishing a fire in area 2, given that the smoke alarm had been activated (*P40ABB*); the probability that the fire department would extinguish a fire in area 2 quickly, given that the potential of a fire there is large and that the smoke alarm had been activated (*P44ABB*); the probability that the employees would succeed in extinguishing a fire in area 6 given that the smoke alarm had been activated (*P84ABB*); the probability of a fire in area 8 being of large

potential ($P105ABB$), and the probability of a fire in area 6 being of small potential ($P81ABB$). Note that since the parameters $P37ABB$ and $P39ABB$ affect each other if one is changed, the other also change only the one that has the stronger effect on the difference in CE is taken into account in the extended decision analysis.

In the extended decision analysis the parameters presented above were represented by triangular distributions. Their maximum, minimum and most likely values are presented in Table 7. The frequency of fire is not presented there since it is not represented by a triangular probability distribution but by a gamma distribution. That distribution is shown in Figure 23.

Table 7 *The maximum, most likely and minimum values for the parameters considered uncertain in the extended decision analysis.*

Parameter	Max	Most likely	Min
The probability of a fire in area 2 being of small potential ($P37ABB$).	0.9	0.8	0.7
The probability of a fire in area 6 being of large potential ($P83ABB$).	0.15	0.1	0.01
The probability that the employees would succeed in extinguishing a fire in area 2, given that the smoke alarm had been activated ($P40ABB$).	0.8	0.6	0.5
The probability that the fire department would extinguish a fire in area 2 quickly, given that the potential of the fire was large and that the smoke alarm had been activated ($P44ABB$).	0.6	0.3	0.2
The probability that the employees would succeed in extinguishing a fire in area 6, given that the smoke alarm had been activated ($P84ABB$).	0.8	0.6	0.4
The probability of a fire in area 8 being of large potential ($P105ABB$).	0.1	0.05	0.01




Figure 24 Tornado diagram showing the effect of epistemic uncertainty on the CE of the sprinkler alternative.

A total of 5000 Monte Carlo simulations were performed for determining the robustness of the decision situation in which the Risk-adjusted net present value was investigated, i.e. the difference in CE between the alternative of keeping the building in its present condition (a_2) and that of investing in a sprinkler system (a_1). The results of the simulations are shown in the histogram in Figure 25. The results there were calculated taking account of the benefits in terms of risk reduction achieved during a 5-year period. The decision maker's preferences for fires occurring at different times were modelled using an r -value (see equation (3.10)) of 0.15.




Figure 25 *Histogram showing the effect of the epistemic uncertainties on the risk-adjusted net present value of investing in a sprinkler system. A period of 5 years was set as the limit for how long the benefits of the risk reduction were to be taken into account.*

In the extended decision analysis, it was decided to treat the risk attitude of the decision maker (ABB) as being risk-neutral. The reason for this was that no investigation of risk-attitude was possible since such an investigation would have required an unreasonable degree of effort on the part of the senior managers at ABB and was outside the scope of the case study.

Figure 25 shows the results of the analysis in a very clear way. The mean value of the Monte Carlo simulations, as shown in the figure, is 8.2 million SEK. If one wishes to use only a single value to describe the

attractiveness of the investment, this is the one to use. However, as one can see in the figure, presenting the decision maker only one value could be misleading, since the decision maker would thus receive no information regarding the epistemic uncertainty of this value. Figure 25 indicates the epistemic uncertainty to be substantial. The robustness index of the analysis is 96%, which means that 96% of the simulated Risk-adjusted values are positive.

It can be useful to provide the decision maker a diagram showing the effect of changing the period of time for which the benefits of the risk reduction due to the investment in the sprinkler is taken into account. Such a diagram is shown in Figure 26. The dashed lines there represent a robustness index of 5% and 95%, respectively.




Figure 26 Diagram showing the effect of changing the number of years for which the benefits of the reduction in risk is taken into account in the analysis. The dashed lines represent a robustness index of 5% and 95%, respectively.

The figure shows that the risk-adjusted net present value is positive and that the decision is robust (the robustness index is greater than 95%) if the period of time for which the risk-reduction benefits of the investment are taken into account is longer than 5 years.

The Avesta Sheffield analysis

Avesta Sheffield¹¹, 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 group was 5.8 billion SEK.

The company has a cold-rolling mill in Nyby (Sweden) that produces approximately 160 thousand tons (figures from 1998 and 1999) of cold rolled steel per year. This constitutes a major part of Avesta Sheffield's annual steel production of approximately 1 million tons. The decision analysis concerns the possible investment in a sprinkler system for the entire cold-rolling mill, which is approximately 15000 m² in size. The investment costs of the sprinkler system were estimated to be 2.5 million SEK and the annual maintenance costs to be 50 thousand SEK. These costs were the only certain costs taken into account in the analysis. The decision alternatives are (a_1) to make an investment in a sprinkler system and (a_2) to keep the building in its present state. In evaluating fires occurring at different times, it was decided by Avesta Sheffield to use an r -value of 0.2 per year (see equation (3.10)).

The analysis here was conducted in the same way as the ABB analysis, that is, through estimating the CE of the decision alternative to invest in a sprinkler system and of the decision alternative to not invest in a sprinkler system.

The losses associated with each of the fire scenarios analysed are both the direct losses and the consequential losses for the *entire* Avesta Sheffield group. These losses were adjudged by personnel from Avesta Sheffield so as to adequately represent the uncompensated losses of a particular fire scenario. If a fire were to destroy the cold-rolling mill, the consequential losses for the other facilities owned by Avesta Sheffield would be substantial. The losses for those other facilities, as well as the consequential losses for the cold-rolling mill itself, need to be taken into account. In sum, the consequential losses for the group if the cold-rolling mill were destroyed would be approximately 1.1 billion SEK per year. The need of accounting for the negative effects of a fire such as that in the cold-rolling mill occurring in other companies within the Avesta

¹¹ In January 2001 Avesta Sheffield merged with Outokumpu Steel and formed a new company, Avesta Polarit.

Sheffield group makes the present analysis somewhat different from the ABB analysis. In the ABB analysis, only negative consequences pertaining to the building of concern were taken into account.

The production process in the cold-rolling mill can be divided into a number of segments, each of which can be treated as involving a separate machine. Both the indirect and the direct costs associated with the destruction of a particular machine are shown in Table 8. Since the production process within the cold-rolling mill is somewhat more complicated than that in the ABB building, the calculation of the total costs, given that a particular area is destroyed, is more complicated. In the Avesta Sheffield case, there are strong dependencies between the different areas, since a product may need to pass through several areas before it is finished. A schematic drawing of the cold-rolling mill is shown in Figure 27, and a drawing showing the flow of material through the different parts of the cold-rolling mill in Figure 28. The numbers shown in the latter figure indicate how large a part of the cold-rolling mill's total steel production passes through the part of the factory in question. All the material (steel) that enters the cold-rolling mill goes through a stage of processing that occurs in production line 60, which is located in area 4. After production line 60, the flow of material is divided into two flows of roughly equal size. One of these flows goes directly to the cutters, located in areas 1 and 2, and the other to area 3, where the steel is rolled up on coils. The material is then cold-rolled in either the old cold-rolling mill, 1, or the new cold-rolling mill, 2. From these cold-rolling mills the material continues on through production line 55, which is located in area 4. In the end, the material goes through one of the cutters. The smoothing roller and the abrasive-belt grinder are used to process 50% and 20%, respectively, of the total production that takes place in the building, *independent* of whether the material goes through cold-rolling mill 1 or 2.




Figure 27 Schematic drawing of the different areas in the cold-rolling mill.




Figure 28 Flow of material within the cold-rolling mill.

Table 8 *Losses associated with the destruction of different machines in the cold-rolling mill. The losses are given in millions of SEK.*

Machine	Direct losses			Consequential losses			
	Min	Most probable	Max	Min	Most probable	Max	
Area 1							
Smoothing roller	100	125	150	450	630	810	C1Av
Cutter 1	32	40	48	225	315	405	C2Av
Cutter 2	60	75	90	225	315	405	C3Av
Area 2							
Cutter 3	80	90	100	450	630	810	C5Av
Area 3							
Cold-rolling mill 1	162	180	198	90	126	162	C6Av
Cold-rolling mill 2	225	250	275	360	504	648	C7Av
Strip coiling machine	72	90	108	450	630	810	C8Av
Area 4							
<i>Production line 60</i>							
Uncoiling capstan, weld	80	100	120	900	1260	1620	C9Av
Cold-rolling mill	160	200	240	900	1260	1620	C10Av
Oven and cooler	40	50	60	540	810	1080	C11Av
Blaster	32	40	48	900	1260	1620	C12Av
Pickling machine	40	50	60	900	1260	1620	C13Av
Stretcher leveller	40	50	60	900	1260	1620	C14Av
Cutter and coiling capstan	80	100	120	900	1260	1620	C15Av
Other (switch room etc.)	88	110	132	900	1260	1620	C16Av
<i>Production line 55</i>							
Uncoiling capstan, weld	80	100	120	450	630	810	C17Av
Oven and cooler	56	70	84	270	405	540	C18Av
Pickling machine 1	56	70	84	450	630	810	C19Av
Pickling machine 2	56	70	84	450	630	810	C20Av
Cutter and coiling capstan	80	100	120	450	630	810	C21Av
Area 5							
Abrasive-belt grinder	120	150	180	180	252	324	C22Av
Oil room	0.5	1	1.5	0	0	0	C23Av
Area 6, Engine room 1	56	70	84	90	126	162	C24Av
Area 7, Machine shop 1	24	30	36	0	0	0	C25Av
Area 8, Engine room 2	80	100	120	360	504	648	C26Av
Area 9, Machine shop 1	24	30	36	0	0	0	C27Av
Oil room	5	7.5	15	0	0	0	C28Av
Pallet storage	0.05	0.1	0.15	0	0	0	C29Av

If an area of the cold-rolling mill is destroyed by fire, a certain share of the total production capacity is lost, as shown in Table 9. If some combination of areas is destroyed, however, it is not easy to determine the cold-rolling mill's remaining production capacity. If areas 1 and 2 are destroyed by a fire, for example, the remaining production capacity cannot be estimated through looking it up in Table 9, an analysis of the product flow needing to be performed instead. For obtaining estimates of the probabilities of different extents of fire spread, given that a fire has begun in a specific area and that it has spread so as to involve the entire area in question, use is made of a computer code. Calculation of the remaining production capacity, given a particular extent of fire spread, is included in that code. The computer code is presented in Appendix F.

Table 9 Share of the total production capacity lost if the area in question is destroyed by a fire.

Area	Share of the production capacity
1	50%
2	50%
3	50%
4	100%
5	20%
6	10%
7	-
8	40%
9	-

In calculating the total consequential losses of a fire, one needs to determine both the share of the production that is lost, and the time it would take to increase production capacity to 100% again. The time it would take to bring production capacity back to a 100% level after a serious fire depends on what equipment was destroyed by the fire. For many of the components in the cold-rolling mill, however, the time required is very long. The Avesta Sheffield personnel estimated the time until production could be back to normal after a serious fire to lie somewhere between 10 and 18 months for many of the machines. In examining the consequential losses shown in Table 8, one can note that if a fire should destroy area 4, stopping all production in the cold-rolling mill, the consequential losses would be extraordinary. The maximum consequential loss in this case would be 1620 million SEK and the maximum direct loss 1330 million SEK, which are to be compared with

the total turnover of the Avesta Sheffield *group* in 1999, which was 5800 million SEK. Note that Avesta Sheffield would be reimbursed by their insurer for a part of the consequential losses mentioned above. This is not of interest in the present analysis, however, since it is the *intrinsic* monetary value which is of concern, i.e. the uncompensated losses, in each of the fire scenarios. The uncompensated losses were adjudged by personnel from Avesta Sheffield to be equal to the sum of the direct losses and the consequential losses reported above. The uncompensated losses due to the less serious fire scenarios are presented in appendix C.

To perform a decision analysis concerning investment in a water sprinkler system for the cold-rolling mill, one needs to create a model for fire spread in the factory. The approach taken was similar to that used in the ABB building. Although the general model described in the section termed “The fire risk analysis model” was made use of, certain aspects of it were changed to take account of the specific circumstances present in the cold-rolling mill. For one thing, the distribution of fires within a given area was explicitly modelled here. This involves estimating the probability that a fire would start in a particular machine, given that a fire had started in the area in question. It also involves estimating the probability of a fire starting in a specific machine, given that a fire has started in the area concerned and in a machine. Event trees concerned with this can be found in Appendix E. The estimates of the probabilities included in the model were performed by the analyst (the author) with help of personnel from Avesta Sheffield. In estimating the probabilities mentioned above, for example, personnel working in the area of concern were asked where fires were most likely to start on the basis of their experience. Unfortunately, no conclusive record of where fires had begun and who had extinguished them were accessible. Thus, only estimates based on the experience of the group just referred to were used in the analysis. Records showing, however, that during a six-year period (1993-1999) a total of 60 fires occurred in the building. Although this might appear to be many, one should realise that processes in the cold-rolling mill involve both high temperatures and combustibles in the form of oil. Oil is, in fact, used to cool the steel during rolling, and that process causes many fires. Most of the fires are small, however. Nevertheless, there has been at least one major fire in the cold-rolling mill. That fire occurred in the abrasive belt grinder, which contains no automatic suppression system, and once both the hydraulic oil and the cooling oil were involved, the personnel were unable to extinguish the

fire, which destroyed most of area 5 and was very close to spreading to nearby areas.

Estimates of the probabilities contained in the model of fire spread are presented in Appendix C, the minimum, maximum and most likely value being included there.

The frequency of fires was estimated using statistics on how many fires had occurred in the cold-rolling mill. In creating a distribution to represent the uncertainty regarding these frequencies, use was made of the fact that in the building a total of 60 fires had occurred during a 6-year period. Assuming a non-informative prior gamma distribution and updating it by use of information on the number of fires that occurred in the building just referred to lead to the posterior distribution shown in Figure 18. That distribution was used to represent the uncertainty regarding the frequency of fires in the cold-rolling mill. The updating procedure is described in chapter 4.

The aim of the analysis was to determine whether the benefits of the risk reduction to be achieved through investing in a water sprinkler system would be sufficient to compensate for the costs associated with it. Since use was made here of extended decision analysis, one needs to relate the epistemic uncertainties regarding the probabilities and the consequences contained in the model to the difference in CE between the two alternatives.

Since over 100 different parameters were considered to be uncertain in the present analysis (see Appendix E), the work of relating the epistemic uncertainty of each to the difference in CE would have been great. Thus, a screening method was employed to identify the parameters that contributed most to the overall epistemic uncertainty regarding the differences in CE. The screening process involved changing each of the parameters, one at a time, from their maximum value to their minimum value (the maximum and minimum values are given in Appendix C) while keeping the other parameters at their most likely value. The effect of this change on the difference between the decision alternatives in terms of CE was noted. If a change in value of a variable result in a large change in the difference in CE, the uncertainty regarding that parameter can be regarded as contributing significantly to the overall uncertainty. When the variables deemed to contribute significantly in this respect had been identified, the uncertainty regarding them could be related to the

difference in CE by use of Monte Carlo simulation. The value of the frequency of fire in the building was varied between 7.3 and 13.3 fires per year. The interval between these values represents an approximate 98% confidence interval for the value of the parameter in question.

The results of the screening process are displayed in terms of the tornado diagram shown in Figure 29, where all parameters that can cause a change of more than 2% (from the minimum to the maximum value) in the difference in CE between the decision alternatives are shown. In that figure, one can note that only a few parameters contribute to any significant degree to the overall uncertainty regarding the difference in CE between the two decision alternatives.

The parameter having the potential to change the difference in CE most is the probability that a fire has the potential to involve the whole area where it started, given that it started in the abrasive-belt grinder in area 5. One can see that several of the parameters with the potential of changing the difference in CE between the decision alternatives significantly pertain to area 5. This makes sense since that area contains the abrasive-belt grinder, a machine in which fires were known to occur frequently (the only major fire in the building occurred there). Furthermore, the machine is not protected by any automatic fire-extinguishing system such as the cold-rolling mills in area 3 are. This implies that if a fire starts in the abrasive-belt grinder and spreads so as to include the oils contained in the machine, there is a good chance that there will be a severe fire. Also, if a fire should grow so as to involve all of area 5, it is likely that it will also spread to some of the other areas, since the fire compartmentation is not very good, especially that between areas 3 and 5. It thus appears reasonable to assume that fires occurring in area 5 would contribute significantly to the overall fire risk in the building, and that epistemic uncertainty regarding parameters pertaining to fires occurring in that area have a strong effect on the difference in CE between the alternatives.




Figure 29 Tornado diagram showing the effect on the difference in CE of changing a parameter from its minimum to its maximum value.

In performing an extended decision analysis concerning investment in a sprinkler system, only parameters that can change the difference in CE between the alternatives by more than 15% (taking the sum of the change in both directions into account) are modelled as probability distributions. Note that some of the parameters are dependent on each other (parameters P101Av and P99Av, for example) and that change in the one thus results in change in the other. In such cases, only the parameter having the potential of changing the difference in CE between the alternatives the most is modelled in the extended decision analysis. The following parameters are treated in the extended decision analysis as being probability distributions: the probability that a fire occurring in the abrasive-belt grinder has the potential of being large (P101Av), the frequency of fires in the building (P1Av), the probability that a fire that occurs in the building will occur in area 3 (P5Av), the probability that a fire starting in a machine in area 5 will start in the abrasive-belt grinder (P97Av and P98Av), the ratio of the cost of a medium-sized fire in a machine to the cost of complete destruction of the machine (*Ratio*), the probability that a fire occurring in the building will occur in area 5 (P7Av), the probability that the employees would succeed in extinguishing a fire in the abrasive-belt grinder in area 5 (P108Av), the probability that the fire department would succeed in extinguishing a fire in the abrasive-belt grinder (P112Av), and the probability that a fire occurring in area 5 will occur in a machine (P96Av).

In performing an extended decision analysis, each of the parameters referred to above is represented by a triangular probability distribution containing the minimum, most likely and maximum value, as shown in Table 10.

Table 10 Maximum, most likely and minimum values for the parameters treated as being uncertain in the extended decision analysis.

Parameter	Max	Most likely	Min
The probability that a fire which occurs in the abrasive-belt grinder has the potential to be large (P101Av).	0.03	0.05	0.1
The probability that a fire that occurs in the building occurs in area 3 (P5Av).	0.25	0.37	0.55
The probability that a fire starting in a machine area 5 will start in the abrasive-belt grinder (P97Av).	0.4	0.5	0.8
The ratio of the cost of a medium-sized fire in a machine to the cost of complete destruction of the machine (Ratio).	0.05	0.1	0.25
The probability that a fire occurring in the building will occur in area 5 (P7Av).	0.15	0.28	0.35
The probability that the employees would succeed in extinguishing a fire in the abrasive-belt grinder in area 5 (P108Av).	0.5	0.75	0.8
The probability that the fire department would succeed in extinguishing a fire in the abrasive-belt grinder (P112Av).	0.6	0.7	0.9
The probability that a fire occurring in area 5 occurs in a machine (P96Av).	0.75	0.97	0.98

It was decided in the extended decision analysis to treat the risk attitude of the decision maker (Avesta Sheffield) as being risk-neutral, since no investigation of risk-attitudes was possible inasmuch as this would have required far too much effort on the part of senior managers in the company and it was also outside the scope of the case study.

In relating the uncertainties regarding the parameters taken up above (the screening process) to the difference between the decision alternatives in terms of CE, use was made of Monte Carlo-simulation. The results of 5000 simulations are shown in Figure 30. In the calculation, the benefits from the risk reduction associated with the sprinkler investment are accounted for during a 5-year period.




Figure 30 Histogram showing the effect of the epistemic uncertainties on the risk-adjusted net present value for the sprinkler investment in the cold-rolling mill. A limit of 5 years has been used for how long the benefits of the risk reduction are to be taken account of.

Note that the Risk-adjusted net present value, used to denote the difference between the decision alternatives in terms of CE, is very high. The mean value of the simulations is 59 million SEK. One can also note that the decision situation is robust, since there is no overlap of the 0-values on the horizontal axis. Thus, according to the extended decision analysis, investment in a sprinkler system is a reasonable decision. Note that the difference between Figure 30 and the figure showing the risk-adjusted net present value contained in [64] (Paper 2) is due to the time period being different (which accounts for most of the dissimilarities) and that there were also minor differences in the risk analysis technique employed in producing the two figures (see the beginning of this chapter).

Instead of only analysing a single time period for the investment in the sprinkler system, one can present the risk-adjusted net present value as a function of the time period considered, as is done in Figure 31. In that figure, the uncertainty regarding the risk-adjusted net present value is shown by the two dashed lines, which represent the boundaries within which 90% of the values obtained in the Monte Carlo-simulation lie. One could say that the lines represent the 5% and the 95% robustness

index. Looking at Figure 31 one can conclude that the robustness of the decision is not sensitive to changing the number of years that the benefits of the risk reduction are taken account of.




Figure 31 *Diagram showing the effects of changing the number of years for which the benefit of the risk reduction is taken into account in the analysis. The dashed lines represent a robustness index of 5% and 95%, respectively.*

Comparison of case studies to fire statistics

To investigate whether the estimates of the probabilities of the different types of fire scenarios presented in the case studies correlate well with available fire statistics, fire statistics from Swedish companies obtained during 1996, 1997 and 1998 were related to the case studies. Note that comparing results of the risk analyses with statistics of apparent relevance does not aim at determining whether either the model or the estimates made here are “wrong”. Since a decision analytical (subjective) framework is employed, one cannot say that an estimate is right or is wrong but can only use such a comparison to strengthen or weaken one’s belief in the estimates performed by the expert(-s). One would expect the results (the probabilities of the different types of fire scenarios considered) to be of about the same order of magnitude as indicated by the industrial statistics available.

Unfortunately, there is not very much detailed information regarding industrial fires that have occurred in Sweden. The information available for this study comes from the Swedish Rescue Services Agency and is presented in [99]. It applies to the years 1996, 1997 and 1998. The information is not sufficiently detailed to allow one to determine the probability of each of the different types of fire scenarios used in the case studies. In [99] a suggestion of a simple type of event tree (shown in Figure 32) that can be used to characterise the fire scenarios to which the statistics apply is provided.




Figure 32 Event tree illustrating different fire scenarios.

In comparing the statistics from the “Metalworking and machine industry” and “Other branches of manufacturing” with the results of the Avesta Sheffield and the ABB decision analyses, respectively, it was assumed that the scenarios in Figure 32 termed “3” and “4” pertaining to the ABB and the Avesta Sheffield cases, respectively, represent scenarios in which the potential of the fire is medium or large and in which neither any fire protection system nor the employees succeed in extinguishing the fire. Although this is a crude approximation, it is the best that can be achieved in view of the quality of the statistical information available. The type of fire scenario referred to above will be termed “a serious fire”.

The information just referred to is available for buildings with and without water sprinkler systems. Since there are few fires reported in buildings with sprinkler systems, however, no estimates of the probabilities of the different types of fire scenarios shown in Figure 32 could be made for buildings of the types involved. Accordingly, a comparison of the analyses carried out for the ABB and the Avesta Sheffield buildings with the fire statistics available was only performed for a design without any water sprinkler system.

In Table 11 and Table 12 estimates are presented of the probabilities of a serious fire in the ABB and the Avesta Sheffield building, respectively, the probabilities being given in terms of their most probable, maximum and minimum values.

Table 11 The probability of a serious fire in the ABB building, given that a fire has occurred in a specific area there.

	Min	Most likely	Max
New PK Workshop	0.008	0.022	0.144
The A workshop	0.020	0.080	0.156
The storage	0.121	0.242	0.486
ABB training center	0.030	0.101	0.204
EMC	0.030	0.080	0.204
The PS workshop	0.040	0.121	0.244
The office	0.010	0.030	0.122
The old PK workshop	0.010	0.040	0.096

The information contained in [99] can be used to estimate the probability of a serious fire, i.e. of scenario 3 or 4 in Figure 32. In doing this, the total number of fires of these types is divided with the total number of fires reported. The total number of fires reported in the “Metalworking and machine industry” is 852, 70 of these being judged to belong to either scenario 3 or scenario 4 in Figure 32. The total number of fires reported in the “Other branches of manufacturing” category is 561, 44 of these being judged to belong to either scenario 3 or scenario 4. This results in an estimate of the probability of a serious fire, given that a fire has occurred, of 0.082 for buildings belonging to the “Metalworking and machine industry” and of 0.078 for buildings belonging to the category “Other branches of manufacturing”.

Table 12 The probability of a serious fire in the Avesta Sheffield building, given that a fire has occurred in a specific machine or in a specific area.

	Min	Most likely	Max
Area 1			
Smoothing roller	0.030	0.060	0.150
Cutter 1	0.015	0.034	0.084
Cutter 2	0.008	0.014	0.060
Other	0.010	0.040	0.075
Area 2			
Cutter 3	0.005	0.020	0.060
Other	0.015	0.034	0.150
Area 3			
Cold-rolling mill 1	0.001	0.004	0.030
Cold-rolling mill 2	0.00002	0.0002	0.007
Strip coiling machine	0.003	0.018	0.060
Other	0.008	0.033	0.090
Area 4			
Weld 55	0.005	0.020	0.090
Cold-rolling mill	0.001	0.006	0.060
Oven 55	0.005	0.024	0.090
Weld 60	0.005	0.020	0.090
Oven 60	0.005	0.024	0.090
Other	0.010	0.036	0.090
Area 5			
Abrasive-belt grinder	0.030	0.050	0.150
Oil-room	0.020	0.060	0.150
Other	0.010	0.036	0.150
Other areas			
Engine rooms	0.011	0.032	0.132
Machine shops	0.010	0.020	0.120

It is reasonable to assume that not all fires belonging to scenario 1 or 2 are reported to the fire department, which would mean their not necessarily being included in the statistical information available. At the same time, it would nevertheless seem reasonable to assume that most of the fires that would have been adjudged to belong to scenario 3 or 4 are included in the material, since these fires are particularly severe, its thus

appearing very likely that the fire department attended the most fires of these types. Accordingly, it is assumed that if all such fires that occurred were included in the statistical material available, estimates of the probability a fire that belonged to scenario 3 or 4 occurring would (as shown in Figure 32) be less than 0.082 and 0.078, respectively.

In comparing the figure 0.082 for scenario 3 or 4 with the estimates made of the probability of a serious fire occurring in the ABB building (see Table 11) one notes that the values termed “most likely” are of about the same order of magnitude as the estimates produced by use of the statistics available. Note that the estimates differ somewhat, depending on the area of the building involved. This is to be expected, since some of the areas differ considerably from the others with respect to fuel configuration, manual fire extinguishing equipment, and the like. For example, the probability of a serious fire, given that a fire has occurred in the storage area, is high due to large amounts of combustible material being stored in that area and there normally being no one there who can quickly initiate extinguishing operations.

In comparing the figure 0.078 for scenario 3 or 4 with the estimates the probability of a serious fire occurring in the Avesta Sheffield building (as presented in Table 12) one can note that the probability estimates presented as the most likely values in the table are all lower than 0.078. For the majority of the probabilities, however, 0.078 is within the uncertainty interval involved (ranging from the max to the min-value). Note that the probability of a serious fire occurring is judged to be lowest in the various cold-rolling mills. This appears reasonable enough, since there are automatic extinguishing systems in these machines that can be expected to reduce the probability of a serious fire in these machines as compared with machines in which there is no automatic fire extinguishing system.

Note in Table 11 and Table 12 that the intervals between the minimum and maximum probability levels seen as plausible are substantial. This is due to the high degree of epistemic uncertainty regarding both the probability that a fire that occurred would be small and that the employees would succeed in extinguishing it and these two probabilities being the ones used in calculating the probabilities of interest here (the probability of the extinguishing systems within some of the machines succeeding in extinguishing a fire was used as well in the Avesta Sheffield case).

As indicated in the beginning of this section, the comparison of the results of the analysis with the statistical information available is not intended to provide a basis for judging whether the results of the decision analysis are “right” or are “wrong”, but should simply be viewed as an attempt to assess whether the results are reasonable. Note that even if the information regarding fires in a specific type of industry was directly applicable to the fire model employed here, i.e. if one knew how many fires had developed in accordance with the different types of fire scenarios considered for the ABB and the Avesta Sheffield building, respectively, it would be difficult to draw any final conclusions regarding the results of the decision analyses presented in the thesis. This is because even though the buildings involved belong to particular branches of industry for which statistical information is available, it is reasonable to assume that there is a great deal of variation in buildings of any particular type. Thus, if one is interested in the probability of a serious fire occurring in a specific building, one cannot say that the estimates one has made are “wrong” simply because of their not agreeing with estimates available for the industrial category to which the building belongs. Since there are many different types of buildings belonging to any such category, the probabilities of interest in one particular building could readily differ from those applying to the type of building to which the buildings belongs.

In comparing estimates of the probability of a serious fire as based on the general statistics available with the estimates made in connection with the case studies here, one can conclude that the estimates in both cases are of basically the same magnitude. This can be regarded as supporting the credibility of the estimates made in the case studies. However, due to the problems referred to above, one should not place undue emphasis on this conclusion, since even if the estimates in the case studies had differed considerably from the values arrived at on the basis of the general information, this would not have meant that the estimates in the case studies were wrong.

Concluding remarks regarding the case studies

The two case studies exemplify the practical application of extended decision analysis as described in chapter 3. In both case studies, use of extended decision analysis led to the recommendation that a sprinkler system be invested in. Note, however, that since these conclusions were

reached in modelling the beliefs and judgements of the decision makers involved, the conclusions reached are not applicable in a general way to buildings other than those that were studied.

The very high risk-adjusted net present value arrived at in analysis of the Avesta Sheffield case might be viewed as being a bit too high. Its being so high due to a combination of the potential damage costs being very high, the maximum total costs being found to be 4400 million SEK (on the basis of the maximum values as reported in Table 8) and the poor fire protection available in the building's original design. The extremely high damage costs, or *monetary equivalent* of the worst fire scenarios, are due to the cold-rolling mill's being one of Avesta Sheffield's major factories, its supplying several other factories with stainless steel. In addition, the machines in the factory take a very long time to replace. The poor fire protection available initially is due to a combination of poor fire compartmentation and the presence of large amounts of combustibles in the form of oil, pallets, and the like. In such an environment, any protection measure able to reduce the probability of a major fire appreciably has a strong impact on the CE of the exposure, and thus has a high risk-adjusted net present value.

In the case of ABB, the company chose not to consider any negative consequences that occurred at other locations within the company than in building 358. This suggests that the risk-adjusted net present value calculated in the ABB study may be too low, since possible negative consequences may have been neglected. Nevertheless, estimates of the losses associated with serious fires in the ABB building are substantial, the total loss for the building in question being equal to a loss of approximately 1500 million SEK. Although the initial fire protection was judged to be somewhat better in the ABB building than in the Avesta Sheffield building, the positive effect of installing a water sprinkler system there was nevertheless estimated to be substantial. The analysis performed showed, however, that the decision to invest in a water sprinkler system in the ABB building was not as robust as a decision of this sort reached for the Avesta Sheffield building. In the analysis of the ABB building the risk-adjusted net present value was found to be 8 million SEK (computed for a time period of 5 years) and the robustness index to be 96%. For the Avesta Sheffield case, the risk-adjusted net present value was found to be about 52 million SEK and the robustness index to be 100%.

No investigation of the risk attitude of the decision makers was included in either of the case studies. Instead, the decision makers' risk attitude was assumed to be that of risk neutrality which means evaluating uncertain situations in terms of their expected value alone. Making this assumption was due to its not being practically possible to investigate the risk attitude of the decision makers through conducting interviews with top management to ask them their preferences regarding choices in risky situations. The only members of management available in the study were the risk managers of the two companies. Although it might have been feasible to investigate their risk-attitudes it was considered that since the results of investigating their risk-willingness in isolation might well not be representative of the company involved, the best thing was to simply assume in both cases that the risk attitude of the decision makers was risk-neutral.

Note that one cannot draw any general conclusions for either the two companies on the basis of these case studies since it is not possible to say, for example, whether investing in a sprinkler system is good generally for buildings of the types involved. The aim of the case studies was simply to show how the methods presented in the thesis could be employed in a real situation, so as to provide the reader a better understanding of the usefulness of the methods.

6 Summary, conclusions and future work

Summary

The thesis is concerned with the evaluation of possible investments in fire safety for specific factories, particularly in cases in which a monetary evaluation of the risk reduction the investment would involve is sought. Previously developed methods examined here in terms of their applicability in evaluating such investments include those concerning expected costs and expected utilities (see chapter 1). It is argued that expected-cost methods are not well suited to analysing fire protection investments in a *specific* factory, due to the numbers of fires expected to occur during the lifetime of most such investments being so low that random effects do not tend to level out, which means that the actual costs of fire are likely to deviate markedly from the expected costs. For this reason, appropriate decision rules were sought within the area of normative decision theory. There, one starts by specifying a set of axioms for decision making that appear intuitively reasonable, and seeks suitable decision rules that can be shown to be in agreement with the axioms postulated initially. The normative decision rule of this sort most commonly employed is the principle of maximising expected utility.

A major concern was to find a normative rule that could serve as the basis for a prescriptive rule in this context, a rule that would *help* the decision maker arrive at well-informed decisions.

Although the maximisation of expected utility could be considered the dominant *normative* decision rule due to its frequent use, it has been criticised substantially both from a descriptive and from a normative standpoint. Some of the criticisms have led to the development of alternative decision rules. Several of the methods based on these are examined critically in chapter 2 and 3 with the aim of determining to what extent they possess features desirable for a prescriptive decision rule to be used in the present context.

It is concluded that in the present context none of those methods are obviously superior to the maximisation of expected utility rule. Nevertheless, since the principle of maximising expected utility has certain drawbacks when employed here, the ideas embodied in various alternative decision rules are utilised in the decision model suggested.

More precisely, the maximisation of expected utility rule requires that the decision maker express his/her assessments in terms of exact probability values or of exact probability distributions to represent these values. Since the information available regarding some of the events important to the development of a fire is often scarce, the decision maker here may feel uncomfortable using exact values or exact probability distributions. It is argued that in order to take adequate account of the possible lack of information regarding various of the probability values involved in the analysis of such investments, what is needed is not some *single* method but rather an evaluative framework involving various methods, three such methods being suggested for use in conjunction with each other. Depending on the “vagueness” of the information available, regarding in particular probabilities, any one of the three methods and the prescriptive decision rule connected with it may appear to be most appropriate. The one method, termed “Traditional decision analysis”, involves assessment of expected utility of each of the decision alternatives. A second method, termed “Extended decision analysis”, involves the evaluation, not simply of expected utilities but also of decision robustness. The latter concerns how likely it is that the alternative found to have the highest expected utility would change if the epistemic uncertainties regarding the probability and utility values involved were to be eliminated (see chapter 3). The third method, called “Supersoft decision analysis”, involves probabilities and utilities being expressed by use of vague statements such as “the probability that the employees will succeed in extinguishing a fire is at least 0,2”. The evaluation of the decision alternatives in a concrete case would involve the appropriate use of three decision criteria, each of them based on the use of expected utilities.

Figure 33 shows the different steps one would take in developing the operational model to be employed, i.e. the model to be used in a practical situation, beginning with the selection of a suitable normative theory, proceeding to the choice of appropriate prescriptive decision rules based on that theory, and concluding with the development of an operational model involving use of these prescriptive decision rules.




Figure 33 The steps that can be taken in developing an operational model for the analysis of possible investments in fire safety.

From a pragmatic standpoint, the perhaps greatest advantage in using the evaluatory framework for assessing investments in fire safety suggested here is that it allows a reduction in risk to be expressed in monetary terms. This enables what is termed here the risk-adjusted net present value of an investment in fire safety to be calculated. It represents the sum of the benefits expressed in monetary terms which the investment provides (the risk reduction being included here) minus the costs of the investment. In calculating this value, use is made of the “primary model” shown in Figure 33.

The evaluatory framework that is suggested can be employed for assessing not only possible investments in fire safety but also *changes* in fire risk. In paper 4 a combination of this framework with Bayesian networks is described. This allows the Bayesian network that is used to be updated by means of frequent measurements being made in the building in question, which in turn allows measures of fire risk in the building to be updated so that changes in fire risk can be evaluated adequately in monetary terms. Since fires do not occur very often, a method for using subjective judgments to update the Bayesian network, such as those provided by experts in connections with annual inspections of the building, is likewise presented.

Two case studies are also included in the thesis for illustrating how the extended decision analysis method described here can be used in practice. The two case studies were performed at the companies ABB and Avesta Sheffield. In both cases, analysing the possible investment in a particular water sprinkler system was involved.

Conclusions

A new method (or methods) for analysing *specific* investments in fire safety in *specific* factories is suggested in the thesis. Compared with previous suggestions of such methods, it provides a new way of estimating the monetary value of the reduction in risk that an investment in fire safety involves. Most importantly, it explicitly addresses epistemic uncertainties and can be used to evaluate decision alternatives even when the magnitude of these uncertainties is considerable.

A number of general conclusions can be drawn on the basis of the work presented:

- The principle of maximising expected utility appears to be the normative decision rule most suitable in the present context.
- Certain additional evaluations are seen as being useful in the practical application of the principle of maximising the expected utility. Two of these, termed Supersoft decision analysis and extended decision analysis, are suggested in the thesis. Together with the original expected utility evaluation, they form the evaluatory framework for assessing investments in fire safety suggested here.
- This evaluatory framework for the analysis of investments in fire safety is very flexible and can be used not only in situations in which one is basically certain regarding the variables of interest (probabilities and consequences) but also in situations in which one is extremely uncertain.
- Supersoft decision analysis and extended decision analysis should not be viewed as competing, but rather as complementary methods, the one being useful when the information at hand justifies epistemic uncertainties being expressed as specific probability distributions, the other being useful when the information at hand is vague and does not justify expressing uncertainties as specific probability distributions.
- Use of this evaluatory framework allows fire risk to be evaluated in monetary terms.

- The evaluatory framework can also be used in combination with Bayesian networks for measuring changes of fire risk, which can likewise be expressed in monetary terms.
- The case studies show the methods suggested to basically be applicable in practical situations. Since the analyses carried out required a great deal of work and effort, however, the procedures employed for estimating the probability of different fire scenarios here would probably need to be simplified in order for the methods to be useful in practical situations.

Future work

Various direction of future work can be suggested.

Investigating relationships between different probabilities

Investigating how different probabilities of relevance in fire risk analysis and decision analysis are related would be of considerable interest. The information this would provide could help to make both risk analysis and decision analysis more credible. Research of this type would be particularly useful in connection with extended decision analysis since that method models the uncertainty regarding the probability values involved explicitly.

In studying relationships between different probabilities, one should also investigate the reliability of other types of protection systems than sprinkler systems and fire detection systems, such as manual fire fighting systems, for example. In such an investigation, identifying the factors that have bearing on the probability that the occupants of a building will succeed in extinguishing a fire is important.

Work on developing improved methods of estimating subjective probabilities

Although a wide variety of methods for estimating subjective probabilities are available (see [68], for example), more should be known of how adequately various of these methods are for use in the context of fire safety engineering. Work on the further development of methods of this sort is important so as to increase the credibility of subjectively estimated probabilities.

Decision analysis and public safety

In a concrete sense, the thesis was concerned above all with fire safety decisions to be made by a company. It would be of interest to investigate to what extent methods similar to those employed here could also be useful in other areas, such as the public safety area, for example, where how satisfactory a specific building design is from a fire-safety standpoint is important.

Investigation of risk-attitudes

Investigating possible relationship between the risk-attitudes typical of companies in various branches of operation and such company characteristics as turnover, profit margin, and the like, would be of interest. The results could make it easier to understand important factors to bear in mind analysing possible investments in fire safety in companies of differing character.

Uncompensated losses

Better methods for estimating uncompensated losses due to fire are needed, indicating the intrinsic monetary value of a given fire scenario. Such methods could be particularly useful for the evaluation of fire safety investments of various types, making it possible to better anticipate the effects of serious fires, for example. One could investigate the relationship between the insured losses and the uncompensated losses incurred in earlier fires, which is important since the insured losses a particular fire scenario would involve are generally easier to estimate than the uncompensated losses.

Probabilities that change with time

The probabilities that different safety systems will work have been dealt with in the thesis as being constant over the course of time. This is an assumption that may be in need of modification since aging might well cause components to have a lower reliability than when they were new. Whether aging would have a noticeable affect on the results of a decision analysis is difficult to say. However, since in performing such an analysis one takes account of the risk reduction during a long period of time, one should at least investigate what effect aging might have on the results of a decision analysis, its hopefully being possible to quantify the effect of aging on reliability, if such effect is found.

7 Acknowledgements

The work presented in the thesis has been supervised by Professor Sven-Erik Magnusson, who has provided valuable support and encouragement. Professor Per-Erik Malmnäs has contributed significantly to the part of the thesis that deals with evaluation of decision situations involving great epistemic uncertainty (Supersoft decision theory). Dr. Robert Goldsmith and Dr. Håkan Frantzich have 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 Grahm (Avesta Sheffield), Olle Österholm (Avesta Sheffield), Bo Sidmar (ABB) and Michael Zeeck (ABB).

In 2001 I spent a year in Boston, USA, where I was working with Dr. Brian Meacham at ARUP Risk Consulting. I would like to thank Brian and the rest of the staff at the ARUP office for making my stay in Boston a very rewarding experience.

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, 103-991 and 105-001. 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 Grahm (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 Administrative Services Agency), Per Nyberg (Skandia), Björn Lindfors (SKF), Jan-Erik Johansson (Stora Risk Management), Liselotte Jonsson (Sycon) and Per-Erik Malmnäs (Stockholm University).

8 References

1. Hall, J. R., Jr., Sekizawa, A., Fire Risk Analysis: General Conceptual Framework For Describing Models, *Fire Technology*, Vol. 27, No. 1, pp. 33-53, 1991.
2. Donegan, H. A., Decision Analysis, *SFPE Handbook of Fire Protection Engineering*, 3rd ed., National Fire Protection Association, Quincy, 2002.
3. Ramachandran, G., *The Economics of Fire Protection*, E & FN Spon, London, 1998.
4. Shpilberg, D., De Neufville, R., Best Choice of Fire Protection: An Airport Study, *Fire Technology*, Vol. 10, No. 4, pp. 5-14, 1974.
5. Cozzolino, J. M., A Method for the Evaluation of Retained Risk, *The Journal of Risk and Insurance*, Vol. XLV, No. 3, 1978.
6. Helzer, S. G., Buchbinder, B., Offensend, F. L., *Decision Analysis of Strategies for Reducing Upholstered Furniture Fire Losses*, National Bureau of Standards, Gaithersburg, 1979.
7. Noonan, F., Decision Analysis for Risk Management of Firesafety Hazards, *Proc. 1st International Symposium on Fire Safety Science*, pp. 621-625, 1986.
8. Van Anne, C., Evaluation of the Risk Problem and the Selection of the Optimum Risk Management Solution, *Proc. 1st International Symposium on Fire Safety Science*, pp. 627-636, 1986.
9. Fitzgerald, R. W., Risk Analysis Using the Engineering Method for Building Safety, *Proc. 1st International Symposium on Fire Safety Science*, pp. 993-1002, 1986.
10. Watts, J. M., Fire Risk Assessment Using Multiattribute Evaluation, *Proc. 5th International Symposium on Fire Safety Science*, pp. 679-690, 1997.
11. Watts, J. M., Criteria for Fire Risk Ranking, *Proc. 3rd International Symposium on Fire Safety Science*, pp. 457-466, 1991.
12. Watts, J. M., Fire Risk Ranking, *SFPE Handbook of Fire Protection Engineering*, 3rd ed., National Fire Protection Association, Quincy, 2002.
13. Gretener, M., *Evaluation of Fire Hazard and Determining Protective Measures*, Association of Cantonal Institutions for Fire Insurance (VKF) and Fire Prevention Service for Industry and Trade (BVD), Zürich, 1973.

14. Budnick, E. K., McKenna, L. A., Watts, J. M., Quantifying Fire Risk for Telecommunications Network Integrity, *Proc. 5th International Symposium on Fire Safety Science*, pp. 691-700, 1997.
15. Takeda, H., Yung, D., Simplified growth models for risk-cost assessment in apartment buildings, *Journal of Fire Protection Engineering*, Vol. 4, No. 2, pp. 53-66, 1992.
16. Beck, V., CESARE-RISK: A Tool for Performance-Based Fire Engineering Design, *2nd International Conference on Performance-Based Codes and Fire Safety Design Methods*, pp. 391-402, 1998.
17. Benichou, N., Yung, D. T., *FiRECAM <TM>: an equivalency and performance-compliance tool for cost-effective fire safety design*, National Research Council of Canada, NRCC-44695, Ottawa, 2001.
18. Johansson, A Bayesian network model for the continual updating of fire risk measurement, *Submitted to Fire Science and Technology (Included in the thesis)*.
19. Van Anne, C., *Managing Risk Through Decision Analysis to Evaluate Firesafety Hazards and Extent of Probable Loss by Use of an Engineering Method*, Worcester Polytechnic Institute, Worcester, 1985.
20. Arrow, K. J., Lind, R. C., Uncertainty and the Evaluation of Public Investment Decisions, *The American Economic Review*, Vol. 60, No. 3, pp. 364-378, 1970.
21. Laughhunn, D. J., Payne, J. W., Crum, R., Managerial Risk Preferences for Below-Target Returns, *Management Science*, Vol. 26, No. 12, pp. 1238-1249, 1980.
22. Spetzler, C. S., Establishing a corporate risk policy, *Readings in Decision Analysis*, (eds.) Howard, R. A., Matheson, J. E., Stanford Research Institute, Stanford, 1977.
23. Keeney, R. L., On the foundations of prescriptive decision analysis, *Utility theories: measurements and applications*, (ed.) Edwards, W., Kluwer Academic Publishers, Boston, 1992.
24. Howard, R. A., In Praise of the Old Time Religion, *Utility Theories: Measurements and Applications*, (ed.) Edwards, W., Kluwer Academic Publishers, Boston, 1992.
25. Coombs, C. H., Dawes, R. M., Tversky, A., *Mathematical Psychology, An Elementary Introduction*, Prentice-Hall, Englewood Cliffs, 1970.
26. Luce, R. D., Raiffa, H., *Games and decisions*, John Wiley & Sons, New York, 1957.

27. Bernoulli, D., Specimen Theoriae Novae de Mensura Sortis, in: *Commentarii Academiae Scientiarum Imperialis Petropolitannae* (1738); translated from Latin into English by L. Sommer, Exposition of a New Theory on the Measurement of Risk, *Econometrica*, Vol. 22, No. 1, pp. 23-36, 1954.
28. von Neumann, J., Morgenstern, O., *Theory of Games and Economic Behaviour*, 2nd ed., Princeton University Press, Princeton, 1947.
29. Savage, L. J., *The Foundations of Statistics*, Wiley, New York, 1954.
30. Fishburn, P. C., *Nonlinear Preference and Utility Theory*, Wheatsheaf Books, Brighton, 1988.
31. Herstein, I. N., Milnor, J., An axiomatic approach to measurable utility, *Econometrica*, Vol. 21, No. 2, pp. 291-97, 1953.
32. Oddie, G., Milne, P., Act and value, *Theoria* LVII, pp. 42-76, 1990.
33. Krantz, D. H., Luce, R. D., Suppes, P., Tversky, A., *Foundations of measurement*, Academic press, New York, 1971.
34. Fishburn, P. C., Foundations of Decision Analysis: Along the Way, *Management Science*, Vol. 35, No. 4, pp. 387-405, 1989.
35. Gärdenfors, P., Sahlin, N.-E., Introduction: Bayesian decision theory – foundations and problems, *Decision, Probability, and Utility*, (eds.) Gärdenfors, P., Sahlin, N.-E., Cambridge University Press, Cambridge, 1988.
36. Clemen, R. T., *Making Hard Decisions – An Introduction to Decision Analysis*, Second edition, Duxbury Press, Belmont, 1996.
37. von Winterfeldt, D., Edwards, W., *Decision analysis and behavioral research*, Cambridge University Press, Cambridge, 1986.
38. Johansson, H., Decision analysis concerned with investments in fire safety, *Proc. 7th International Symposium on Fire Safety Science*, 2002. (Included in the thesis)
39. Ellsberg, D., Risk, ambiguity, and the Savage axioms, *Quarterly Journal of Economics*, Vol. 75, No. 4, pp. 643-669, 1961.
40. Kahneman, D., Tversky, A., Prospect theory: An analysis of decision under risk, *Econometrica*, Vol. 47, No. 2, pp. 263-291, 1979.
41. Bell, D. E., Regret in decision making under uncertainty, *Operations Research*, Vol. 30, No. 5, pp. 961-981, 1982.
42. Loomes, G., Sugden, R., Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty, *The Economic Journal*, Vol. 92, No. 368, pp. 805-824, 1982.
43. Hodges, J. L., Lehmann, E. L., The use of Previous Experience in Reaching Statistical Decisions, *Annals of Mathematical Statistics*, Vol. 23, No. 3, pp. 396-407, 1952.

44. Quiggin, J., A theory of anticipated utility, *Journal of Economic Behaviour and Organization*, Vol. 3, pp. 323-343, 1982.
45. Yaari, M. E., The dual theory of choice under risk, *Econometrica*, Vol. 55, No. 1, pp. 95-115, 1987.
46. Tversky, A., Kahneman, D., Advances in Prospect Theory: Cumulative Representation of Uncertainty, *Journal of Risk and Uncertainty*, Vol. 5, pp. 297-323, 1992.
47. Malmnäs, P.-E., *Methods of Evaluation in Supersoft Decision Theory*, Department of Philosophy, Stockholm University, Stockholm, 1995.
48. Johansson, H., Malmnäs, P. -E., Application of Supersoft decision theory in fire risk assessment, *Accepted for publication in Journal of Fire Protection Engineering. (Included in the thesis)*
49. Danielsson, M., *DDT – The delta decision tool*, Department of Computer and Systems Sciences, Stockholm University, Stockholm, 1996.
50. Frantzich, H., *Risk Analysis in Fire Safety Engineering*, Report 1016, Dept. of Fire Safety Engineering, Lund University, Lund, 1998.
51. Kaufmann A., Gupta M. M., *Introduction to Fuzzy Arithmetic*, Van Nostrand Reinhold, New York, pp. 1-77, 1985.
52. Saaty, T. L., *The Analytic Hierarchy Process: Planning, setting priorities, resource allocation*, RWS Publications, Pittsburgh, 1990.
53. Eppel, T., Matheson, D., Miyamoto, J., Wu, G., Eriksen, S., Old and new roles for expected and generalized utility theories, *Utility theories: Measurement and applications*, (ed.) Edwards W., Kluwer, Boston, 1992.
54. Malmnäs, P.-E., *Foundations of Applicable Decision Theory*, Department of Philosophy, Stockholm University, Stockholm, 1999.
55. Malmnäs, P. -E., (ed.), *Hypermjuk beslutsteori och ekonomisk optimering av det industriella brandskyddet*, Thales, Stockholm, 2002.
56. Laplace, P. S., *Essai Philosophique sur les Probabilites*, 5th ed., Paris, Translation by Dover, New York, 1952.
57. Wald, A., *Statistical Decision Functions*, Wiley, New York, 1950.
58. Savage, L. J., The Theory of Statistical Decision, *Journal of the American Statistical Association*, Vol. 46, No. 253, pp. 55-67, 1951.
59. Hurwicz, L., *Optimality Criteria for Decision Making Under Ignorance*, Discussion paper No. 370, Cowles Commission, 1951.
60. Winkler, R. L., Decision modelling and rational choice: AHP and utility theory, *Management Science*, Vol. 36, No. 3, pp. 247-248, 1990.

61. Malmnäs, P.-E., Axiomatic Justifications of the Utility Principle: a Formal Investigation, *Synthese*, Vol. 99, pp. 233-249, 1994.
62. Keeney, R. L., Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley & Sons, New York, 1976.
63. Farquhar, P. H., Utility Assessment Methods, *Management Science*, Vol. 30, No. 11, pp. 1283-1300, 1984.
64. Johansson, H., Investment appraisal using quantitative risk analysis, *Journal of Hazardous Materials*, Vol. 93, pp. 77-91, 2002. (Included in the thesis)
65. Oberkampf, W. L., DeLand, S. M., Rutherford, B. M., Diegert, K. V., Alvin, K. F., Error and uncertainty in modeling and simulation, *Reliability engineering and system safety*, Vol. 75, No. 3, pp. 333-357, 2002.
66. International Atomic Energy Agency (IAEA), *Evaluating the Reliability of Predictions Made Using Environmental Transfer Models*, Safety Series No. 100, Vienna, 1989.
67. Goldsmith, R. W., Assessing probabilities of compound events in a judicial context, *Scand. J. Psychol.*, Vol. 19, pp. 103-110, 1978.
68. Spetzler, C. S., Stael Von Holstein, C.-A. S., Probability Encoding in Decision Analysis, *Management Science*, Vol. 22, No. 3, pp. 340-358, 1975.
69. Kahneman, D., Slovic, P., and Tversky, A., *Judgement under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge, 1982.
70. Kaplan, S., Garrick, B. J., On the Use of a Bayesian Reasoning in Safety and Reliability Decisions – Three Examples, *Nuclear Technology*, Vol. 44, pp. 231-245, 1979.
71. Ang, A. H-S., Tang, W. H., *Probability Concepts in Engineering Planning and Design, Volume 2 – Decision, Risk and Reliability*, John Wiley & Sons, New York, 1984.
72. Sui, N., Apostolakis, G., Uncertain Data and Expert Opinions in the Assessment of the Unavailability of Suppression Systems, *Fire Technology*, Vol. 24, No. 2, pp. 138-162, 1988.
73. Sui, N., Apostolakis, G., Modeling the Detection Rates of Fires in Nuclear Plants: Development and Application of a Methodology for Treating Imprecise Evidence, *Risk Analysis*, Vol. 6, No. 1, pp. 43-59, 1986.
74. Sui, N., Apostolakis, G., A Methodology for Analyzing the Detection and Suppression of Fires in Nuclear Power Plants, *Nuclear Science and Engineering*, Vol. 94, pp. 213-226, 1986.

75. Apostolakis, G., Probability and Risk Assessment: The Subjectivistic Viewpoint and Some Suggestions, *Nuclear Safety*, Vol. 19, No. 3, pp. 305-315, 1978.
76. Olsson, F., *Tolerable Fire Risk Criteria for Hospitals*, Report 3101, Dept. of Fire Safety Engineering, Lund University, Lund, 1999.
77. Frantzich, H., *Fire Risk Analysis of a Hotel*, Report 3091, Dept. of Fire Safety Engineering, Lund University, Lund, 1997.
78. Jönsson, R., Lundin, J., *The Swedish Case Study, Different Fire Safety Design Methods Applied on a High Rise Building*, Report 3099, Dept. of Fire Safety Engineering, Lund University, Lund, 1998.
79. Apostolakis, G., Kaplan, S., Garrick, B. J., Dughity, R. J., Data Specialization for Plant Specific Risk Studies, *Nuclear Engineering and Design*, Vol. 56, pp. 321-329, 1980.
80. Mosleh, A., Apostolakis, G., The Assessment of Probability Distributions from Expert Opinions With an Application to Seismic Fragility Curves, *Risk Analysis*, Vol. 6, No. 4, pp. 447-461, 1986.
81. Morris, P.A., Combining Expert Judgements: A Bayesian Approach, *Management Science*, Vol. 23, No. 7, pp. 679-693, 1977.
82. Apostolakis, G., Mosleh, A., Expert Opinion and Statistical Evidence: An Application to Reactor Core Melt Frequency, *Nuclear Science and Engineering*, Vol. 70, pp. 135-149, 1979.
83. Winkler, R. L., The Consensus of Subjective Probability Distributions, *Management Science*, Vol. 15, No. 2, pp. B61-B75, 1968.
84. Wilks, S. S., *Mathematical Statistics*, John Wiley & Sons, New York, 1962.
85. Rutstein, R., The estimation of the fire hazard in different occupancies, *Fire Surveyor*, Vol. 8, No. 2, pp. 21-25, 1979.
86. Lindley, D.V., *Introduction to Probability and Statistics From a Bayesian Viewpoint, Part 2: Inference*, Cambridge University Press, Cambridge, 1965.
87. Jensen, F. V., *Bayesian Networks and Decision Graphs*, Springer-Verlag, New York, 2001.
88. Särndqvist, S., *Demand for Extinguishing Media in Manual Fire Fighting*, Lund University, Department of Fire Safety Engineering, Lund 2000.

-
89. Tillander, K., Keski-Rahkonen, O., The Influence of Fire Department Intervention to the Fire Safety of a Building Assessed Using Fire Risk Analysis, *Performance-Based Codes and Fire Safety Design Methods*, 3rd International Conference, (eds.) Jönsson, R., Almand, K., Maskas, J., pp. 247-256, Lund, 2000.
 90. Rutstein, R., Gilbert, S., *The performance of sprinkler systems*, Home Office Scientific Advisory Branch Memorandum 9/78, London, 1978.
 91. Young, R. A., Costs and benefits of automatic sprinkler systems, *Commission of the European Communities Symposium, "Fires in Buildings"*, Luxembourg, 1984.
 92. Stirland, C., *Sprinklers and the building regulations – the case for trade-offs*, British Steel Corporation Report T/RS/1189/22/81C, Teesside, 1981.
 93. Marryatt, H.W., Sprinkler trade-offs and the Australasian experience, *Fire Journal*, Vol. 75, No. 1, pp. 59-62, 1981.
 94. Maybee, W. W., American survey shows high success rate for sprinklers, *Fire Prevention*, No.170, pp. 30-33, 1984.
 95. BSI Draft for development DD240: *Fire safety engineering in buildings. Part 1: Guide to the application of fire safety engineering principles*, British Standards Institution, London, 1997.
 96. *Fire Engineering Guidelines*, First edition, Fire Code Reform Center Limited, Sydney, 1996.
 97. Hiles, A., Marketing Protection: a justification for funding of Total Asset Protection programmes?, *The Definite Handbook of Business Continuity Management*, (eds.) Hiles, A., Barnes, P., John Wiley & Sons, Chichester, pp. 67-73, 1999.
 98. Keski-Rahkonen, O., Rahikainen, J., Determination of Ignition Frequency of Fires in Different Premises in Finland, *Fire Engineers Journal*, Vol. 58, No. 197, pp. 33-37, 1998.
 99. Johansson, H., *Decision Making in Fire Risk Management*, Report 1022, Dept. of Fire Safety Engineering, Lund University, Lund, 2001.
 100. Fitzgerald, R.W., *The anatomy of building firesafety, Volume 2, The framework*, Draft 5, V2, WPI Center for Firesafety Studies, Worcester, 2002.

Papers

Paper 1: Decision analysis concerned with investments in fire safety

Johansson, H., *7th International Symposium on Fire Safety Science*, Worcester, USA, 2002.

Decision Analysis Concerned With Investments in Fire Safety

Henrik Johansson

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

Abstract

Decision analysis concerned with investments in fire safety is discussed. Particular attention is directed at the treatment of uncertainty, the evaluation of consequences, and the choice of a decision rule for use in this context. An approach involving use of a decision rule based on the principle of maximising expected utility, together with a complementary evaluation of the decision alternatives, is described, the latter involving analysis of the uncertainty regarding the probabilities and consequences of different fire scenarios.

Keywords: *Decision analysis, Bayesian methods, uncertainty, fire safety.*

1. Introduktion

The management of an organisation has the obligation towards the shareholders and other interested parties, of managing effectively any risks that can threaten the organisation's goals. This involves making decisions concerning risk-reducing investments such as investments in fire safety. The present paper deals with various aspects of decision analysis concerned with investments in fire safety, both a decision rule and a method for performing such an analysis being suggested. The paper will focus on the choice between different fire protection alternatives for a given building. Note that what is of interest here is the choice between decision alternatives, not the attempt to determine whether a given decision alternative *possesses* certain necessary or desired properties, as would be the case if one employed decision analysis to investigate whether an alternative met the requirements of the building codes that apply. It is assumed that all fire protection alternatives that are considered comply with the building codes that are applicable.

In the following section, the connection between risk analysis and decision analysis is discussed. The practical benefits are pointed out of

using a risk analysis of a particular building as a point of departure when performing a decision analysis regarding possible fire protection measures for a building. Section 3 deals with the treatment of uncertainty. Section 4 is concerned with estimating and evaluating consequences within a decision analysis. Section 5 deals with decision rules for the analysis of different fire protection alternatives. In section 6, the decision method suggested is presented, together with a real-world example. In section 7, finally, a number of conclusions are drawn regarding the use of decision analysis for evaluating fire safety investment alternatives.

2. Risk analysis

It is assumed that the decision analysis is based on a quantitative risk analysis. If the general framework for fire-risk analysis outlined by Hall and Sekizawa [1] is employed, for example, the “fire risk” involved or the “outcome measure of fire risk”, is defined according to Eq. 1. The term $g(s')$ in Eq. 1 is a function that transforms the severity measure s' into the measure of interest in the risk analysis. For example, if s' is the monetary loss due to a specific fire scenario and the measure of interest is monetary losses, then $g(s') = s'$. $P(s = s')$ in Eq. 1 refers to the probability that the severity measure s' will occur.

$$Risk = \int_{-\infty}^{+\infty} g(s')P(s = s') ds' \quad (1)$$

Since in practice it is likely that the risk measure will be based upon a finite set of fire scenarios, Eq. 1 can be replaced by Eq. 2, in which n is the number of fire scenarios that is deemed to be *relevant* in the building in question.

$$Risk = \sum_{i=1}^n g(s_i)P(s = s_i) \quad (2)$$

If one performs a risk analysis using Eq. 2 one must have a number of different fire scenarios that have been defined, together with the outcome measure and the probability of occurrence for each scenario.

In using risk analysis as a point of departure for the decision analysis to be carried out, there are (at least) three areas in which difficulties are

likely to be encountered. These are the evaluation of consequences for each of the different fire scenarios, treatment of the uncertainty in the probability estimates and in the estimates of consequences, and the choice of a decision rule. The choice of a decision rule is dependent upon the methods one elects to use in performing the other tasks, and it is also that aspect of decision analysis with which the present paper is most concerned. Because of its dependence on the other aspects of decision analysis, it will be taken up last.

3. Managing uncertainty

The first difficulty in connection with decision analysis to be discussed here is that of dealing with uncertainty regarding both the consequences and the probabilities associated with them. When probabilities are used to describe uncertainty, it is necessary to first define how the probability concept is to be conceived. The interpretation of probability with respect to risk analysis has been discussed in [2], and with respect to fire-risk analysis in [3]. Both authors involved suggest use of a subjective interpretation of probability, meaning that probability is regarded as a measure of degree of belief. In the present paper, the subjective interpretation of probability will be adopted. The reason for this is (1) that this interpretation is used in the Bayesian decision theory, which is the theory employed in the present paper and (2) it gives a flexibility to use other kinds of information than purely empirical information, such as expert judgement, for example.

In endeavouring to estimate the probabilities of each of the uncertain events that affect the outcome of a fire, it is often difficult to assign precise values to the probabilities in question. This is because one usually does not have sufficient information regarding any given probability to feel comfortable in expressing one's degree of belief as a single value. Instead, using a set of plausible values or an interval may seem more adequate. From a Bayesian point of view, uncertainty regarding a specific probability value is expressed as a probability distribution representing one's degree of belief regarding the different probability values (See [4], for example). In expressing one's belief regarding a particular probability as a probability distribution, one can use Bayes' theorem to incorporate new information into one's initial belief.

Bayesian methods can also be used to help the decision maker incorporate information from other sources than those of his/her own

judgements into the analysis, such as expert judgements or fire statistics (see, [5], for example). In many cases, this helps considerably in reducing the uncertainty (making the distribution less broad) regarding the value of the probability in question. Although a large reduction in uncertainty can be achieved by use of Bayesian methods, one still ends up with a distribution of probability values that one needs to somehow make use of in the decision analysis. How such probability distributions are dealt with in decision analysis is discussed in the section concerned with decision rules.

4. Evaluation of consequences

The consequences of a fire can be expressed in many different ways, such as the number of people whose health was affected by it, the value of the physical property that the fire destroyed, or whatever. In analysing different fire protection alternatives in a building that belongs to some particular organisation it is often convenient, assuming that all alternatives comply with the building codes, to endeavour to assess the damage due to a fire in terms of the intrinsic (negative) monetary value of the consequences as viewed by the decision maker. Methods of differing degrees of sophistication can be used to arrive at this intrinsic monetary value. One could use multi-attribute utility theory (see e.g. [6]) for example, to arrive at the intrinsic monetary value of each possible set of consequences, or one could settle for less formal models and simply try to evaluate the intrinsic monetary value for each fire scenario directly, without use of any formal approach to the problem. One reason for using intrinsic monetary values to obtain measures of relative preference for the different possible sets of consequences is that the monetary scale is one that people are accustomed to, its thus providing an effective means of communicating how good or bad the decision maker judges a particular outcome of a fire to be.

In practice, one needs to decide which losses that should be part of the evaluation. Obviously, monetary losses the decision maker is reimbursed for in case of fire should not be treated as losses in the decision analysis. However, one needs to be careful in considering the effects of insurance. Even with good insurance coverage there may be losses the decision maker will not be reimbursed for. In [7] it is indicated that only some 40-60% of the actual losses due to a disaster are covered by insurance. Although the amount of the losses covered by insurance obviously depends upon the building involved and the insurance covering it, it is

important to remember that considerable losses for the owner may occur, even if the building has good insurance coverage.

A term that can be used to denote all losses due to a fire, including losses due to business interruption, that the owner eventually has to defray is that of *uncompensated losses*. Such losses can include lost market shares, fines, negative reputation, and the like. It is very difficult to provide any general guidelines for the types of losses to be included in the calculations. Rather, that needs to be investigated in the specific case. Once the uncompensated losses have been identified, one needs to estimate their intrinsic monetary value. In doing so it is very likely, just as it was for the probabilities discussed above, that one will feel uncertain regarding the value to use. Instead of expressing the value as a precise number, it may be better to use an interval or a probability distribution to represent one's belief regarding the plausibility of the different values.

In working with practical applications, it is not always feasible to perform a complete analysis of uncompensated losses, since this could involve disproportionate work efforts in relation to the importance of the decision. It is useful, therefore, to distinguish between different levels of analysis, the level chosen depending on how thoroughly the uncompensated losses are to be investigated. A suggestion for how these levels of analysis can be defined is provided in [8]. As indicated there, an analysis of fire safety investment can be performed on at least three levels, that of (1) ignoring the increase in safety and of basing an 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 (and benefits) at level 1 and adding to this the valuation of the risk reduction achieved by using a subset of the uncompensated losses in the consequence estimations or any other losses for which the relation they have to the uncompensated losses can be assessed, or (3) taking all losses at level 1 into account and attempting to estimate all the uncompensated losses of importance.

Although which of the levels required depends on the problem at hand, it could be wise to start an analysis at level 1 and then increase the level of analysis if it is deemed necessary, since a higher level of analysis generally requires more work. A higher level of analysis tends to "favour" decision alternatives representing safety investments, since such investments generally decrease the probability of some of the fire

scenarios that have serious consequences and generally includes large uncompensated losses.

5. Decision rules

Having discussed some of the major problems and some of the possible ways of estimating the probabilities and consequences involved in a decision analysis, one needs to also consider the basis for evaluating the different decision alternatives.

In order to find a suitable decision rule, Bayesian decision theory will be examined to see whether that theory can prove useful in the present fire engineering context. The applicable decision rule for Bayesian decision theory is the principle of maximising expected utility. This is a principle that has been used extensively in the context of engineering (see [9], for example) and it has also been used in fire engineering (see [10], for example).

Modern decision theory has its roots in work performed by Ramsey [11], von Neumann and Morgenstern [12] and Savage [13], in particular. In these references, axiomatic systems for comparing preferences for different acts with uncertain outcomes have been formulated. The basic approach taken in constructing such axiomatic systems is to formulate a number of rules (axioms) that seem intuitively reasonable for comparing preferences between different acts with outcomes that are uncertain. From these axioms, a number of important results can then be derived, such as the principle of maximising expected utility (MEU). The MEU principle implies that a person who is willing to follow these axioms in his/her decision making will evaluate decision alternatives according to their expected utility and choose the decision alternative with the highest expected utility. Of the authors referred above, Savage has been called the principal founder of modern decision theory [14], which is also termed Bayesian decision theory. A review of the various theories of this type and of major aspects of modern decision theory have been provided by Fishburn [14].

Before discussing whether the MEU principle is reasonable to employ in the present context, it is useful to review some of the criticism that have been directed against axiomatic systems of the type that Savage proposes. The criticism that are discussed here are of two types: criticisms based on empirical investigations and criticisms directed against the logical foundations of the MEU principle.

In the first category, criticisms based on empirical investigations, the perhaps most famous criticisms of Bayesian decision theory are those made by Allais [15] and by Ellsberg [16]. Of these two authors, the one whose criticism is most relevant in the present context of fire safety would seem to be Ellsberg. Ellsberg's basic criticism is that in making choices between decision alternatives with uncertain outcomes, uncertainty regarding the probabilities and the value of the outcomes appears to influence how people choose. This type of uncertainty should, according to the Bayesian decision theory, not matter for a decision. According to this theory the uncertainty regarding probability values and consequence values should be expressed in terms of probability distributions representing the decision maker's belief regarding these values. In evaluating decision alternatives within a Bayesian framework uncertainty regarding the probabilities (how spread the distribution representing one's belief is) does not affect the decision, the only thing used in the evaluation of decision alternatives being the expected value (mean) of the distribution representing one's degree of belief. In the present context, nevertheless, it is desirable to be able to distinguish between situations in which a decision maker is very certain regarding his/her probability estimates and one in which he/she is not. For this reason, the term "robust decision" is introduced. Robust decisions will be discussed shortly.

The second category of criticism concerns the logical foundation of the MEU principle. Malmnäs [17] shows that the axiomatic systems proposed by Savage [13], among others, is too weak to imply the MEU principle. This is a serious criticism since it suggests that there are other decision criteria besides the MEU principle that satisfy the axioms and that the MEU principle is thus not a logical consequence of having accepted the axioms. Malmnäs undermines in this way one of the strongest arguments for using the MEU principle as a decision rule, namely that by accepting the axioms as rules for one's decision making one will then act as if one were evaluating decision alternatives according to their expected utility. In another paper, Malmnäs [18] examines the extent to which it is possible to provide the MEU principle support in a different way, that of showing that the rule does not give rise to counter-intuitive choices to any appreciable extent, counter-intuitive in the sense of a decision rule's evaluating an decision alternative with uncertain outcomes in a way not supported by human intuition. Although Allais [15], for example, has provided examples of situations in which the MEU principle generates counter-intuitive choices, Malmnäs

concludes that any simpler rule than the MEU principle gives rise to counter-intuitive choices to a greater extent than the MEU principle does and that "...the prospects for finding an evaluation [decision rule] that is much better than $E(A,f)$ [MEU] are not particularly bright."

As was indicated above, although considerable criticism has been directed against the MEU principle, this decision rule still appears to be a strong candidate for being a decision rule that can be used in connection with a quantitative risk analysis and at the same time is *practical* for use in the present context. The rule may possibly be in need of slight modification or require a complimentary evaluation of decision alternatives so that those alternatives involving uncertainty regarding the probabilities and consequences can be recognised. One way of doing this is to first evaluate all the decision alternatives using the MEU criterion, so as to find the decision alternative with the highest expected utility, which can be termed "the MEU alternative". When this decision alternative has been identified, it should be compared with the other decision alternatives in terms of the uncertainty connected with the estimates of probabilities and of consequences. One way of doing this would be to relate the uncertainty regarding the probabilities and consequences to the value of the expected utility. *Relating the uncertainty regarding the probabilities and consequences to the value of the expected utility involves the expected utility no longer being expressed as a single value but as a probability distribution.* Thus, comparing the MEU alternative with the other decision alternatives involves comparing probability distributions rather than precise values.

In comparing the alternatives in terms of the uncertainty connected with the estimates of probabilities and of consequences one is interested in the difference in expected utility. Since the expected utility of a decision alternative is expressed as a probability distribution, the difference in expected utility between two decision alternatives is also a probability distribution. Expressing the difference in expected utility in this way makes it possible to visualise the uncertainty regarding the value the difference has, and to take account of this in the decision to be made. If the *major part* of the mass of the probability distributions illustrating the difference in expected utility between the MEU alternative and the other decision alternatives indicates the MEU alternative to be best, then the decision is said to be *robust*, its otherwise being deemed *not robust*. What the "major part" in the above sentence means is up to the individual decision maker to decide. He/she might assume, for example,

that a decision is robust if 95% of the resulting distribution representing the difference in utility between two decision alternatives indicates the MEU alternative to be best. The concept of a *robust decision* is introduced here to provide an indication of how likely it is that the recommended decision alternative (the MEU alternative) will change if a plausible degree of change in the probabilities and the consequences should be made. To exemplify such an approach, consider a choice between three fire protection alternatives for which the uncertainty regarding the values of the probabilities and of the consequences in the model is expressed as distributions that represent the decision maker's belief regarding their values. Assume in addition that the result when calculating the expected utility of the different decision alternatives is that alternative 1 has the highest expected utility, followed by alternative 2 and alternative 3 in that order. Thus, according to Bayesian decision theory, alternative 1 is the decision alternative the decision maker should choose. Assume, however, that there is not much that differs between alternative 1 and alternative 2, and that in comparing the two decision alternatives in terms of the difference in the expected utility ($E(U_1) - E(U_2)$) and expressing the difference as a probability distribution, one can see that a slight change in the decision maker's belief could lead to alternative 2 being the best decision alternative, as shown in the distribution termed A in Fig. 1. In that figure, the area of the probability distribution to the left of the 0 value on the horizontal axis implies that alternative 2 is best, since $E(U_1) - E(U_2)$ is negative there. In this case, the decision to choose alternative 1 would probably not be considered to be robust since a large part of the probability distribution termed A in Fig. 1 implies that alternative 2 is best. In contrast, if one looks at the distribution representing the difference in expected utility between alternative 1 and alternative 3, as shown in the distribution termed B in Fig. 1, one notes that the situation is quite different. There, the whole probability distribution representing the difference in expected utility between the decision alternatives ($E(U_1) - E(U_3)$) is within the positive region on the horizontal scale, so that if alternative 1 and 3 are the only decision alternatives to choose between, deciding for alternative 1 would be considered a robust decision.




Fig. 1 Probability distribution representing the difference in expected utility between alternative 1, 2 and 3.

6. Summary of the approach

The approach suggested here for decision analysis concerned with investments in fire safety is based on the extension of Bayesian decision theory presented in the previous section. The treatment of uncertainty and the quantification of consequences were discussed in section 3 and 4.

A real-world decision analysis will be used to exemplify the approach taken. The analysis in question was conducted in 1998 at a firm called Asea Brown Boveri (ABB). It concerned the possible investment in a sprinkler system for a building belonging to the company. At the time, ABB was producing circuit cards for use in their robots and automation systems in the building. The building was approximately 55000 m² in size. Since the analysis was quite an extensive one, involving more than 150 different fire scenarios, only selected parts of it will be discussed. See Ref. 8 for a more comprehensive account.

The first step in conducting the type of decision analysis described here (see Fig. 2, step 1) is to identify the decision alternatives involved and decide upon the time period of concern. In the ABB case, there were only two decision alternatives: (1) keeping the building in its current state and (2) investing in a water sprinkler system for the building as a whole. The time period decided upon was one of 40 years. To determine whether choice of this particular time period had any effect on which alternative was deemed best, the same analysis was conducted for

periods of 5 years, 10, etc., its being concluded that the length of the time period had no effect.

The next step is to determine what level of analysis to employ (step 2 in Fig. 2). This involves deciding which losses are to be treated as uncompensated ones. In the ABB case, it was decided that the total costs of the equipment destroyed and of the interruption in business that a particular fire scenario entailed would be considered as uncompensated losses. Although ABB would later be reimbursed for the loss of equipment and for a part of the costs of the business interruption this sum was judged to be an appropriate measure of the total uncompensated losses as seen in monetary terms. The analysis as a whole was performed in accordance with the level 2 definition given in Section 4.

1. Define the alternatives and specify the time period of interest.
2. Choose the level of analysis (i.e. decide which losses to regard as uncompensated losses).
3. Determine the costs of each of the alternatives.
4. For each alternative:
 - Estimate the probability of the different fire scenarios.
 - Determine the uncompensated losses for each fire scenario.
 - Estimate the frequency of fires in the building.
 - Calculate the expected utility (or Certainty Equivalent) for the time period of interest.
5. Determine which alternative is the MEU alternative (i.e. the one with the highest expected utility).
6. Determine how the uncertainty of the probabilities and of the consequences is related to the difference between the alternatives in terms of expected utilities (or Certainty equivalent).
7. Determine robustness.

Fig. 2 The method of decision analysis suggested for decisions concerning investments in fire safety.

Since the sprinkler system was the only investment considered in the ABB case, only costs associated with that system needed to be included in the analysis. The sprinkler system was estimated to cost \$1,000,000 and annual maintenance of it \$10,000 (step 3 in Fig. 2).

The next step is to perform a risk analysis of each of the decision alternatives with the aim of identifying a set of fire scenarios, their respective probabilities of occurrence and their consequences in terms of uncompensated losses. This could be achieved, for example, by use of an event tree technique in which the uncertain events judged to affect the outcome of the fire are modelled. Which events to include in the event tree depends very much on the building at hand and the level of detail aimed at. In the ABB case, for example, events involving the sprinkler system and the fire detection system, as well as the building occupants and the fire department were used in the event trees. Evaluating the uncompensated losses involves estimating a monetary value that is seen as equal to each of the consequences. In the ABB case, this was accomplished by having the analyst explain a particular fire scenario, in terms of the extent of fire spread, to people from ABB and having them estimate the effect of such a fire in terms of uncompensated losses. Examples of uncompensated losses associated with some of the fire scenarios considered in the ABB case are given in Table 1. One can see that, so as to express the uncertainty involved, numbers are given there representing the most likely, the minimum and the maximum value respectively, for the consequences in question. These values are used to create triangular probability distributions to represent the decision maker's beliefs regarding the losses to be expected.

Table 1 The uncompensated losses associated with different fire scenarios in the ABB case. The fire scenarios apply to a fire compartment in which an electronic workshop is located.

Fire scenario	Uncompensated losses associated with the fire scenario in question (\$ thousand)		
	Minimum	Most likely	Maximum
A fire is extinguished either by employees or by the sprinkler system.	5	10	15
A fire of limited scope is extinguished by the fire department.	25	50	75
An extensive fire is extinguished by the fire department.	533	1067	1933
A fire completely destroys the fire compartment.	25920	32000	38720

The next step is to evaluate the different *decision alternatives*. The basis for doing so was discussed in the previous section, where it was concluded that the maximisation of expected utility is the decision rule applicable here. The expected utility of a decision alternative can be used to calculate the Certainty equivalent (CE) of it. The CE is a monetary sum equal in value to that of some particular situation involving uncertainty (see [19]). In this case, the CE is the (negative) monetary amount equal in value to choosing a particular fire protection alternative, including the costs of the alternative and the possibility of having one or more fires in the building during the time period of interest (see Fig. 3). The CE can be considered to be a better unit than “expected utility” for comparing decision alternatives, since it is expressed in terms of monetary value and people are more likely to feel comfortable using monetary sums than using expected utilities for comparison purposes. Note that whether expected utility or CE is used for comparing the decision alternatives should not affect the end result, the alternative being recommended being the same in both cases.

In order to calculate a CE here, one needs first to estimate how frequently fires will occur. Estimating the frequency of fires in the ABB building involved use of Bayesian methods, utilising the information that four fires altogether had occurred there during the years of 1996, 1997 and 1998. Using the estimate of fire frequency arrived at, together with

the event tree presenting the different fire scenarios, the costs of the sprinkler alternative, and the uncompensated losses associated with each fire scenario made it possible to estimate the CE for each of the two decision alternatives (step 4. in Fig. 2). This involves calculating the expected utility of one fire and then multiplying this value with the expected number of fires during the time period of interest. Note that one can regard losses occurring late in the period of interest as being less severe than those occurring earlier. This is discussed in detail in [20].




Fig. 3 Illustration of the calculation of the Certainty Equivalent of a decision alternative.

The results of the CE calculations indicated the decision alternative of investing in a sprinkler system to have the highest CE, and thus the highest expected utility (step 5. in Fig. 2).

In order to determine whether the decision of choosing the sprinkler alternative was robust, the question of how the knowledge uncertainty concerning the probabilities and the consequences was related to the difference in CE between the MEU alternative (the sprinkler alternative in the ABB example) and the other decision alternative was investigated. Since many of the probabilities and consequences used to calculate the CE (Expected utility) are uncertain, the CE is also uncertain. In the analysis, there were over 100 probabilities and consequences with a significant uncertainty regarding their values. One of those was the probability that the occupants would extinguish a fire in a particular storage room given that the smoke detection system functioned as intended. This conditional probability was estimated to be somewhere between 0,2 and 0,6, with a most likely value of 0,4. If the smoke detection system did not function as intended the probability was estimated to be somewhere between 0,1 and 0,3 with a most likely value

of 0,2. This kind of uncertainties was modelled using triangular distributions representing the probabilities when the calculation of the CE was performed. By use of Monte Carlo simulation (5000 iterations), the histogram presented in Fig. 4, showing the differences between the two decision alternatives in terms of CE could be obtained (step 6. in Fig. 2).

Figure 4 illustrates that the decision to invest in a sprinkler system for the ABB building is robust (compare Fig. 4 to Fig. 1) since all the values from the Monte Carlo simulation indicate that the sprinkler alternative has the highest CE (step 7. in Fig. 2). Figure 4 also shows that the difference between the two decision alternatives, in terms of CE, is substantial. The mean value of the difference is approximately \$3.1 million. Two reasons for the large difference in CE is that a serious fire in the building would cause significant losses for ABB (if the whole building is destroyed the uncompensated losses would be in the order of several hundred million dollars) and that the standard of the fire protection in the buildings original design were poor.




Fig. 4 The differences in CE in the ABB example. The CE of alternative 2 (not to invest in a sprinkler system) is subtracted from the CE of alternative 1 (to invest in a sprinkler system).

7. Conclusions

Decision analysis as applied to problems in which a decision maker is to decide between different fire protection alternatives for a particular building have been discussed.

In connection with estimating the possible consequences of a particular fire scenario, the concept of uncompensated losses was defined as the losses that the decision maker or organisation in question eventually have to defray. Since it is often practical to express such losses as monetary consequences, the intrinsic monetary value of the uncompensated losses generally needs to be estimated and to be used in the analysis.

The question of what decision rule should be used in the present context was discussed. Use of the criterion of maximising expected utility (MEU) being recommended. Various of the, major criticisms of use of this criterion in the present context were presented. In view of this criticism, it was considered to be advantageous to complement the MEU criterion with an evaluation of the robustness of the decision. A robust decision alternative was defined as an alternative that in terms of the MEU criterion remained the preferred one for most of the combinations of plausible probability and utility values that could be identified.

A real-world problem involving the evaluation of an investment in a water sprinkler system was presented. The building in which the investment was considered belonged to the company ABB. The investment in a water sprinkler system was found to be the best alternative. It was also concluded that the decision to invest in a sprinkler system was robust.

Acknowledgements

The author would like to thank The Swedish Fire Research Board (BRANDFORSK) for funding the research on which the present paper is based.

References

1. Hall, J.R., Sekizawa, A., "Fire Risk Analysis: General Conceptual Framework For Describing Models", *Fire Technology* 27: 33-53 (1991).
2. Apostolakis, G., "Probability and Risk Assessment: The Subjectivistic Viewpoint and Some Suggestions", *Nuclear Safety* 19: 305-315 (1978).
3. Noonan, F., Fitzgerald, R., "On the Role of Subjective Probabilities in Fire Risk Management Studies", *Fire Safety Science – Proceedings of the Third International Symposium*, International Association for Fire Safety Science, 1991, pp. 495-504.
4. Lindley, D.V., *Introduction to Probability and Statistics From a Bayesian Viewpoint, Part 2: Inference*, Cambridge University Press, Cambridge, 1965.
5. Sui, N., Apostolakis, G., "Uncertain Data and Expert Opinions in the Assessment of the Unavailability of Suppression Systems", *Fire Technology* 24: 138-163 (1988).
6. Keeney, R.L., Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, New York, 1976.
7. Hiles, A., "Marketing Protection: a justification for funding of Total Asset Protection programmes?", *The Definite Handbook of Business Continuity Management*, Hiles A., Barnes P. (eds.), John Wiley & Sons, Chichester, 1999, pp. 67-73.
8. Johansson, H., *Decision Making in Fire Risk Management*, Report 1022, Department of Fire Safety Engineering, Lund University, Lund 2001.
9. Benjamin, J.R., Cornell, C.A., *Probability, Statistics, and Decision for Civil Engineers*, McGraw-Hill Book Company, New York, 1970.
10. Shpilberg, D., De Neufville, R., "Best Choice of Fire Protection: An Airport Study", *Fire Technology* 10: 5-14 (1974).
11. Ramsey, F.P., "Truth and Probability", *The Foundations of Mathematics and Other Logical Essays*, Routledge and Kegan Paul, pp. 156-198, London, 1931. Reprinted in H. E. Kyburg and H. E. Smokler (Eds.), *Studies in subjective probability*, Wiley, New York, 1964, pp. 61-92.
12. Von Neumann, J., Morgenstern, O., *Theory of Games and Economic Behaviour*, Princeton University Press, NJ, 1944.
13. Savage, L.J., *The Foundations of Statistics*, Wiley, New York, 1954.
14. Fishburn, P.C., "Foundations of Decision Analysis: Along the Way", *Management Science* 35: 387-405 (1989).

15. Allais, M., “Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine”, *Econometrica* 21: 503-546 (1953).
16. Ellsberg, D., “Risk, ambiguity, and the Savage axioms”, *Quarterly Journal of Economics* 75: 643-669 (1961).
17. Malmnäs, P.-E., “Axiomatic Justifications of the Utility Principle: a Formal Investigation”, *Synthese* 99: 233-249 (1994).
18. Malmnäs, P.-E., *Methods of Evaluation in Supersoft Decision Theory*, Department of Philosophy, Stockholm University, 1995.
19. Clemen, R.T., *Making Hard Decisions – An Introduction to Decision Analysis*, 2nd ed., Duxbury Press, Belmont, 1996, p. 469.
20. Johansson, H., “Investment Appraisal Using Quantitative Risk Analysis”, *Journal of Hazardous Material*, Accepted for publication.

Paper 2: Investment appraisal using quantitative risk analysis

Johansson, H., *Journal of Hazardous Materials*, Vol. 93, pp. 77-91, 2002.

- The information contained in this paper can also be found in chapter 3.
- Note that equation (1) was incorrect in the original publication and has been revised here.

Investment appraisal using quantitative risk analysis

Henrik Johansson

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

Abstract

Investment appraisal concerned with investments in fire safety systems is discussed. Particular attention is directed at evaluating, in terms of Bayesian decision theory, the risk reduction that investment in a fire safety system involves. It is shown how the monetary value of the change from a building design without any specific fire protection system to one including such a system can be estimated by use of quantitative risk analysis, the results of which are expressed in terms of a Risk-adjusted net present value. This represents the intrinsic monetary value of investing in the fire safety system. The method suggested is exemplified by a case study performed in an Avesta Sheffield factory.

Keywords: *Decision analysis, risk analysis, investment appraisal, fire protection, Bayesian updating.*

1. Introduction

Making a decision of whether to install a particular fire protection measure in a factory can be difficult, particularly if the measure is not required for meeting the demands of the building code in question. In such a situation, a method is needed for comparing the benefits the fire protection measure would provide with the costs of investing in it. Decision-making problems of this type are traditionally solved using some capital investment method, e.g. net present value or rate of return, in order to calculate the profitability of the investment, and it would be beneficial if a similar method could be used in the present context.

How should such a traditional investment appraisal method be employed in the present context in a way allowing the reduction in risk that the investment implies to be taken into account? One way is to evaluate the risk reduction in terms of its *intrinsic* monetary value, treating it as “income” from the investment in question. Estimating the intrinsic monetary value of the risk reduction a specific fire safety investment provides can be based on the use of decision theory. This involves investigating the decision maker’s preferences towards risk, identifying

fire scenarios that are representative for the building in question, and employing some form of quantitative risk analysis (QRA) in which estimates of the probabilities and the consequences of the different fire scenarios that have been identified are made.

The present paper proceeds with a short presentation of decision theory in the context of decisions on fire protection measures and a discussion of how the uncertainties concerning the probabilities involved can be handled in a decision analysis. An account is provided of how one can model the frequency of fire as well as the different fire scenarios that can occur in a given building so as to be able to estimate the intrinsic monetary value of the risk reduction achieved by investing in a specific fire safety measure. How the uncertainty here can be reduced by the use of fire statistics is also taken up. The paper concludes with the presentation of a practical application of the suggested method in a case study involving investment in a sprinkler system in the cold-rolling mill of the Avesta Sheffield plant in Nyby, Sweden.

2. Decision analysis

In this section a brief account of decision analysis in the present context of fire safety is provided (for a more detailed description, see Johansson [1]). The concept of *certainty equivalent* will be considered in some detail because of its importance to the model for the investment appraisal of fire safety measures suggested here.

Modern decision theory has its roots in particular in the work performed by Ramsey [2], Von Neumann and Morgenstern [3] and Savage [4], who have developed axiomatic systems for comparing preferences for different acts with uncertain outcomes. The basic approach taken in constructing such axiomatic systems has been to formulate various rules (axioms) that seem intuitively reasonable for comparing preferences between different acts with uncertain outcomes. From these axioms, a number of important principles can then be derived, such as the principle of maximising expected utility (MEU). The MEU principle implies that a person who is willing to follow these axioms in his/her decision making will evaluate decision alternatives in terms of their expected utility and choose the alternative for which the expected utility is highest. Of the authors just referred to, Savage is the one often regarded as the principal founder of modern decision theory [5], or of what is also termed Bayesian decision theory. A review of various theories of this type has been provided by Fishburn [5].

Decision analysis, as described in a general way in [6], for example, involves the derivation of a utility function defined in terms of one or more attributes (such as monetary consequences, for example) that the decision maker wishes to take account of. The utility values obtained can be seen as measures of the decision makers' preferences, a consequence with a higher utility value being preferred to one with a lower value. Techniques for eliciting utility functions are summarised in [7].

When the decision maker's utility function has been determined one can calculate the expected utility of the different decision alternatives on the basis of the probabilities of the different consequences and their respective utility values. In the present context the word "disutility" might be considered more appropriate, since in most cases it is the utility of losses one is interested in. Nevertheless, the term utility will be used throughout. One should bear in mind, however, that it is usually a negative utility value that is meant when the expected utility of a fire is referred to.

It is important to note that the consequences of a fire are of a multi-attribute character. A serious fire can involve loss of sales, loss of market shares, getting a negative reputation, etc. Losses of this sort that the decision maker is not compensated for will be termed *uncompensated losses*. It can be useful to express these in terms of their *intrinsic* negative monetary value ([1], [8]). This allows measures of relative preference for the different possible sets of consequences to be obtained, and it gives the decision maker an effective means of communicating how good or bad he/she judges a particular outcome of a fire to be, since the monetary scale is one that people are accustomed to. The technique used for estimating the intrinsic negative monetary value of a specific loss can vary. It has been suggested that the technique adopted involve analysis at different levels that may differ considerably in the effort they require ([1], [8]). In the approach advocated in the present paper, no general evaluation of the different attributes is made, evaluations being performed instead on a scenario basis, the decision maker expressing his/her preferences within the framework of each fire scenario.

In the present context, calculating the monetary value for the decision maker of the reduction in risk that a particular fire safety investment involves is of interest. In carrying out a quantitative fire-risk analysis for a building, one estimates the probability of each of the possible consequences both before and after the investment under consideration

has been made, and expresses the consequences as utility values. This allows the expected utility, given that a fire has occurred in the building, to be calculated. This value, in turn, can be translated into a *certainty equivalent* (CE), which in the present case is the monetary value the decision maker is prepared to pay in order to escape the effects of an occurrence of fire in the building. A formal definition of CE is provided in equation (1), in which $u(CE)$ is the utility value corresponding to the monetary amount CE , $u(c_i)$ is the utility value corresponding to the consequence c_i , and n is the number of possible consequences. A general definition of CE is "...the amount of money that is equivalent in your mind to a given situation that involves uncertainty" [6]. Assume that CE has been calculated both for the alternative in which the building is equipped with the fire safety measure under consideration and for the building in its current state. If, in addition to this, one has an estimate of the annual frequency of fire (expected number of fires per year), one can also estimate, for any given time period, the intrinsic monetary value of the reduction of risk that the investment involves.

$$u(CE) = \sum_{i=1}^n p(c_i) \cdot u(c_i) \quad (1)$$

To illustrate how this can be done, assume that a decision maker has two alternatives to choose between for the fire protection to be found in a particular factory, the first alternative being to keep the factory in its present state and the second alternative being to invest in a certain type of fire safety measure. Both alternatives can be regarded as "lotteries". The difference between this situation and that of an ordinary lottery is that here the number of "drawings" is uncertain, in that the number of fires that will occur during the period which the analysis is concerned with is not known at the time of the decision and that in this "lottery" there are no prizes, only losses. Despite these dissimilarities, thinking of an alternative in terms of a "lottery" is helpful, although the term "exposure" can be considered more appropriate in the present context. A particular fire exposure is defined here as an uncertain situation in which the number of fires that will occur in the building (or whatever) in question during a specific period of time is unknown and the consequences of any particular fire is uncertain.

Although there is considerable uncertainty regarding the outcome of a certain type of exposure in a particular building, i.e. the number of fires

that will occur and their severity, it is possible to analyse the situation in such a way that exposures of different types (for example different building designs) can be compared. In evaluating different types of exposures, the concept of certainty equivalent (*CE*) is helpful. As explained above, *CE* is the monetary value a particular uncertain situation is seen to possess, which means that for a particular type of exposure for which *CE* is calculated the decision maker should be willing, in terms of Bayesian decision theory, to pay any amount that is less than *CE* in order to avoid that type of exposure.

The crucial question here is how much money the decision maker would be willing to pay in order to *change* his/her exposure from that which the current building design involves to the type of exposure that would result from the decision maker having invested in additional fire safety measures. This monetary amount can be assessed by calculating *CE* for each of the two types of exposure and determining the difference between them. This value then is the *intrinsic* monetary value of the risk reduction that the fire safety investment involves.

3. Time preference

The discussion of decision theory above has dealt with *risk* preferences, such as in connection with *certainty equivalents* and with preferences for different outcomes, as represented by the utility values of the possible consequences. There is one additional preference that is of importance in the present context, that of time preference.

A time preference can involve, for example, receiving a given sum of money today being regarded as better than receiving the same amount a year later. This is a matter dealt with by the methods for investment appraisal that are commonly employed such as the Present worth-method, the Annual worth-method, and the Future worth method (see Canada and White [9]). Time preferences are considered important in the present context, a method similar to the discounting technique as used in the Present worth-method, for example, being suggested for representing the decision maker's preferences regarding the time at which a consequence occurs.

In order to calculate the certainty equivalent (*CE*) for a certain type of exposure during a particular time period of interest, one needs first to calculate the expected utility associated with this exposure (see equation (1)). This requires that certain assumptions be made about the decision

maker's preferences regarding the occurrence of more than one fire during a given period of time. In particular, it is assumed that the expected utility of k fires during a given period of time, each of them with the expected utility $E(u)$, is $kE(u)$. This implies that the utility of any given fire not being affected by how many other fires occurred during the period in question.

The assumption just referred to enables one to calculate the expected utility of a particular type of fire exposure during a given time period j , its likewise being assumed that occurrence of the fires can be described by a Poisson process (see section 6). Equation (2) is used to calculate the expected utility ($E(u_j)$) for the type of fire exposure involved, λ being the frequency of fire (in fires per year), t_j the length of the time period considered (in years), $P(k)$ the probability of k fires occurring during this period, and $E(u)$ the expected utility of any given fire.

$$E(u_j) = \sum_k E(u) \cdot k \cdot P(k) = E(u) \cdot \sum_k k \cdot P(k) = E(u) \cdot \lambda \cdot t_j \quad (2)$$

As discussed above, discounting methods used in traditional investment appraisal (see [9], for example) are employed to take account of the time preferences. Such methods involve the loss of a particular monetary amount five years from now, for example, being seen as less severe than the loss of the same amount at present. The intrinsic monetary value (x) of a loss that occurs n years from now is assumed to be equal to a loss of $x/(1+i)^n$ today, i being the interest rate that corresponds to the decision maker's time preferences. Dividing the period of time which is planned for into shorter time periods enables one to discount to the present level the intrinsic monetary value of the consequences that occur during each of these time periods. Usually, time periods of one year each are employed. This means that the utility of a fire that causes losses having the intrinsic monetary value of x during the j th year of the period planned for is calculated by discounting x to the present and then calculating the utility of the discounted amount. Equation (3) is used to calculate the utility ($u(x_{s,j})$) of the loss ($x_{s,j}$), in the case of fire scenario s , occurring during the j th year of the time period of interest, i being the interest rate used to discount the monetary values in question.

$$u(x_{s,j}) = u\left(\frac{x_s}{(1+i)^j}\right) \quad (3)$$

The expected utility of a given type of exposure can be calculated for a particular period of time by use of equation (4), in which $E(u_E)$ is the expected utility of a particular type of exposure, n is the number of years considered and m is the number of fire scenarios taken account of in the building in question. The fire frequency here (λ) represents the expected number of fires per year, the time period (t_j) likewise being expressed in terms of years. When each time period is a year, the term t_j in equation (4) can be disregarded.

$$E(u_E) = \sum_j \left(\lambda \cdot t_j \cdot \sum_s u(x_{s,j}) \cdot P(x_{s,j}) \right) \quad \begin{matrix} j = 1, \dots, n \\ s = 1, \dots, m \end{matrix} \quad (4)$$

The equation indicates that the expected utility of the type of exposure that is considered is calculated by summarising the expected utility over the years in question to yield the expected utility for the period as a whole that is planned for. The expected utility can readily be translated then into the certainty equivalent, enabling the monetary value for the type of exposure in question to be calculated.

4. Uncertain estimates

Calculating the expected utility for a given type of exposure is not easy, however. The considerable uncertainty associated with the occurrence and spread of fire is a major reason for this. Various methods can be employed to deal with this uncertainty, quantitative risk analysis being a fruitful point of departure. In quantitative risk analysis based on the definition of risk proposed by Kaplan [10], one aims at specifying the accident scenarios that are representative for the building in question and at assessing their respective consequences and probabilities. In doing this, it is common to combine the probabilities of various events in an event tree, such as that “The sprinkler system succeeds in extinguishing the fire”, in such a way that the probability of a given fire scenario can be obtained.

Although uncertainty can be represented in ways other than by probabilities, such as by fuzzy measures [11], for example, use of probability measures seems to be the most fruitful approach in the present context [1]. In using probabilities to represent uncertainty, it is important to take account of the interpretation of probability that one explicitly or implicitly adopts. It has been argued that the subjective interpretation of probability is particularly useful in risk analysis [12].

Such an interpretation is the one adopted in the present paper. A subjective interpretation means a probability being regarded as a degree of belief in some proposition or event. Use of this interpretation provides considerable flexibility when a risk analysis is performed, and can also be considered as essential for the practical application of the methods suggested here.

In performing a quantitative risk analysis of a factory of some sort, one is very likely to feel uncertain about the estimates of various parameters, such as probabilities and frequencies. From a Bayesian standpoint, ambiguity regarding a probability or frequency estimate should be represented by a probability distribution defined over all possible values of the parameter in question (see [13], for example). An example of such a distribution will be given shortly. In Bayesian decision theory, however, ambiguity of this sort is assumed to not affect which decision alternative is best, or how much the decision maker should be willing to pay (the certainty equivalent) in order to avoid a particular type of exposure. According to that theory, the expected value of the distributions are the only values needed to determine the certainty equivalent of a given type of exposure. The author has argued [1], however, that in a context such as the present a Bayesian evaluation based on expected utilities is in need of being complemented by a further evaluation, one aimed at determining whether the choice of which decision alternative is best is *robust*. In brief, the concept of robustness implies that if a plausible degree of change in the assessment of the consequences and the probabilities is made, the alternative regarded as best will not change. The key to determining whether a decision alternative is robust is to relate the uncertainty of the probabilities and of the utilities of the consequences to the difference in expected utility between the decision alternative in question and the other alternatives.

To exemplify the approach suggested, consider a choice between three fire protection alternatives for which the uncertainty regarding the values of the probabilities and of the utilities of the consequences as assessed in the model can be expressed as distributions that represent the decision maker's belief regarding their values. Assume in addition that the result, when the expected utility of the different alternatives is calculated is that alternative 1 has the highest expected utility, followed by alternative 2 and alternative 3 in that order. According to Bayesian decision theory, alternative 1 is thus the alternative that the decision maker should choose. Assume, however, that the difference between alternative 1 and

alternative 2 is only slight. If one expresses the difference between the expected utility of the two alternatives as a probability distribution, it could look like the one in Fig. 1. One can see there that most of the mass of the probability distribution denoting the difference in probability ($E(U_1)-E(U_2)$) is located on the positive part of the horizontal axis, indicating alternative 1 to have the highest expected utility. However, there is also a significant part of the probability distribution located within the negative region, indicating alternative 2 to have the highest utility. This would imply, loosely speaking, that a reasonable change in the assessments of the probabilities and of the consequences could result in the alternative with the highest expected utility changing. This is a situation in which the alternative regarded as the best (alternative 1) is not deemed to be robust. If, on the other hand, the decision maker only had alternative 1 and alternative 3 to choose between, the choice of alternative 1 would likely have been considered robust, since if one looks at the distribution showing the difference in expected utility between alternative 1 and alternative 3 (the distribution illustrated in Fig. 2) one can see that the entire mass of the distribution is located in the positive region along the horizontal axis.




Fig. 1. Probability distribution representing the difference in expected utility between alternative 1 and 2.




Fig. 2. Probability distribution representing the difference in expected utility between alternative 1 and 3.

In practical applications of the method just discussed, a decision is generally deemed to be robust if 95% of the distribution representing the difference in utility between two alternatives indicates one and the same alternative to be best. This approach is only one that is recommended, its being up to the individual decision maker to choose a value that he/she feels comfortable with.

5. Fire scenarios and fire frequency

Through quantitative risk analysis, one can estimate the frequency of fire in a particular building and arrive at a plausible set of possible fire scenarios together with their respective conditional probabilities of occurrence (conditional on the event that a fire has occurred). The technique for doing this can vary considerably, the method described here being one found to be useful in two real-world analyses the author has carried out.

The basic idea of the method to be described is to divide up the building in question into suitable areas, preferably coinciding with the various fire compartments of the building. For each such area, a model of how a fire might develop needs to then be created. In the two real-world analyses referred to, an event tree technique was used to indicate the different fire scenarios that were considered suitable in the buildings in question and to calculate the conditional probability of each scenario, given that a fire had occurred in that area. In the event tree, different events that could

mitigate or affect in some other way the spread of a fire were included. The events can be considered roughly to be of five different kinds, those pertaining to fire potential, to employees, to active systems (such as sprinklers), to passive systems (such as fire compartments) and to the actions of the fire department. Fire potential concerns such matters as the fact that if a fire occurs in an area where the amount of combustible material is limited it might consume all the material there and be extinguished before causing any significant damage. All the relevant events that can mitigate or in any other way affect the development of a fire must be included in the event tree. Examples of such trees are given in [8].

The next step, after the model have been created, is to estimate the probabilities of the different events. As has already been indicated, these probabilities can sometimes be very difficult to estimate, particularly when there is only limited information about the events and the events are concerned with phenomena that are difficult to create models for. An example of such an event is “Those employed in the building succeed in extinguishing the fire”. Since the decision maker is likely to feel very uncertain in estimating probabilities of this sort it is advantageous to employ a decision analytic framework that allows probabilities to be expressed in an imprecise way. As will be shown in the next section, Bayesian methods can be used to reduce the uncertainty regarding the frequencies and probabilities considerably.

When the model for the development of a fire in the building is complete, one can create a list of all relevant fire scenarios, their consequences and conditional probabilities. Besides having the list described above, one needs to also have a model of how often a fire can be expected to occur, or of the frequency of fire. A good point of departure in estimating the frequency of fire in a building is to consider the results of investigations of the frequency of fire in buildings of different categories, as presented in for example Fontana et al. [14]. In this reference, estimates of the frequency of fire per square meter in buildings of various types are given. This information can help one arrive at an estimate for a specific building. Fire statistics from the building in question can also contribute to this. Bayesian methods for the incorporation of new evidence into estimates through use of Bayes’ theorem are useful here. This will be discussed and exemplified in the next section.

6. Bayesian updating

As already indicated, one easily feels uncertain about the value of a probability of an event that affect the outcome of a fire. Accordingly, instead of assigning a precise value to the probability in question, it may be better to employ a probability distribution to represent one's belief regarding the value the probability has. One benefit of doing this, besides its enabling the decision maker to express his/her uncertainty in a more adequate way, is that it enables information from different sources to be combined in estimates made by use of Bayesian methods. In the present context, such information can be information regarding a limited number of fires that have occurred in the building of interest, for example. Whereas this information alone is usually not sufficient to serve as a basis for estimating the different probabilities in the model, it becomes much more useful if combined with other sources of information, such as expert judgement and the like. How different types of information can be combined in this way in situations of different kinds that are likely to arise in a context such as the present one of decision analysis with respect to fire safety has been discussed by the author in [8]. Here, only one of these possible situations will be discussed, that of estimating the frequency of fire in a particular building.

The basic principles of Bayesian methods employed when incorporating information from different sources into a probability assessment have been dealt with in detail in [13]. Stated briefly, one begins with a *prior* probability distribution, one that represents the decision maker's belief regarding the uncertain parameter *before* any of the evidence has been taken into account. This prior distribution is updated then using the information in question, which could include information regarding a particular fire in the building, for example. This updated distribution, termed the *posterior* probability distribution, is obtained by use of Bayes' theorem.

In a case study carried out concerning a cold-rolling mill belonging to the company Avesta Sheffield, a study in which the methods discussed in this paper were employed, one of the uncertain parameters was the frequency of fire in the building. The only information available regarding this parameter was the number of fires that had occurred during the past six years. In order to obtain an estimate of the frequency of fire in the building, a so-called diffuse prior distribution was employed. This is a distribution that represents there being no strong belief in any particular value of the parameter in question. Although in

the case study both a discrete and a continuous prior distribution were used to represent the frequency of fire, only results involving use of the continuous distribution will be presented here. This distribution was chosen from the class of Gamma distributions since such distributions are flexible and are the conjugate family of distributions when the parameter of interest is the expected number of occurrences of some uncertain event, such as a fire, and when the number of events per time interval can be described by a Poisson distribution.

If fires can be assumed to occur independently of each other and to occur with a constant intensity, the Poisson distribution can be used to calculate the probability that some given number of fires will occur during a specified period of time. Since both these assumptions appeared reasonable the Poisson distribution was used to calculate the probability that a given number of fires would occur in the cold-rolling mill, given a particular frequency of fire in the building. The diffuse prior distribution employed was a Gamma distribution in which both parameters, s and m , were equal to 0. According to Lindley [13], this is the Gamma distribution to be used for representing vague prior knowledge.

The information contained in the fire statistics from the cold-rolling mill indicated that during a period of 6 years there had been 60 fires, which might be seen as many in a building of the type and size of the present one (see section 8 for a description). However, most of the fires were very small and were extinguished quickly. Many of them occurred in the machines, where they could be extinguished by automatic suppression systems. Nevertheless, all such incidents were counted in estimating the frequency of fire in the building. Using this information to update the diffuse prior distribution resulted in a posterior distribution that looked like the one shown in Fig. 3. It is a Gamma distribution with the parameters $s = 60$, and $m = 6$.

This distribution represents the decision maker's belief regarding the frequency of fire in the mill after the information contained in the fire statistics had been taken into account.

Looking at the figure showing the posterior distribution, one can draw the rough conclusion that the frequency of fire in the mill is likely to be somewhere between 6 and 14 fires per year.




Fig. 3. *The posterior distribution of the frequency of fire (λ) in the cold-rolling mill.*

One can use the posterior distribution shown in Fig. 3 as a prior distribution in a later updating procedure if additional evidence becomes available. Note that the updating procedure just exemplified can be used for all uncertain probabilities contained in the model of the different fire scenarios in any given building. The only difference as compared with the example just described is that instead of a Poisson distribution some other distribution might be needed, depending upon the information one makes use of. The author has discussed and exemplified some of the most common situations likely to be encountered in a context such as the present one in [8]. Note that Bayesian methods can also be employed in connection with expert judgements to incorporate them in a formal way into assessments of probability here. Thus, even without any statistical evidence regarding fire in the building of interest, uncertainty regarding the parameters of concern can be reduced by use of expert judgement.

7. Investment appraisal

The method used here for the investment appraisal of a fire safety investment is based on various of the methods discussed above. It is similar to the Present worth-method in that cash flows are discounted to the present, i.e. to the time when the decision is made, so that cash flows occurring at different times are comparable.

That which is aimed at is an estimate of the Risk-adjusted net present value of the investment, which in turn involves taking account of the risk

reduction that the investment provides, as well as costs of a more fixed or certain character, such as those for maintenance. The Risk-adjusted net present value being defined as the monetary value equivalent to making the investment in question.

Obtaining this estimate requires that the certainty equivalent of installing the fire protection system be estimated. As indicated above, this involves first choosing the period of time to which use of the fire protection measure is to apply and then performing a risk analysis in which the probability of each of the possible fire scenarios, as well as the consequences of each, are calculated, both for the building in its present state and when equipped with the fire safety system.

Performing a decision analysis using a risk analysis as a point of departure takes account of the decision maker's preferences with respect not only to the possible consequences of a fire, but also to risk in general, as well as to the occurrence of the consequences in question at differing times during the period that is planned for.

From this, finally, one can obtain the Risk-adjusted net present value of the investment, or the monetary value the investment has when both the intrinsic value of the risk reduction and more certain or fixed costs such as those of maintenance are taken into account.

8. Case study: Avesta Sheffield

In 1999 an analysis using the methods described above was performed for a cold-rolling mill belonging to the company Avesta Sheffield. The mill had a production capacity of approximately 100000 ton of stainless steel per year. The factory was approximately 15000 m² in size. The analysis that was performed concerned investment in a water sprinkler system for the entire building.

In carrying out the analysis, two event trees were created, each representing the fire scenarios that had been identified in the building. The two event trees represented the building with and without the sprinkler system, respectively. Using these event trees in combination with estimates of the probabilities of the different events in the event trees and assessments of the consequences resulting from the different fire scenarios enabled the intrinsic monetary value of the risk reduction that an investment in a sprinkler system would involve to be calculated.

Since an investigation of the company's risk tolerance was outside the scope of the study, assumptions had to be made regarding it. As it turned out, however, using a risk neutral utility function or an exponential utility function (signifying risk aversion), different values for risk tolerance being inserted into it, did not change the alternative found to be best. The results reported here were obtained using the risk-neutral utility function.

The probabilities and consequences employed in the model for fire spread were uncertain and were thus expressed as probability distributions (see section 4 and 6), which meant that the Risk-adjusted net present value was also uncertain. Monte Carlo-simulation involving 5000 iterations was employed for estimating the distribution describing the uncertainty regarding the Risk-adjusted net present value. The results of these iterations are shown in the histogram in Fig. 4. The two dashed lines in the figure represents the boundaries between which 90% of the values obtained by the Monte Carlo-simulation are located. The Risk-adjusted net present values at these boundaries are \$10.2 million and \$22.5 million, respectively. The monetary sums given in this section, originally in Swedish crowns (SEK), were converted to US dollars at the rate of \$1 to 10 SEK.




Fig. 4. *The probability distribution representing the Risk-adjusted net present value for the investment in a sprinkler system in the cold-rolling mill.*

In calculating the Risk-adjusted net present value, the costs taken into account were the initial investment costs, estimated to be \$250000, and

the annual maintenance costs, estimated to be \$5000 per year. In the primary analysis, it was assumed that a period of 40 years was planned for in connection with the sprinkler system and a discount rate of 20% was employed.

Since the basis for the evaluation is the Bayesian decision theory, the value that should be used to evaluate the investment is the expected value of the resulting distribution, which in this case means that a good approximation for this value would be the mean value of the result from the 5000 Monte Carlo-simulations. This value is \$15.6 million, which indicates that the investment is very “profitable” and should be made. No considerations of price changes were taken in the calculations.

To determine whether the decision was robust, it was decided that that would be the case if 95% of the Risk-adjusted net present values from the Monte Carlo-simulation indicated that the investment should be made. Looking at the histogram in Fig. 4, it is clear that the decision is robust since all of the values are on the positive region of the horizontal scale.

Although the planning period and the discount rate were provided by the decision maker (Avesta Sheffield), it was considered useful to perform a sensitivity analysis of these parameters. Fig. 5 presents the results of the sensitivity analysis for the period in question. The two dashed lines represent the boundaries between which approximately 90% of the resulting distribution of the Risk-adjusted net present value lies. The figure shows the Risk-adjusted net present value to be nearly the same for all periods longer than approximately 20 years. In the case of planning for shorter periods of time, the value is less but is still positive and robust (the lower dashed line is in the positive region along the vertical axis), implying that use of a shorter time period of interest would not affect the attractiveness of the investment.




Fig. 5. *The effect on the Risk-adjusted net present value of changing the time period of interest.*

The same type of sensitivity analysis was performed for the discount rate. It showed that if a lower discount rate than the one used in the primary analysis (20%) was employed, the Risk-adjusted net present value would be greater than in the primary analysis. If the discount rate was set at a level higher than 20%, the Risk-adjusted net present value was also positive for all the values employed (i.e., up to 50%). It was thus concluded that the result was stable with respect to the discount rate as well.

Note that the high Risk-adjusted net present value of the sprinkler system does not mean that the investment's "pay-back time" is short. It is possible, though unlikely, that the sprinkler system will never need to be used during the 40-year period that is planned for. In such a case, of course, the investment would be a bad one because of never having been needed. Since when one makes the decision, however, it is impossible to know whether the system will be needed or not, one has to rely on estimates of fire frequency and on the modelling of fire spread. The model described in this paper gives the result that the sprinkler can be expected to be very useful as a risk reducing measure in the Avesta Sheffield building. Since a large fire can cause considerable losses (in the order of hundreds of million dollar), lowering the probability of such a consequence slightly has a large affect on the result. This is why the Risk-adjusted net present value is so high.

In comparing the investment appraisal described here with a similar one performed in a factory belonging to the company Asea Brown Boveri (ABB) one can note that investment in the sprinkler system located in the present case in a cold-rolling mill is more “profitable” than the investment was in the case of the ABB factory [8]. For the ABB building the Risk-adjusted net present value was calculated (by use of the same method) to be \$3.1 million. This difference in Risk-adjusted net present value can be explained by the fact that the passive fire protection in the cold-rolling mill (fire-rated walls, etc.) is not as good as that in the ABB building. This means that if a fire grow large in a fire compartment it is more likely to spread to other compartments in the cold-rolling mill than in the ABB building. In addition, the losses associated with fires are smaller in the ABB building than in the cold-rolling mill, and in the ABB building there are other kinds of fire protection systems (automatic smoke detection in the entire building, for example) which the cold-rolling mill does not possess. Because of these differences, the relative increase in safety which investing in a sprinkler system would provide is greater for the Avesta Sheffield than for the ABB building.

9. Conclusions

Use of quantitative risk analysis for the appraisal of fire safety investments, using methods based on Bayesian decision theory, has been discussed. Particular attention has been directed at the problem of evaluating losses due to fires that occur at different times, use of a method similar to that of the discounting of cash flows being suggested for modelling the decision maker’s time preferences. Taking account of the decision maker’s time preferences, risk preferences, and preferences regarding various monetary consequences of fire, as shown by the corresponding utility functions, enables the Risk-adjusted net present value of an investment in fire safety, or the assessed monetary value of having made the investment, to be calculated. If and only if this latter value is positive, should the investment be made.

Calculating the Risk-adjusted net present value of an investment in fire safety is based on quantitative risk analysis. Since many of the probabilities used here are uncertain, the Risk-adjusted net present value obtained is uncertain. According to Bayesian decision theory, uncertainty of this sort should not affect the decision made. In the present context it has been judged to be beneficial, however, to relate the uncertainty regarding both the probabilities and the consequences to the Risk-adjusted net present value, the latter being represented as a

probability distribution, so as to indicate how certain the decision maker is regarding the Risk-adjusted net present value of the investment.

A method of reducing the uncertainty regarding the occurrence and spread of fire is dealt with in the paper. The method employs Bayesian methods for integrating specific information concerning the building of interest with other types of information, such as expert judgement, general statistics, and the like.

A real-world problem dealing with investment in a water sprinkler system for a cold-rolling mill was also analysed, showing in practical terms how the approach described can be applied and how the results can be presented in a meaningful way to the decision maker.

Acknowledgements

The author would like to thank The Swedish Fire Research Board (BRANDFORSK) for funding the research on which the present paper is based.

References

- [1] H. Johansson, Decision Analysis Concerned with Investments in Fire Safety, 7th International Symposium on Fire Safety Science, Worcester, USA, 2002.
- [2] F.P. Ramsey, in: R. Braithwaite (Ed.), The Foundations of Mathematics and Other Logical Essays, Routledge and Kegan Paul, London, 1931, p. 156, Reprinted in: H.E. Kyburg and H.E. Smokler (Ed.), Studies in subjective probability, Wiley, New York, 1964, p. 61.
- [3] J. Von Neumann and O. Morgenstern, Theory of Games and Economic Behaviour, Princeton University Press, Princeton, 1944.
- [4] L.J. Savage, The Foundations of Statistics, Wiley, New York, 1954.
- [5] P.C. Fishburn, Management Science, Vol. 35 (1989) No. 4.
- [6] R.T. Clemen, Making Hard Decisions – An Introduction to Decision Analysis, 2nd edn., Duxbury Press, Belmont, 1996.
- [7] P.H. Farquhar, Management Science, Vol. 30 (1984) No. 11.
- [8] H. Johansson, Decision Making in Fire Risk Management, Department of Fire Safety Engineering, Lund University, Lund, 2001.

- [9] J.R. Canada and J.A. White, Capital Investment Decision Analysis for management and engineering, Prentice-Hall, Englewood Cliffs, NJ, 1980.
- [10] S. Kaplan, Risk Analysis, Vol. 17 (1997) No. 4.
- [11] G.J. Klir and T.A. Folger, Fuzzy Sets, Uncertainty, and Information, Prentice-Hall, Englewood Cliffs, New Jersey, 1988.
- [12] G. Apostolakis, Nuclear Safety, Vol. 19 (1978) No. 3.
- [13] D.V. Lindley, Introduction to Probability and Statistics From a Bayesian Viewpoint, Part 2: Inference, Cambridge University Press, Cambridge, 1965.
- [14] M. Fontana, J.P. Favre, C. Fetz, Fire Safety Journal, 32 (1999).

Paper 3: Application of Supersoft decision theory in fire risk assessment

Johansson, H., Malmnäs, P-E., *Accepted for publication in Journal of Fire Protection Engineering.*

Application of Supersoft Decision Theory in Fire Risk Assessment

Henrik Johansson

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

Per-Erik Malmnäs

Department of Philosophy, Stockholm University
SE-106 91 Stockholm, Sweden

Abstract

The application of Supersoft decision theory (SSD) to fire safety problems, and of decision analysis in general to decisions involving a high degree of epistemic uncertainty, are discussed. SSD and two traditional decision analytic methods employed earlier within the context of fire engineering are compared, particularly regarding how uncertainties are dealt with and the robustness of decisions - robustness concerning the likelihood that the alternative adjudged to be best will change when a reasonable degree of change in assessments of either the probabilities or the utilities involved occurs. Substantial differences between the three methods in decision robustness were noted. It was found that, since traditional decision analysis involving precise probability and utility values gives no indication of robustness, it can lead to wrong conclusions, making it unsuitable in the present context. It is argued that methods not providing the decision maker with information on decision robustness are unsuitable in situations involving a high degree of epistemic uncertainty. A procedure involving use of Supersoft decision theory and extended decision analysis for decision problems involving choice between different fire protection alternatives in a concrete case involving a specific building is suggested.

Keywords: *Decision analysis, fire risk analysis,
epistemic uncertainty, Supersoft decision theory*

Introduction

Performing a quantitative fire risk analysis for a particular building involves dealing with various uncertainties. First, one needs to address the uncertainty regarding the outcome of a fire, which is usually done by constructing a model of various fire scenarios, for example by use of

event trees. Uncertainty of this type, termed here aleatory uncertainty, has also been conceptualized in terms of irreducible uncertainty, inherent uncertainty, variability or stochastic uncertainty [1]. Secondly, one has to deal with uncertainty regarding the values of the variables used in the model of the different fire scenarios, the probability values, for example. Uncertainty of this type, based on lack of knowledge or information, is termed here epistemic uncertainty. It has been conceptualised as well in terms of reducible uncertainty, subjective uncertainty or cognitive uncertainty [1]. The present paper is concerned primarily with epistemic uncertainty in decision analysis concerned with investments in fire safety.

Note that the focus here is not on decision analysis concerning a category of buildings, such as, for example, analysis of strategies for reducing residential fire loss generally (see [2], for example). Instead, decision analysis of potential investments in fire safety for a *specific* building is of concern. The difference between analysing decisions for a category of buildings and doing so for a specific building is usually substantial in terms of the amount of information available regarding various parameters of interest. One usually has some information about fires that have already occurred in buildings of a particular category, whereas one may very likely have little or no information about any previous fires in a specific building. Because of this, a decision analysis pertaining to a specific building is likely to involve a high degree of epistemic uncertainty, especially as regards extreme or catastrophic events which even in a large group of buildings occur very seldom. The question is how this uncertainty will affect a decision analysis in a specific building and what method or methods can be used to deal with the large epistemic uncertainties involved. In discussing the analysis of such uncertain decision situations here, an application of a decision analysis method called Supersoft Decision Theory (SSD) is presented. This is a method specifically designed to deal with decision situations involving large epistemic uncertainties. The major aims of the paper are to present the conceptual framework of SSD, show how the SSD method can be applied to problems of fire safety, compare SSD with two alternative methods for decision analysis concerned with fire safety employed earlier and provide some general suggestions on how to evaluate decision situations in which an extraordinary degree of epistemic uncertainty exists.

The paper begins with a brief discussion of different criteria that could be useful in evaluating various alternatives in decision analyses concerned with investments in fire safety in a particular building. The SSD method, its theoretical framework and how it can be applied within the context of fire safety engineering are then taken up. The paper continues with a presentation of two examples of how SSD can be used to analyse decision problems concerned with fire safety, each involving a choice between different fire protection alternatives for a particular building. The first example concerns in a basic way the use of SSD, whereas the second example aims at clarifying differences between the use of SSD and of more traditional decision analysis methods in this context. The paper concludes with a general discussion of decision analysis involving a high degree of epistemic uncertainty.

Decision analysis concerned with investments in fire safety

In decision analysis one distinguishes between decision making under *risk* and decision making under *uncertainty*. Decision making under risk is characterised by the decision maker's knowing the probabilities of the outcomes of the various decision alternatives exactly, whereas decision making under uncertainty involves the decision maker's having no information at all about the probabilities of the different possible outcomes. Thus, in terms of epistemic uncertainty, decisions under risk involve no epistemic uncertainty whereas decisions under uncertainty involve the maximum epistemic uncertainty possible.

The most common decision criterion in making decisions under risk is the principle of maximising expected utility (MEU), which has been applied to fire safety problems earlier (see [3] and [4], for example). In making decisions under uncertainty, there are a number of decision criteria one could choose between. Donegan [5] discusses four such criteria: the Laplace paradigm [6], the Wald paradigm [7], the Savage paradigm [8] and the Hurwicz paradigm [9]. Since it is assumed that the decision maker is completely ignorant with respect to the probabilities of the various outcomes that are possible for the different decision alternatives, each of these criteria involve some form of valuation of the outcomes themselves. In the present context, the decision rule suggested by Laplace implies that the decision maker should choose the decision alternative that minimises the expected loss, its being assumed that each outcome considered is equally likely to occur. The Wald paradigm involves choosing the alternative for which the loss in case the worst outcome occurs will be lowest (this rule is also called the Maximin rule).

The decision rule Savage suggested involves choosing the alternative that would result in the lowest loss possible if the best outcome should occur (this rule is also called the Maximax rule). The Hurwicz decision rule is a combination of the Maximax and Maximin rules. In choosing between different fire protection alternatives for a specific building, none of these decision rules can be considered suitable, however, since they ignore any differences between the alternatives in terms of the probabilities of the different consequences. Since investing in fire protection aims in part at reducing the probability of a serious fire, the benefits of such an investment are not taken into account by any of the decision rules just referred to. An investment in a sprinkler system, for example, reduces the probability of a serious fire but does not reduce the negative effects of the worst possible *consequence*, namely the complete destruction of the building. Thus, use of the Maximin rule, which simply focuses on the worst possible consequence, would never lead to the recommendation that one makes a fire-safety investment.

The problem of performing a decision analysis concerned with alternative designs for fire protection in a particular building is likely to lie somewhere between decision making under risk and decision making under uncertainty.

Supersoft Decision Theory (SSD) was chosen for use in the present, fire-engineering context because of its readily being used in conjunction with a quantitative risk analysis (event trees are used in the paper), and also because its enabling one to compare in a clear way the results obtained with the results of a more conventional decision analysis. Both of these more conventional decision analysis methods are based on Bayesian decision theory, which involves use of the principle of maximising expected utility as the decision rule. One of these two methods will be termed traditional decision analysis. It involves probabilities and utilities being assigned as precise values. Use of this method allows different decision alternatives to be compared on the basis of expected utilities, the alternative having the highest expected utility being the alternative deemed best. Thus, only *one* value, the expected utility, is used to compare the different decision alternatives. That value (V_T) can be calculated using equation (1), in which n is the number of possible outcomes of choosing a specific decision alternative that have been identified, P_i is the probability of outcome i occurring, and U_i is the utility associated with the occurrence of the outcome in question.

$$V_T = \sum_{i=1}^n (P_i \cdot U_i) \quad (1)$$

The other method of more conventional character is termed extended decision analysis. It involves probabilities and utilities being expressed as probability distributions. In comparing different decision alternatives, it is the expected utilities that are compared, the alternative with the highest expected utility being deemed best. This is almost the same decision rule as that employed in traditional decision analysis, the difference being that in extended decision analysis calculating the expected utility requires taking account of the epistemic uncertainty regarding the probability and utility values. The value (V_E) employed in comparing one alternative with another by use of extended decision analysis can be calculated using equation (2). There, $f_i(P_i)$ is the probability density function representing the epistemic uncertainty regarding the probability value P_i , and $g_i(U_i)$ is the probability density function representing the epistemic uncertainty regarding the utility value U_i .

$$V_E = \int \dots \int_D \left(\sum_{i=1}^n (P_i \cdot U_i) \cdot f_1(P_1) \cdot \dots \cdot f_n(P_n) \cdot g_1(U_1) \cdot \dots \cdot g_n(U_n) \right) dP_1 \dots dP_n dU_1 \dots dU_n \quad (2)$$

$$D = \left\{ (P_1, \dots, P_n, U_1, \dots, U_n); \sum_{i=1}^n P_i = 1 \right\}$$

In addition, however, extended decision analysis also involves evaluation of the effect which the epistemic uncertainties have on the expected utility. The idea here is that, since probabilities and utilities are expressed as probability distributions that represent the degree of confidence one has in different values for these, one can also express one's degree of confidence in different expected utility values. Thus, one can relate the epistemic uncertainty pertaining to probabilities and utilities to the expected utility of a decision alternative, which can be expressed as a probability distribution. This probability distribution can be used then to compare different decision alternatives in terms of decision robustness. Robustness has to do with how likely it is that the decision alternative considered best will change if the estimates of the probabilities and the utility values should change. Decision robustness is one of the key topics in this paper and will be taken up shortly.

Note that extended decision analysis can be viewed as being in many respects equivalent to what is termed Bayesian analysis (see [10], for example). In Bayesian analysis, epistemic uncertainties are represented by probability distributions and decision alternatives are evaluated on the basis of the expected utility (see equation (2)). Bayesian analysis differs from extended decision analysis in only the value of the expected utility being used in the evaluation of decision alternatives, whereas in extended decision analysis the effect which epistemic uncertainties have on the expected utilities, and thus the robustness of the decision, also being taken into account.

All three methods considered above (traditional decision analysis, extended decision analysis and Supersoft decision analysis) utilise evaluation of expected utilities in one way or another. The use in the present context of expected utility for the evaluation of decision alternatives seems reasonable in view of results that Malmnäs [11] has presented. Malmnäs concludes that any rule that is simpler than that of expected utility performs worse as an evaluator of uncertain decision alternatives than expected utility does. More advanced methods¹ cannot be expected to be substantially better than use of expected utilities, especially when decision situations are involved in which probabilities and utilities cannot be expressed precisely.

Having decided to use expected utilities in evaluating uncertain decision situations leads to the question of how the expected utilities of decision alternatives should be calculated. To do so, one needs to be able to estimate the probability of each possible outcome of an alternative, as well as to assign a utility value to each of the outcomes in such a way that the utility value arrived at represents the decision maker's preference for the (uncertain) outcome in question. The first problem one encounters here is that of estimating the probabilities of such events as whether the employees in the building will succeed in extinguishing a fire, whether the fire department will succeed in extinguishing it, and the like. Since observations of events of this type are usually rare in any particular building, information regarding them is likely to be scarce, making estimates of the probabilities involved difficult to make. There are different approaches one can take in dealing with the problem of

¹ The decision rules suggested by Hagen [19], Fishburn [20], Loomes and Sugden [21], Green and Jullien [22], Quiggin [23], and Yaari [24] are evaluated in [11].

probabilities involved being based on such limited information. The three methods referred to above will be compared in this respect. The first method to consider is that of traditional decision analysis. It involves use of precise values for probabilities and utilities, and consequently of exact values for the expected utilities. In comparing two decision alternatives using this method one compares two expected utilities, one for each alternative, the alternative with the highest value being regarded as best (see example 1 in Figure 1). When dealing with problems of the present type, however, it is questionable whether expressing estimates of probabilities as precise values is suitable. This method is taken up in the paper nevertheless in order to compare it with the other methods. Note that it provides the decision maker no information at all on how any uncertainties regarding the probability or utility values affect the decision. The second method to be considered, that of extended decision analysis, involves expressing the uncertainty one has concerning the probability and utility values by use of probability distributions. In comparing two alternatives by use of this method, one compares two distributions of expected utilities, the distribution with the highest expected value representing the decision alternative that is best (see example 2 in Figure 1). Although the decision rule employed for determining the best alternative involves use of only the expected value of each of the two distributions, the form and position of the distributions provides the decision maker information regarding the robustness of the decision, which the first method does not. This method, is described in greater detail in [4]. Note that in analysing different investments for a specific building, it may not be possible to use specific probability distributions to represent probability and utility values due to the lack of information. Therefore, one may need to employ a method that can deal with a high degree of epistemic uncertainty regarding the probabilities and utilities. The third method, which is called Supersoft Decision Theory (SSD) [12] and is described in greater detail in the next section, is a method able to do this. In comparing alternatives by use of SSD, one compares the maximum and the minimum values of the expected utilities. Besides these two values, one should also take account of a value termed the Average, as will be described in the next section. The maximum and the minimum value form an interval in which the value of the expected utility lies (see example 3 in Figure 1).




Figure 1 Illustration of the results each of three different decision analysis methods provides in comparing two decision alternatives. $E(U_1)$ is the expected utility of alternative 1 and $E(U_2)$ the expected utility of alternative 2.

Later in the paper, the three methods described above will be applied to the same decision problem, the results of the analyses of that decision problem providing an additional illustration of the difference between them.

Supersoft Decision Theory

Supersoft Decision Theory (SSD) [12] allows the decision maker to utilise vague assessments of the values of the probabilities and consequences of interest. “The probability must be somewhere between 0.2 and 0.8” and “The consequence c_1 is at least twice as good as the consequence c_2 ” are examples of such vague assessments. Vague expressions of this sort are interpreted as inequalities. Thus, the representation of the probability just referred to could be $0.2 < p < 0.8$. Even when utilising such imprecise statements, one can still make use of

the same basic model for how a fire in a building can be expected to develop as one does in performing a quantitative risk assessment. The event tree technique, which is useful for modelling possible fire scenarios in a building, will be used here for exemplifying how SSD can be employed for evaluating different fire protection alternatives.

In evaluating a decision situation in terms of SSD, one needs to create a representation of it in terms of a *decision frame*. This consists of the following: the different alternatives that can be chosen (a_1, \dots, a_n), a list of the possible consequences C_i for each alternative, a list of utility statements U_i that pertain to these consequences, and a list of conditional probability statements P_i . The items on the list of consequences could be of the type “Areas 1 and 2 are completely destroyed” and those on the list of utilities could be of the type “The utility of consequence 1 is at least 20 times as high as the utility of consequence 2”. The items on the list of probabilities can be statements of the type “The probability of event 2 is highly likely given event 1”. The event trees (T_1, T_2 , etc.), which indicate how the uncertain events are connected with the consequences, represent the last component of the decision frame. Thus, the decision frame can be summarised as consisting of (a_i, C_i, P_i, U_i, T_i). In practice, one should start by clarifying which alternatives are possible to choose between and then to identify for each of the alternatives the various events that can influence the outcome of the decision. The relationship between the occurrence of these events and the consequences should then be described (this can be done by using event trees) and the probability statements for the different events be formulated.

To evaluate the different alternatives, so as to identify which one is best, the qualitative statements of the decision frame need to be transformed into quantitative ones. A qualitative statement of the type “Event 1 is highly likely, given event 2” can be translated, for example, into a quantitative statement of the type “ $0.85 \leq P(E_1|E_2) \leq 0.95$ ”. Note that SSD does not prescribe any rules for how qualitative statements regarding probabilities are to be translated into intervals. Considering the empirical evidence (see [13], for example) indicating a great between-subject variability in the probability values assigned to verbal statements, use of such fixed transformation rules is probably not a very good idea. Using the verbal statements here only as points of departure in determining the intervals of the probabilities can be suggested instead. In doing this, the analyst and the decision maker needs to work together to

find suitable intervals to represent the verbal statements made. Note that in determining the intervals of the probabilities, the decision maker is asked to *exclude* probability and utility values he/she considers too unlikely to be worth considered. This makes SSD different from other types of probability estimation techniques, in which either a single probability value (examples of such methods are provided in [14]) or a single interval containing the most likely values is to be estimated. The analyst might ask a decision maker who states “Event 1 is highly likely, given Event 2” whether it would be possible to exclude probability values of less than 0.05 for this conditional event. If this seems reasonable to the decision maker, one can continue and ask him/her whether it would be reasonable to exclude values of less than 0.1 and so on. In the end, an interval is established, that can be used to represent as adequately as possible the probability in question. For a more comprehensive discussion of this matter, see [12] and [16].

Examples of the types of statements that can be employed and of the respective inequalities are shown in Table 1.

Table 1 Applicable statements and their corresponding mathematical representations.

Statement	Representation
The probability of event E is equal to the probability of event F .	$P(E) = P(F)$
The probability of event E is less than x .	$P(E) < x$
The probability of event E is greater than x .	$P(E) > x$
The probability of event E lies between x and y .	$y < P(E) < x$
The probability of event E is at least i times as probable as event F .	$i * P(E) > P(F)$
The utility of consequence c_1 is higher than that of consequence c_2 .	$U(c_1) > U(c_2)$
The utility of consequence c_1 is at least i times as high as that of consequence c_2 .	$U(c_1) > i * U(c_2)$
The utility of consequence c_1 is equal to that of consequence c_2 .	$U(c_1) = U(c_2)$

Evaluation of alternatives in Supersoft Decision Theory

In employing SSD to evaluate the different decision alternatives, use is made of their expected utilities. However, since the decision frame contains statements in which the probabilities and the utilities are not assigned precise values or single probability distributions, the decision

criterion of maximising expected utility cannot be employed directly. Instead, the evaluation of alternatives in SSD is based on three criteria presented in equations (3) to (5). $E(U,P)$ is the expected utility of the alternative in question (which is a function of the probabilities, P , and the utilities of the consequences, U).

$$\text{Min}(E(U,P)) = \text{Min}_{P,U} \left(\sum_{i=1}^n (P_i \cdot U_i) \right) \quad (3)$$

$$\text{Max}(E(U,P)) = \text{Max}_{P,U} \left(\sum_{i=1}^n (P_i \cdot U_i) \right) \quad (4)$$

$$\text{Average}(E(U,P)) = \frac{\int \dots \int \left(\sum_{i=1}^n (P_i \cdot U_i) \right) dP_1 \dots dP_n dU_1 \dots dU_n}{\int \dots \int 1 dP_1 \dots dP_n dU_1 \dots dU_n} \quad (5)$$

The $\text{Min}(E(U,P))$ and $\text{Max}(E(U,P))$ criteria are the lowest and highest expected utility values that satisfy the decision frame. Satisfying the decision frame means that a solution to the inequalities is found within the decision frame. For example, assume that in the decision frame it is stated that the probability of a particular sprinkler system extinguishing a fire is somewhere between 0.8 and 0.9 ($0.8 \leq P(\text{Sprinkler}) \leq 0.9$), and also that this probability is at least 2 times as great as the probability that the employees will extinguish the fire ($2 \cdot P(\text{Employee}) \leq P(\text{Sprinkler})$). If the probability that the employees will extinguish the fire were 0.6, there would be no solution within the decision frame, since this implies that the probability of the sprinkler system's extinguishing the fire would be higher than 1, which is impossible. However, if the probability that the employees will extinguish the fire were 0.42, then the decision frame would be satisfiable, since this implies that the probability that the sprinkler system will extinguish the fire is greater than 0.84, which satisfies the inequality $0.8 \leq P(\text{Sprinkler}) \leq 0.9$. Throughout the paper, Min will be used to denote $\text{Min}(E(U,P))$ and Max to denote $\text{Max}(E(U,P))$.

For simple problems, the calculation of Min and Max is not very complicated. One such simple problem is illustrated in Figure 2. The problem involves there being two possible outcomes if a particular

decision alternative is chosen, one of them with a utility value of between 0.9 and 1 (u_1) and the other with a utility value between 0 and 0.5 (u_2). The best consequence (u_1) occurs with the probability p_1 , which is judged to be somewhere between 0.85 and 0.95, and the other consequence with the probability $1-p_1$.




Figure 2 Illustration of an uncertain situation.

In order to evaluate the decision alternative shown in Figure 2, one needs to analyse the expression for the expected utility ($E(U,P)$), as given in equation (6). The *Min*-value of equation (6), given the constraints (O, in Figure 2), is found by setting p_1 , u_1 and u_2 to their lowest values. The resulting value is 0.765. The *Max*-value is calculated using the same procedure but setting the parameters (p_1 , u_1 and u_2) at their highest values. This yields a value of 0.975.

$$E(U, P) = p_1 \cdot u_1 + (1 - p_1) \cdot u_2 \quad (6)$$

Although calculating the *Max* and *Min*-values may be a simple task in the case of such limited problems as that illustrated in Figure 2, as the scope of the problem increases the expression for the expected utility becomes more complicated, making the calculations much more complex. The complexity is due to the fact that the problem of calculating *Min* and *Max* is a nonlinear multivariable optimisation problem with a set of inequalities as constraints. $E(U,P)$ is the objective function that one seeks to minimise or maximise, the inequalities found in the decision frame represents the constraints. Such a problem, except for one of the simplest type, is difficult to solve by hand, but there are computer programs that can solve them².

The *Max* and *Min*-criteria alone are not ideal for comparing decision alternatives, since they are very sensitive to changes in values near the

² In order to solve the first example that is found later in the present paper, the “fmincon” function in the Optimization Toolbox for MATLAB [17] was used.

edges of the decision frame, these being the values the decision maker is most likely to be uncertain about. It is useful, therefore, to also employ the $Average(E(U,P))$ -criterion shown in equation (5). The $Average(E(U,P))$ -criterion can be seen as the expected value of $E(U,P)$ when the probability and utility values are treated as being uniform distributions that extend between their maximum and minimum values. An example of an evaluation using the $Average(E(U,P))$ -criterion is given in equation (7), where calculation of the $Average(E(U,P))$ -value of the decision alternative shown in Figure 2 is presented. Throughout the paper, $Average$ will be used to denote $Average(E(U,P))$.

$$Average(E(U,P)) = \frac{\int_0^1 \int_0^1 p_1 \cdot u_1 + (1-p_1) \cdot u_2 \, dp_1 du_1 du_2}{\int_0^1 \int_0^1 1 \, dp_1 du_1 du_2} = \frac{0.0044}{0.005} = 0.88 \quad (7)$$

In calculating the $Average$ -value for decision alternatives involving only parameters (probabilities and utilities) that are independent of each other, which was the case in the example presented above, one can utilise the fact that the expected value of the sum of two stochastic variables is equal to the sum of the expected value of each, as well as that the expected value of the product of two *independent* stochastic variables is equal to the product of the expected value of each of the variables. Using this method, the $Average$ -value that was calculated in equation (7) can be calculated instead using equation (8), without any integrals needing to be solved. This simpler form of calculation becomes increasingly useful as the scope of the problem increases and the integrals in equation (5) becomes more cumbersome to solve. Note that when analysing uncertain situations represented by event trees having chance nodes with three or more branches, the probabilities of the different branches are not independent and therefore the problem becomes more difficult to solve. It is possible, however, to evaluate even this type of problem using SSD, a computer algorithm has been developed so as to help the decision maker in doing this [15].

$$\begin{aligned} Average(E(U,P)) &= E(p_1) \cdot E(u_1) + (1 - E(p_1)) \cdot E(u_2) = \\ &= 0.9 \cdot 0.95 + 0.1 \cdot 0.25 = 0.88 \end{aligned} \quad (8)$$

Although the *Average*-criterion is also sensitive to changes near the edges of the decision frame, this criterion, if used in isolation, is nevertheless a natural candidate for use as the general decision criterion. It can be reasonable, however, to employ a set of criteria by combining all three criteria. Thus, if one wanted to evaluate the decision alternative which the situation shown in Figure 2 represents, one would use the three values $Min(E(U,P)) = 0.765$, $Max(E(U,P)) = 0.975$ and $Average(E(U,P)) = 0.88$ and compare these with the comparable values obtained in analysing whatever other decision alternatives involved.

In the original account of SSD [12], the quantitative evaluations which the paper takes up (equations (3) to (5)) are conceived as being employed in combination with qualitative methods. Also, the method described in [12] does not prescribe the preferred alternative needing to be best in terms of *all* of the criteria presented above, although in the present paper we do treat an alternative as being best only if it is the best in terms of all three criteria.

Analysing a small decision example using SSD

In order to exemplify the use of SSD within the context of fire risk management, a simple hypothetical example will be used. The aim of this example is to show how a decision problem can be analysed using SSD.

The hypothetical example concerns the decision of whether a particular investment in fire safety should be made, one rather modest in cost. Assume that the risk manager of a company has found there to be a room containing electrical equipment in which a fire might readily start, a fire that could be very severe in its effects. The risk manager wants to determine whether the decision to invest in a CO₂ system to be installed there would be a good one. Since both the occurrence of fire in that room and the spread of fire from it if a fire should occur are very uncertain, the risk manager decides to use SSD to evaluate the different alternatives.

Denote the expected utility of the alternative of investing in a CO₂ system as $E_1(U,P)$ and the expected utility of the alternative of not investing as $E_2(U,P)$. Although the analysis of the decision here will only involve use of the *Min* and *Max* criteria, $Min(E_1(U,P)-E_2(U,P))$ and $Max(E_1(U,P)-E_2(U,P))$ will be examined, rather than $Min(E_1(U,P))$ being compared with $Min(E_2(U,P))$ and $Max(E_1(U,P))$ with $Max(E_2(U,P))$. This allows account to be taken of the fact that the probability of a

severe fire occurring in the room and the probability that such a fire, if it does occur, will be contained there (within the room of origin) are estimated to be the same for both alternatives. If both $\text{Min}(E_1(U,P) - E_2(U,P))$ and $\text{Max}(E_1(U,P) - E_2(U,P))$ give positive values, it can be concluded that alternative 1 is best, whereas if both evaluations yield negative values, it can be concluded that alternative 2 is best. Note that the approach of evaluating the *difference* in expected utility between two alternatives resembles that of the Delta-method [16].

The first thing to do is to set up the decision frame. Two alternatives have been identified, one involving the company's investing in the CO₂ system (a_1) and the other the company's not investing in it (a_2). Assume, so as to simplify the problem, that a fire in the room can only have three possible consequences: (1) its being too small to have any appreciable impact on the company (consequences $c_{1,1}$, $c_{1,2}$ and $c_{2,1}$ in the event trees shown in Figure 3 and Figure 4), (2) the fire destroys everything in the room of origin but is contained in that room (consequences $c_{1,3}$ and $c_{2,2}$ in the event trees shown in the two figures), and (3) the fire is spreading from the room of origin and destroying the entire factory (consequences $c_{1,4}$ and $c_{2,3}$ in the event trees shown in the same two figures).

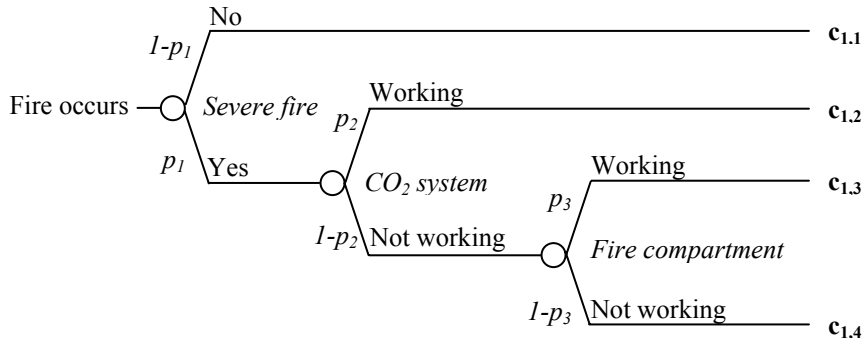


Figure 3 Event tree representing the possible consequences of a fire in the electrical equipment room, given that alternative 1 is chosen.




Figure 4 Event tree representing the possible consequences of a fire in the electrical equipment room, given that alternative 2 is chosen.

In order to continue with the analysis, the decision maker needs to make a statement regarding the probability that a fire with the potential to become severe will occur in the room of interest within the period planned for, which is assumed to be ten years. The risk manager estimates this probability as being lower than 0.2 but not lower than 0.05. For simplicity, assume that during the period planned for only one severe fire can occur. This probability and the remaining statements regarding probabilities are summarised in Table 2.

Table 2 The probability statements and their representation in the form of inequalities.

Statement	Representation
The probability of a potentially severe fire within the next 10 years is between 0.05 and 0.2.	$0.05 \leq p_1 \leq 0.2$
The probability that the CO ₂ system will be working and will extinguish the fire is between 0.7 and 0.95.	$0.7 \leq p_2 \leq 0.95$
The probability that the fire will be contained in the room of origin given that the fire has not been extinguished, is between 0.5 and 0.95.	$0.5 \leq p_3 \leq 0.95$

To arrive at statements regarding the utility values, it can be helpful to begin by visualising the relative positions of the various consequences on a utility scale. Figure 5 shows the utilities of the different consequences ($u_{1,1}, \dots, u_{2,3}$). Note that the distances between the utilities of the different consequences, as shown in the figure, are not correct, only their relative positions being correct. The utility statements and the inequality representations of the utility values are shown in Table 3. Note that the reason for consequence $c_{1,1}$ and $c_{1,2}$ not being as good as

$c_{2,1}$, despite none of them involving any serious fire occurring, is that $c_{1,1}$ and $c_{1,2}$ involve the company's having invested in a CO₂ system, whereas $c_{2,1}$ does not. The same reasoning applies to the difference found between $c_{1,3}$ and $c_{2,2}$. However, $c_{1,4}$ and $c_{2,3}$ can be judged to be equally bad due to the costs of the CO₂ system being small compared with the costs of a total loss of the factory.




Figure 5 Diagram of the utilities of the different consequences. Note that the value distances as shown are not correct, only the relative positions of the consequences on the utility scale being correct.

Table 3 The utility statements and their representation in the form of inequalities and equalities.

Statement	Representation
Consequence $c_{2,1}$ is the best consequence.	$u_{2,1} = 1$
Consequence $c_{1,4}$ is the worst consequence.	$u_{1,4} = 0$
Consequence $c_{1,1}$ and consequence $c_{1,2}$ are equally good.	$u_{1,1} = u_{1,2}$
Consequence $c_{1,4}$ and consequence $c_{2,3}$ are equally bad.	$u_{1,4} = u_{2,3}$
The utility difference between $c_{1,1}$ and $c_{1,4}$ (X_2) is at least 1000 times as great, and not more than 10000 times as great, as the difference between $c_{2,1}$ and $c_{1,1}$ (X_1).	$0.9999 \geq u_{1,1} \geq 0.999$
The utility difference between $c_{2,2}$ and $c_{1,4}$ (Y_2) is at least 100 times as great, and not more than 1000 times as great, as the difference between $c_{2,1}$ and $c_{2,2}$ (Y_1).	$0.999 \geq u_{2,2} \geq 0.99$
The utility difference between consequence $c_{1,3}$ and consequence $c_{2,2}$ is equal to the difference between consequence $c_{1,1}$ and $c_{2,1}$ or less, and the utility of consequence $c_{2,2}$ is equal to or better than that of consequence $c_{1,3}$.	$u_{2,2} \geq u_{1,3} \geq (u_{2,2} - (u_{2,1} - u_{1,1}))$

In calculating the expected utilities of the two alternatives, the event trees shown in Figure 3 and Figure 4 are employed. These result in equations (9) and (10), which represent expressions for the expected utilities of the two alternatives.

$$E_1(U, P) = (1 - p_1) \cdot u_{1,1} + p_1 \cdot p_2 \cdot u_{1,2} + p_1 \cdot (1 - p_2) \cdot p_3 \cdot u_{1,3} + p_1 \cdot (1 - p_2) \cdot (1 - p_3) \cdot u_{1,4} \quad (9)$$

$$E_2(U, P) = (1 - p_1) \cdot u_{2,1} + p_1 \cdot p_3 \cdot u_{2,2} + p_1 \cdot (1 - p_3) \cdot u_{2,3} \quad (10)$$

One can calculate the difference in expected utility between the two alternatives by use of equation (11). Several of the expressions appearing in the column labelled “Representation” in Table 3 have been used in arriving at equation (11).

$$E_1(U, P) - E_2(U, P) = u_{1,1} - p_1 \cdot p_2 \cdot p_3 \cdot u_{1,3} + p_1 \cdot p_2 \cdot u_{1,1} + p_1 \cdot p_3 \cdot (u_{1,3} - u_{2,2} + u_{1,4}) + p_1 \cdot (1 - u_{1,1} - u_{1,4}) - 1 \quad (11)$$

Equation (11) represents the difference in expected utility between alternative 1 and alternative 2. This is the equation one seeks to calculate *Min* and *Max* for. Equation (11) can be regarded as the objective function in a nonlinear multivariable optimisation problem. The constraints of the problem are given by the decision frame. The constraints are presented in the “Representation” columns in Table 2 and Table 3.

Although the present optimisation problem is relatively simple to solve, and can readily be solved using only hand calculations, a function called “fmincon” in the Optimization Toolbox for MATLAB [17] was used to solve it. The results obtained were that $Min(E_1(U, P) - E_2(U, P))$ is $7.84 \cdot 10^{-4}$ and $Max(E_1(U, P) - E_2(U, P))$ is $9.59 \cdot 10^{-2}$. Since both the *Min* and the *Max*-evaluations result in positive values, one can conclude that the best alternative is to invest in a CO₂ system (alternative 1) without performing an evaluation of the *Average*-value (see case 1 in Figure 9). Note that, despite one’s not knowing the exact value of the expected utility of either alternative 1 or alternative 2, it can be shown that the expected utility of alternative 1 is higher than that of alternative 2. Thus, the decision to invest in a CO₂ system is robust. This implies that the decision alternative is the best, regardless of which values of the

uncertain variables are “correct” (assuming the values to be contained within the decision frame). This is an important principle, one that will be discussed in detail later in the paper.

Comparing SSD with Bayesian decision analysis: An example of a real-world example

The aims of the example that follows are to show how SSD can be applied to a real-world decision problem and to compare the results of the SSD analysis with those of two other decision analysis methods. The two methods used for purposes of comparison are those referred to in a previous section as traditional decision analysis and extended decision analysis.

The decision problem here concerned the question of whether an investment in a water sprinkler system for a factory should be made. Since the production in the factory, which belonged to the firm ABB, was very important for the company, a serious fire in the building would have had extremely negative consequences. The decision problem was analysed earlier by use of extended decision analysis, an analysis described in greater detail in [3] and [4]. The alternatives the decision maker (the company) could choose between were to invest in a water sprinkler system (a_1) and to not invest in it (a_2). Evaluation of the alternatives was performed by use of an event tree technique aimed at modelling different fire scenarios that were possible. Since the building to which the analysis referred was large (55,000 m²), an extensive decision analysis was required. In order to simplify the presentation of the problem and the comparison of the results, it was decided here to only carry out a comparative analysis for one of the fire compartments in the building (that was approximately 5,500 m² in size). This fire compartment was treated as if it were a separate building, one for which the decision maker was to decide whether to invest in a water sprinkler system.

In the original analysis, the consequences were expressed in terms of monetary losses and, since the monetary losses associated with any given fire scenario were uncertain, they were expressed by use of triangular probability distributions. The same approach was employed for the probabilities used in the model. The minimum, the most likely and the maximum values for the probability distributions are presented in Table 4 and in Table 5. These probability distributions were arrived at in discussions between the analyst and personnel both from ABB and from

the fire department. The monetary sums given in this section, originally in Swedish crowns (SEK), were converted to US dollars at the rate of \$1 to 10 SEK.

The event tree used to describe each of the fire scenarios considered is presented in Figure 6 (note that in the alternative of there being no sprinkler system in the building the probability that the water sprinkler system will succeed in extinguishing the fire is considered to be 0). Using this event tree in combination with the different estimates of the probabilities and the consequences allows one to calculate the expected utility of investing in a water sprinkler system. Assume exactly 1 (one) fire to be the number of fires expected to occur during the time for which the decision maker wishes to take account of the benefits the sprinkler system would provide. Certain assumptions are also made concerning the decision maker's preferences regarding the occurrence of more than one fire during a given period of time. In particular, it is assumed that the expected utility of k fires occurring during a given period, each fire having the expected utility of $E(u)$, is $kE(u)$. This implies that the utility of any given fire is not affected by how many other fires occurred during the period in question. The assumptions just mentioned allow the expected utility for each of the two decision alternatives to be calculated by multiplying the expected utility of one fire by the number of fires expected to occur during that time period. In the present case, 1 fire is expected to occur.

Note that one could discount losses occurring in the future in the same way as is done in capital investment analysis. Since the focus here, however, is on the decision rules used in the different methods and on how to deal with epistemic uncertainty, no attempt is made to discount future losses. For further information on this matter, see [18].

For simplicity, assume that the decision maker's preferences with respect to uncertain monetary outcomes can be described by a risk-neutral utility function, which means that he/she evaluates uncertain monetary outcomes exactly according to their monetary values. Since in the present context only potential losses are being analysed, it might be regarded as being more appropriate to use the term disutility in discussing the decision maker's preferences with respect to different losses. Nevertheless, the term utility will be used here, the worst consequence being assigned a utility value of -1 and the best consequence a utility value of 0. In assigning utility values to the

different consequences, we use the monetary outcomes reported in Table 5, assigning a utility value of -1 to the consequence involving a loss of \$38,820,000 (loss of the entire factory, including the sprinkler system) and a utility value of 0 to a loss of \$0. Since we assume that the decision maker is risk neutral, each of the other monetary outcomes can easily be translated into utility values of between 0 and -1.

In addition, we assume that the sprinkler system costs \$100,000, which is a reasonable assumption in view of the fact that in the original analysis the total cost of the sprinkler system for the building as a whole was approximately \$1,000,000. This \$100,000 cost is taken account of when one calculates the expected utility of the alternative of investing in a sprinkler system.

Table 4 *The minimum, the most likely and the maximum value of the different probabilities used in the model. All the probabilities are conditional upon the event that a fire has occurred in the building.*

Probability	Abbreviation	Min	Most likely	Max
The probability that the growth potential of the fire will be low.	P(Low)	0.55	0.80	0.85
The probability that the fire detection system will detect the fire.	P(Alarm)	0.90	0.98	0.99
The probability that the employees will succeed in extinguishing the fire given that the fire detection system has detected it.	P(Emp. Alarm)	0.70	0.80	0.95
The probability that the employees will succeed in extinguishing the fire given that the fire detection system has <i>not</i> detected it.	P(Emp. NoAlarm)	0.20	0.50	0.60
The probability that the fire department will succeed in extinguishing the fire before it destroys the fire compartment, given that the fire detection system has detected the fire.	P(Fire dept. Alarm)	0.50	0.70	0.85
The probability that the fire department will succeed in extinguishing the fire before it destroys the fire compartment, given that the fire detection system has <i>not</i> detected the fire.	P(Fire dept. NoAlarm)	0.20	0.30	0.50
The probability that the water sprinkler system will extinguish the fire.	P(Sprinkler)	0.90	0.95	0.98

Table 5 The minimum, the most likely and the maximum value of the different monetary consequences used in the model.

Fire scenario	Abbreviation	Monetary loss associated with a particular fire scenario (\$ thousand)		
		Minimum	Most likely	Maximum
The fire will be extinguished by the employees or by the sprinkler system.	C ₁	5	10	15
A fire with low growth potential will be extinguished by the fire department.	C ₂	25	50	75
A fire with high growth potential will be extinguished by the fire department.	C ₃	533	1067	1933
The fire will completely destroy the fire compartment.	C ₄	25920	32000	38720

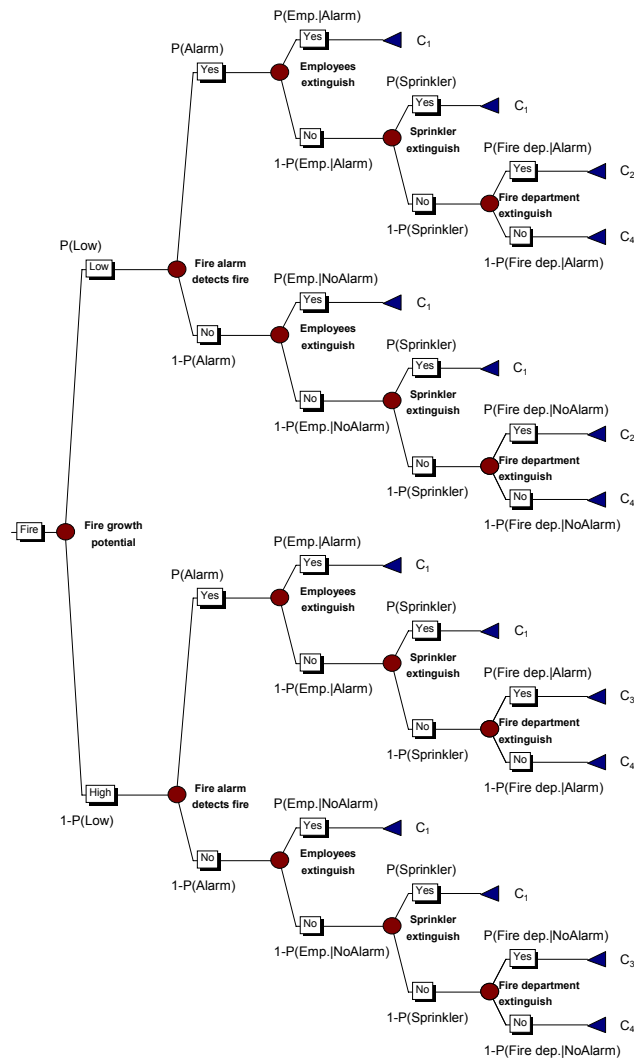


Figure 6 Event tree showing the different fire scenarios considered in the analysis of whether to invest in a sprinkler system.

One aim of this example is to clarify the relationship between decision analyses performed using Bayesian decision theory, as described in [3] and [4], and one performed using Supersoft Decision theory. In one of the analyses using Bayesian decision theory (the extended decision analysis method) two Monte Carlo-simulations were carried out so as to be able to relate the uncertainty contained in the probability and

consequence estimates to the expected utility of the different alternatives. In each of the two Monte Carlo-simulations, 5000 iterations were performed. The distribution representing the expected utility of a particular alternative can then be compared with the results generated by analysing the alternative by use of SSD. The results of the Monte Carlo-simulations, expressed in terms of expected utility, and of the analysis using SSD are presented in Figure 7 and Figure 8. The vertical lines are the results of the SSD evaluation using the *Max* and *Min* values given in Table 4 and Table 5. The line associated with the highest utility is the result of the *Max*-evaluation (see equation (4)), the line associated with the lowest expected utility being the result of the *Min*-evaluation (see equation (3)). The line in the middle is the result of the *Average*-evaluation. The *Average*-evaluation was performed using the expected values of the probabilities and the utilities, assuming there to be a uniform distribution between the highest and lowest values, and utilising the same technique as illustrated by equation (8).




Figure 7 Results of the analysis of the alternative of keeping the building in its present condition. Note that the horizontal scale is not the same as the one used in Figure 8.




Figure 8 Results of the analysis of the alternative of investing in a water sprinkler system for the fire compartment in question. Note that the horizontal scale is not the same as the one used in Figure 7.

The results of the SSD-evaluation are presented in Table 6. A decision analysis using exact values for the probabilities and utilities (traditional decision analysis) was also performed. The values for the probabilities and the consequences as given in the “Most likely” column in Table 4 and 5 were used to calculate the expected utilities of the two decision alternatives. The results are presented in Table 7.

Comparing the results of evaluating the two decision alternatives allows one to conclude that, in terms of an extended decision analysis, the alternative of investing in a sprinkler system is best, since the mean value of the distribution shown in Figure 8, which represents the alternative of investing in the sprinkler system, is higher than the mean value of the distribution shown in Figure 7, which represents the alternative of not investing in the sprinkler system. The results SSD provides allow the same conclusion to be drawn, the alternative of investing in the sprinkler system being best, since the *Min*-evaluation gives a higher expected value for the sprinkler alternative, as well as for the *Max*-evaluation and the *Average*-evaluation. The results obtained using traditional decision analysis (see Table 7) imply the same

conclusion to be reached there as in the other two analyses, namely that the decision alternative of investing in a sprinkler system is best.

Table 6 Results of the SSD-evaluation.

Alternative	Expected utility		
	Min	Average	Max
Sprinkler	-0.0232	-0.0067	-0.0028
No sprinkler	-0.2023	-0.0641	-0.0065

Table 7 Results of the evaluation using exact probability and utility values.

Alternative	Expected utility
Sprinkler	-0.00559
No sprinkler	-0.05536

Decision robustness

There are differences between the three approaches to evaluation in terms of decision robustness. Here, decision robustness is used to denote how likely it is that the best alternative would change if a reasonable degree of change were to be made in the estimation of either the probabilities or the utilities. Since the probabilities and utilities are expressed as exact values when traditional decision analysis is employed, that method provides no information concerning robustness.

One way of evaluating the robustness of a decision is to compare the resulting distributions of the expected values for all of the alternatives and to note whether these distributions overlap to an appreciable extent. Looking at the results of the extended decision analysis of the ABB-case, it is clear that the decision to invest in the sprinkler system can be considered robust, since the expected utility distributions for the two alternatives under consideration do not overlap. Although an SSD evaluation does not result in a probability distribution, one can investigate whether the intervals defined by *Min* and *Max* overlap. In the ABB-case, one finds that the intervals for the two alternatives do overlap. The fact that the two approaches to evaluation differ in robustness (i.e. in the extent to which the intervals or distributions overlap) is due to the SSD approach being conservative in terms of its manner of dealing with robustness. More specifically, since the endpoints of the interval representing the possible values for the

expected utility of an alternative obtained by use of SSD usually has little credibility in terms of extended decision analysis, they are not part of the distributional results presented in Figure 7 and in Figure 8. Comparing the results of the two approaches highlights the fact that SSD is more coarse in its treatment of robustness than the extended decision analysis method is.

The robustness of a decision is very important when a high degree of epistemic uncertainty is involved, since taking account of robustness can enable the decision maker to reach a definite conclusion there nevertheless regarding which alternative is best. When the robustness of a decision problem has been analysed, such a conclusion can be drawn, provided one of the decision alternatives has the highest expected utility and its expected utility is clearly separated from those of the other alternatives (i.e. if there is no overlap between the distributions or intervals). Note that to conclude by use of this method that a particular alternative is the best requires (1) that one accepts expected utility as being the basis for evaluation and (2) that the decision frame (i.e. the basis for analysis) contains all the plausible values of both the probabilities and the utilities. How the decision situation is structured when the expected utilities are separated from each other can be illustrated by case 1 in Figure 9, in which the results are presented in the manner typical for SSD.

If the two alternatives are not completely separated in terms of the plausible values of the expected utilities assigned to them, one needs an evaluation criterion appropriate for determining the best alternative under such conditions. Two such criteria have been presented in the paper, the maximisation of expected utility (MEU) and the SSD criteria (equation (2) to (5)). An extended decision analysis uses the MEU criterion but also utilises the distributions of the expected utilities in determining the degree of robustness, i.e. how much the distributions of the expected utilities overlap. The situation in which the two distributions of expected utilities are not entirely separated and an appropriate decision rule for determining the best alternative under such conditions is needed is illustrated by case 2 in Figure 9. There, in terms of SSD, alternative 1 is the best since each of the three criteria (equations (3) to (5)) gives a higher value ($E(U_1)$) for alternative 1 than for alternative 2 ($E(U_2)$).




Figure 9 Illustration of three different results of an evaluation of two decision alternatives using SSD.

One can end up, however, with a situation in which the evaluation criteria taken up in the present paper cannot provide any recommendation of which alternative to select. This can happen when the possible values of the expected utilities for the different alternatives are too close together and none of the intervals of possible expected utility values are clearly higher in terms of expected utility values. This is illustrated by case 3 in Figure 9, in which the results of an SSD evaluation of two decision alternatives are presented. Since the *Average* and *Max* values for alternative 1 are higher than the corresponding values for alternative 2, one might think that alternative 1 is better. However, since the *Min* value of alternative 1 is lower than that of alternative 2, one cannot conclude that alternative 1 is best. Instead, one needs to obtain more information regarding the problem, so as to

hopefully reduce the uncertainties sufficiently to arrive at a clear conclusion. Note that traditional decision analysis utilising exact values of the probabilities and the utilities could not have distinguished between the three cases shown in Figure 9. That method would have indicated alternative 1 being the best in all three cases.

Discussion on the application of decision analysis to a possible fire protection investment in a specific building

In the previous sections, three methods for decision analysis were compared in terms of their applicability to a problem involving a high degree of epistemic uncertainty. One can conclude that these methods differ substantially in how precise one needs to be when estimating the probabilities and evaluating the consequences. It is doubtful whether traditional decision analysis, which requires that the parameters involved be assigned exact values, has any practical usefulness in the present context, since in dealing with possible fire protection investments in a particular building one is not likely to have the amount of information needed to assign exact values. Using such a method in the present context could in fact be very misleading, since the results could give the impression that those decision alternatives which are *not* identified as being the best are *clearly* inferior to the alternative deemed best whereas in reality the obtaining of additional information might well lead to the results one arrives at changing easily. For a context such as the present one, therefore, a method which involves expressing the probabilities and utilities as being uncertain is more appropriate.

From a practical standpoint, methods such as SSD which involve interval statements seem attractive. In a practical decision situation there may be several stakeholders and thus several “decision makers”. Under such conditions, it may be impossible for the stakeholders to agree upon a specific distribution for the probability and utility values to be employed. Instead, each stakeholder could assign the parameters of interest a maximum and a minimum value. One could then employ the lowest of the stakeholders’ minimum values together with the highest of the stakeholders’ maximum values in the analysis to be carried out. This would result in a decision frame that includes all the stakeholders’ estimates. If the analysis resulted in a robust decision, all the stakeholders should be satisfied with the decision alternative that was recommended (provided they accepted expected utility as a reasonable means of evaluating the decision alternatives).

Another important aspect of decision analysis in the present context is that in many applications one is only interested in determining whether a particular alternative is better or worse than the other alternatives, not in *exactly* how much better or worse it is. This can be exemplified by the two examples included in the present paper in which the question was which of two alternatives one should choose. In the first example, it was shown that to answer this question one did not need to know exactly how much better or worse the alternatives were in comparison to each other. Instead it was sufficient to show, within the decision frame at hand, that for all plausible values of the evaluation criteria (in this case, expected utility) one alternative was better than the other. This was almost true, but not quite, in the second example, where an extended decision analysis showed there to be no overlap between the distributions of expected utility, but evaluation by SSD showed there to be some possible values of the probabilities and utilities for which the decision alternative to be recommended changed.

This brings up one point concerning differences between SSD and extended decision analysis. As was mentioned earlier, SSD is conservative with respect to robustness. This means that, even if SSD indicates there to be an overlap of the intervals of the expected utilities, the decision situation may very well be considered robust in terms of extended decision analysis since the endpoints indicated by an SSD analysis have so little credibility, the values involved being so unlikely that in practice the decision can be considered robust. It is clear, therefore, that in decision situations in which the information available regarding the probabilities, for example, is sufficient to justify expressing epistemic uncertainty by use of specific probability distributions, extended decision analysis is better to use than SSD. However, the salient argument in favour of SSD is that, since it requires no probability distributions for representing epistemic uncertainty, it can be used for decision analysis in cases in which the information regarding the probabilities, for example, is too vague to justify using specific probability distributions. In such cases, extended decision analysis cannot be employed. Thus, SSD and extended decision analysis should not be viewed as competing, but rather as complementary methods, the one being useful when precision in terms of robustness is sought and when the information at hand justifies epistemic uncertainties being expressed as specific probability distributions, the other being useful when the information at hand is vague and does not justify expressing uncertainties as specific probability distributions. Also, note that SSD

and extended decision analysis will provide results that are very similar when use is made of the Average evaluation in connection with SSD (equation (5)) and the expected utility evaluation in connection with extended decision analysis (equation (2)), the results of the two being identical, in fact, if uniform distributions are used to represent the epistemic uncertainty involved in the case of extended decision analysis.

It could appear that SSD involves more complicated and time-consuming calculations in evaluating decision alternatives than extended decision analysis does. Use of computers, however, can make the calculations SSD would require no more cumbersome than those involved in Monte Carlo-simulations, SSD calculations probably being faster, in fact, than Monte Carlo-simulations because of one's not having to define probability distributions in SSD.

In practice, the choice of a method for analysing different fire protection alternatives for a specific building is not always an easy one, its depending very much on the situation at hand. One can conclude, however, that traditional decision analysis, if used in isolation, is clearly unsuitable for decision problems involving a high degree of epistemic uncertainty, due to its inability to provide or utilise information regarding the uncertainty of the results. This can be seen as applicable to Bayesian analysis as well. There, although probabilities and utilities are expressed as probability distributions, only one value of the expected utility is used when alternatives are compared. In many situations, extended decision analysis is probably useful, since it provides information regarding the robustness of a decision in a way that is readily grasped (yielding a distribution of expected utilities). The decision maker who finds it difficult to assign probability distributions to uncertain parameters or lacks the time to do so can use the SSD method instead.

In practice, a possible procedure in analysing a decision problem would be to start using a rough model involving use of SSD. If the results indicate that the decision problem to be of type 1 or 2 as shown in Figure 9, one can readily conclude which alternative is best. One could then take the analysis one step further if one wished, using extended decision analysis (if the information justified its use) to analyse the robustness of the decision as adequately as possible. If, on the other hand, the results of the initial SSD evaluation indicate the decision problem to be of type 3, one would need to collect more information about the problem at hand

so as to either be able to conclude which alternative was best through use of SSD or use extended decision analysis for determining the robustness of the decision problem.

Note that the analysis just presented of the differences between the three different methods discussed also applies to situations in which determining which alternative is best is based not on expected utilities but on other measures of evaluation. If decision alternatives are to be screened, for example, by excluding from further analysis any alternatives for which the probability of an extreme event occurring is too high, one can make use of exact probabilities, probability distributions or SSD, the latter two approaches allowing one to analyze the robustness of the screening process.

Conclusions

Decision analysis of fire safety decisions applying to the choice of a possible fire protection investment was discussed, an investment applying not to buildings in general or some basic category of buildings but to a specific building, particular attention being directed at situations in which one can expect there to be a considerable degree of epistemic uncertainty. A new decision analysis method termed Supersoft Decision Theory (SSD) was introduced and was applied to problems of fire safety. Two concrete applications of SSD in a fire safety context were described. The first case discussed was a hypothetical decision situation of limited scope aimed at illustrating some of the calculations to be made when evaluating alternatives by use of SSD. The second case involved part of a real-world decision problem that had been analysed earlier by use of Bayesian decision theory. The SSD analysis of the second case was compared with two other types of decision analysis utilising Bayesian decision theory, one of them termed traditional decision analysis and the other termed extended decision analysis. The traditional decision analysis utilised exact values of the probabilities and utilities involved, whereas the extended decision analysis used probability distributions to represent the uncertainty regarding the probabilities and utilities.

It was concluded that a decision analysis involving use of precise values for probabilities and utilities can easily be misleading, since the results obtained would provide no indication of the robustness of the decision, i.e. of how readily the alternative judged to be best could change if a

reasonable degree of change in the probability or utility values should occur.

It was also argued that the robustness of a decision is an important consideration when the decision situation involves a high degree of epistemic uncertainty. Methods in which no evaluation of the robustness of the decision is provided are not suitable for analysing decision situations in which there is a high degree of epistemic uncertainty.

Supersoft decision analysis and extended decision analysis are not viewed as competing methods. Rather, they are seen as complementing each other, the one being able to deal better than the other with situations of certain types. Use of extended decision analysis is appropriate when epistemic uncertainty can be quantified in terms of specific probability distributions, whereas Supersoft decision analysis is appropriate when no such quantification is possible. Extended decision analysis provides more precise information on the robustness of the results (since it provides a distribution for each expected utility), whereas Supersoft decision analysis provides only an interval for the expected utilities. In practice, one could start analysing a problem by use of a rough model based on SSD. Depending on the results of this initial analysis, one could then, if desired and regarded as possible, continue with a more refined analysis involving use of extended decision analysis.

The examples presented indicate Supersoft Decision Theory (SSD) to be a tool that can be used for analysing practical decision problems. The main advantage of using SSD is that it allows the decision maker to employ imprecise statements concerning the probability and utility values used in the model. This also makes SSD particularly suitable for difficult decision problems in which the decision maker is not a single person, but a group of persons. The use of imprecise statements allows the interval used to denote the probabilities and utilities to be sufficiently broad to encompass the estimates and views of all the group members.

Acknowledgements

The authors would like to thank The Swedish Fire Research Board (BRANDFORSK) for funding the research on which the present paper is based.

References

1. Oberkampf, W. L., DeLand, S. M., Rutherford, B. M., Diegert, K. V., Alvin, K. F., "Error and uncertainty in modeling and simulation", *Reliability engineering and system safety*, Vol. 75, No. 3, 2002, pp. 333-357.
2. Gomborg, A., Buchbinder, B., Offensend, F. J., *Evaluating alternative strategies for reducing residential fire loss – The fire loss model*, U.S. Department of Commerce, Washington, National Bureau of Standards, NBSIR 82-2551, 1982.
3. Johansson, H., *Decision Making in Fire Risk Management*, Report 1022, Lund, Dept. of Fire Safety Engineering, Lund University, 2001.
4. Johansson, H., "Decision analysis concerned with investments in fire safety", *7th International Symposium on Fire Safety Science*, not yet published.
5. Donegan, H. A., "Decision Analysis", *SFPE Handbook of Fire Protection Engineering*, 3rd ed., Quincy, National Fire Protection Association, 2002.
6. Laplace, P. S., *Essai Philosophique sur les Probabilites*, 5th ed., Paris, Translation by Dover, New York, 1952.
7. Wald, A., *Statistical Decision Functions*, New York, Wiley, 1950.
8. Savage, L. J., "The Theory of Statistical Decision", *J. Amer. Statist. Assoc.*, Vol. 46, 1951, pp. 55-67.
9. Hurwicz, L., "Optimality Criteria for Decision Making Under Ignorance", Discussion paper No. 370, *Cowles Commission*, 1951.
10. Berger, J. O., *Statistical Decision Theory – Foundations, Concepts and Methods*, New York, Springer-Verlag, 1980.
11. Malmnäs, P.-E., *Foundations of Applicable Decision Theory*, Stockholm, Department of Philosophy, Stockholm University, 1999.
12. Malmnäs, P.-E., *Methods of Evaluation in Supersoft Decision Theory*, Stockholm, Department of Philosophy, Stockholm University, 1995.
13. Wallsten, T. S., Budescu, D. V., Rapoport, A., Zwick, R., Forsyth, B., "Measuring the vague meanings of probability terms", *Journal of Experimental Psychology: General*, Vol. 115, 1986, pp. 348-365.
14. Spetzler, C. S., Stael Von Holstein, C.-A. S., "Probability Encoding in Decision Analysis", *Management Science*, Vol. 22, No. 3, 1975, pp. 340-358.
15. Malmnäs, P.-E. (ed.), *Hypermjuk beslutsteori och ekonomisk optimering av det industriella brandskyddet*, Stockholm, Thales, 2002.

16. Danielsson, M., *Computational Decision Analysis*, Doctoral Thesis, Stockholm, Department of Computer and Systems Sciences, Stockholm University and Royal Institute of Technology, 1997.
17. Coleman, T., Branch, M. A. and Grace, A., *Optimization Toolbox User's Guide*, Version 2, The Math Works Inc., 1999.
18. Johansson, H., "Investment appraisal using quantitative risk analysis", *Journal of Hazardous Materials*, Vol. 93, 2002, pp. 77-91.
19. Hagen, O., "Towards a Positive Theory of Preferences Under Risk", *Expected Utility and the Allais Paradox*, Allais, M. and Hagen, O. (eds.), Dordrecht: Reidel, 1979, pp. 271-302.
20. Fishburn, P., "Transitive Measurable Utility", *Journal of Economic Theory*, Vol. 31, 1983, pp. 293-317.
21. Loomes, G., Sugden, R., "Regret Theory: an Alternative Theory of Rational Choice Under Uncertainty", *The Economic Journal*, Vol. 92, 1982, pp. 805-824.
22. Green, J., Jullien, B., "Ordinal Independence in Nonlinear Utility Theory", *Journal of Risk and Uncertainty*, Vol. 1, 1988, pp. 355-387.
23. Quiggin, J., "A Theory of Anticipated Utility", *Journal of Economic Behaviour and Organization*, Vol. 3, 1982, pp. 323-343.
24. Yaari, M., "The Dual Theory of Choice Under Risk", *Econometrica*, Vol. 55, 1987, pp. 95-115.

Paper 4: A Bayesian network model for the continual updating of fire risk measurement

Johansson, H., *Submitted to Fire Science and Technology.*

A Bayesian network model for the continual updating of fire risk measurement

Henrik Johansson

Department of Fire Safety Engineering
Lund University
P.O. Box 118
S-221 00 Lund, Sweden
Tel. +46462224850
Fax. +46462224612
Email: henrik.johansson@brand.lth.se

Abstract

A risk measure based on decision theory is suggested for use in factories in which continual updating of fire risk assessments is aimed at. It is argued that, when a specific building is involved, this measure has certain advantages compared with the expected value measure more commonly employed. It is shown how use of this measure can be combined with use of a Bayesian network for measuring the fire risk present over a period of time. Although a technique termed fractional updating can be employed for updating the Bayesian network as new information relevant to the fire risk in the building is received, the amount of information of this sort that any fires that have occurred in the building will provide will be only slight if few fires occur. A model for updating the Bayesian network under such conditions through use of expert estimates is provided. An example of a real-world analysis is presented to exemplify use of this method.

Keywords: *Bayesian networks, fire risk analysis, decision analysis, continual updating.*

1. Introduction

Consider a firm interested in managing its fire risks. In order to manage risk one needs to be able to measure it. It would be desirable for a risk manager to have a measure of fire risk that was continually updated so that an increase or a decrease in fire risk could be registered. Such changes in fire risk are likely to occur since the conditions in a building relevant to the development of a fire are likely to change over time. For example, the walls may be moved or penetrated by cables, employees

may receive training in manual fire fighting, such active fire fighting systems as the sprinkler system may not receive proper maintenance, and the like. Intuitively, such changes in the conditions relevant to the development of fire are likely to affect the fire risk in a building, but the question is how much?

The aim of the present paper is to address the issue of the continually revised measurement of fire risk in a specific building by suggesting an integration of a decision analysis method developed earlier (see [1], for example) with the use of Bayesian networks. The idea of using Bayesian networks is that they can easily be updated when new information arrives, and can thus easily be used for the continually revised measurement of risk. The decision analytical framework is used to provide an operationally meaningful interpretation of the measure of risk. The paper also describes how expert judgement can be used in the continual updating of the risk measure.

Note that the paper is concerned primarily with industrial buildings in which the consequences of a fire can be evaluated in monetary terms. To extend the use of the method suggested here so that the consequences of a fire can be evaluated in terms of lost lives, for example, is theoretically straightforward but requires difficult value judgements regarding human lives, which is outside the scope of the paper.

2. Measures of fire risk

Hall and Sekizawa [2] suggest a risk-measure called “outcome measure of fire risk” which is the expected value of a severity measure applied to fire. A risk-measure following the suggestion Hall and Sekizawa have made and of possible use in the present context, in which industrial buildings are of concern, is that of expected loss (EL). The loss could be expressed in this case in monetary terms, such as in \$ per year. The definition of the EL is presented in equation (1). λ is the expected number of fires annually, t is the number of years considered, $C(s_i)$ is the loss, expressed in monetary terms, associated with fire scenario s_i , $P(s_i)$ is the probability of fire scenario s_i given that a fire has occurred, and n is the number of fire scenarios considered.

$$EL = \lambda \cdot t \cdot \sum_{i=1}^n P(s_i) \cdot C(s_i) \quad (1)$$

A risk measure based on an expected *value* measure, such as that of expected loss, does not consider the risk-attitude of the firm that is exposed to the risk. In the present context, using a risk measure that takes account of this is important since firms may differ in their assessment of a given risk due to differences in their risk-attitude. Whereas for a large corporation the loss of \$100,000, for example, may not be very serious, a loss of this size can be catastrophic for a small firm. In addition, since the concern here is with a *specific* building, use of expected value as a risk measure is not very appropriate. If one had been dealing with a large group of buildings in which many fires could be expected, the expected value of these fires would have been a reasonable risk measure, since through the number of fires being large, random effects could be expected to level out, meaning that the actual outcome should be close to the expected one. In contrast, in a single building in which few fires, if any, are expected to occur during the building's lifetime, use of expected loss as a risk measure cannot be supported by the arguments presented above. Instead, one needs to find some other measure of risk or support expected loss as a risk measure using other arguments.

The approach taken here is to use decision theory as a basis for a measure of fire risk. In decision theory the preferences and the risk attitude of the decision maker (in this case of the firm) are modelled explicitly, the result being a recommendation of which of a set of different decision alternatives, which may be risky prospects, the decision maker should choose. A risky prospect is an uncertain situation in which the decision maker can suffer negative consequences, without any possibility of positive consequences occurring instead.

Classical decision theory, also called Bayesian decision theory (see [1] for a discussion of the application of Bayesian decision theory to fire safety problems), is based on the decision maker's accepting a set of axioms for his/her decision making. Various versions of the axioms have been suggested. Luce and Raiffa [3], for example, provided the following six axioms:

Ordering. A set of outcomes can be ordered using a "preference or indifference" ordering.

Reduction of compound lotteries. A decision maker is indifferent between a complicated compound lottery and an equivalent lottery

involving only a simple uncertain event, the equivalence of which is determined on the basis of standard probability manipulations.

Continuity. A decision maker is indifferent between the outcome of a particular lottery and the outcome of an equivalent lottery involving only the best and the worst outcome of the first lottery.

Substitutibility. A decision maker is indifferent between any lottery involving outcome o_i and a lottery in which o_i is replaced by a lottery that is judged to be equivalent to o_i .

Transitivity. Preference and indifference between uncertain situations are transitive relations. This means that if a decision maker prefers alternative L_1 to alternative L_2 , and alternative L_2 to alternative L_3 , then he/she must prefer alternative L_1 to alternative L_3 .

Monotonicity. A lottery involving only two possible outcomes is preferred to a similar lottery involving the same two outcomes but where the probability of the better outcome is less than in the preferred lottery.

It can be shown (see [3], for example) that if a decision maker makes decisions in accordance with the principle of maximising expected utility, he/she acts in accordance with the axioms just referred to. Thus, if a decision maker accepts the axioms as rational and wants to adhere to them, he/she can choose to make all his/her decisions according to the principle of maximising expected utility and can be certain of not violating the axioms.

If the comparison of two sources of risk (two buildings, for example) in terms of the risk involved is considered as being a choice between two risky prospects (its being necessary to choose between them), the expected utility criterion can be used to determine which risk source would be best to choose. That risk source could then be interpreted as involving the lesser risk of the two. Viewing risk in terms of this decision analytical approach provides an operationally meaningful interpretation of what a higher risk of one alternative than another means:

For a decision maker who is determined to make decisions in accordance with the principle of maximising expected utility, a risky prospect for which the risk is higher than it is for another

prospect is seen as less preferable than the prospect for which the risk is lower.

As this implies, the decision analytical meaning of a reduction in risk in a particular building from one year to another means that, if the decision maker owning the building follows the principle of maximising expected utility, he/she would prefer (if a choice of this sort were possible) exposing himself/herself to the fire risk the building has today than the fire risk it had the year before.

Adopting this view of risk makes it possible to speak meaningfully of differences in risk and of what such differences means to a decision maker in concrete terms. It is possible to estimate, for example, what is termed the certainty equivalent of an uncertain situation. This is the monetary value the decision maker regards as equivalent to the value of the uncertain situation, the decision maker thus being indifferent between the alternative of losing this certainty equivalent and that of suffering the consequences of the uncertain situation.

In a given choice situation which is uncertain and in which the probabilities involved are “objective” the certainty equivalent can differ with the decision maker, depending on the risk-attitude the latter has. The risk-attitude of the decision maker can be of any of three different types, these being a risk-averse, a risk-neutral risk-attitude and a risk-seeking attitude towards risk.

A risk-neutral decision maker is one who evaluates a risky situation on the basis of its expected value. This implies the decision maker’s being indifferent between the risky prospect and its expected (negative) outcome. In the present context, this is the same as evaluating the “risk” involved in terms of expected loss (see equation (1)).

In contrast, a risk-seeking decision maker evaluates a risky situation as being worth more than the expected value of the outcome. Thus, given the choice between paying the expected value of the risky situation and exposing himself/herself to its consequences, the decision maker would choose to expose himself/herself to the consequences of the risky situation. Just the opposite is true of a risk-averse decision maker, who would rather prefer to pay the expected value of a risky prospect rather than exposing himself/herself to its consequences.

The certainty equivalent (CE) of a risky situation can be calculated by use of equation (2), where $U(s_i)$ is the utility associated with consequence s_i (which is expressed in monetary terms), and $P(s_i)$ is the probability of the consequence in question occurring. The utility of the best and of the worst consequence of a risky prospect can be assigned utility values 1 and 0 respectively.

$$U(CE) = \sum_{i=1}^n P(s_i) \cdot U(s_i) \quad (2)$$

The utility of different consequences can be determined by letting the decision maker state his/her preferences for various lotteries involving the consequences in question (see [4], for example). If the consequences are expressed in terms of their monetary value, one can create a utility function for money. The utility function indicates the decision maker's risk-attitude, i.e. whether he/she is risk-neutral, risk-seeking or risk-averse, and in the case of the latter two to what extent. As noted above, different risk-attitudes lead to different evaluations of risky situations and thus to differing evaluations of the risk associated with the alternatives in question.

Consider a risky situation (such as one involving the possible consequences of a fire) in which three consequences are possible. The first consequence (s_1) involves nothing happening, the monetary evaluation of this consequence being (\$0). The second consequence (s_2) is evaluated as being equivalent to a loss of \$1,000,000. The third consequence (s_3), finally, is equivalent to a loss of \$30,000,000. The first consequence is assigned a utility value of 1 and the third a utility value of 0. For values of between -\$30,000,000 and \$0, the form of the utility function will depend on the risk-attitude of the decision maker. Figure 1 shows three different utility functions for that segment of scale. According to a decision maker whose risk-attitude is represented by the risk-seeking utility curve in the figure the utility of the second consequence (s_2) is 0.983. According to a decision maker whose risk-attitude is represented by the risk-averse utility function the same consequence has a utility value of 0.934, and according to a decision maker who is risk-neutral the utility of this consequence is 0.967.




Figure 1 Three different types of utility functions.

Assume that the probability of the first consequence (s_1) is 0.9, that of the second consequence (s_2) is 0.09 and that of the third consequence (s_3) is 0.01. This implies that the expected utility of the risky situation is 0.988 for the risk-averse decision maker, 0.987 for the risk-neutral decision maker and 0.984 for the risk-seeking decision maker. The expected utilities can be translated into monetary values in the form of certainty equivalents, in accordance with equation (2). Those represent the monetary amount the decision maker in question regards as being equivalent to the value of the risky situation. The resulting certainty equivalents for the three different decision makers are presented in Table 1. There, one can see the effect of different types of risk attitudes on the evaluation of the risky situation. Since the expected value of the risky situation is -\$390,000, the decision maker who is risk-neutral evaluates the situation as being equivalent to this amount, the one who is risk-averse evaluates it as being equivalent to -\$727,000, his/her thus being willing to pay more to avoid the situation in question than the risk-neutral decision maker would, the risk-seeking decision maker, finally evaluating the situation as being equivalent to -\$239,000.

Table 1 *Certainty equivalents for three different decision makers under conditions described in the text.*

Decision maker	Certainty equivalent (\$)
Risk-averse	-727,000
Risk-neutral	-390,000
Risk-seeking	-239,000

Using a decision analytic approach, one can express a given change in fire risk in a building in monetary terms in the same way as demonstrated above. Note that the abstract matter of “risk” is expressed here in a concrete way by use of a scale (the monetary scale) that the decision maker is familiar with.

One can employ various assumptions regarding the certainty equivalent of the occurrence of more than one fire. In practice, this value will depend on the decision maker’s preferences and therefore it is difficult to provide any general way of calculating it. It is assumed here that if the decision maker regards the occurrence of a single fire as being equal to $\$x$, then $\$k \cdot x$ is the monetary value being deemed equivalent to the occurrence of k fires. Note that it is possible to investigate in detail, for a specific decision maker, the value that is seen as being equivalent to a set of fires using Multi-objective utility theory, which is discussed in [5].

The assumption referred to above enables one to construct a measure of risk resembling that suggested by Hall and Sekizawa [2]. The measure of risk suggested here is the certainty equivalent of the risky situation involving the decision maker’s exposing himself/herself to the consequences of fire in the building in question during the period of time specified (for example, a year). This measure provides an answer to the question “How much should the decision maker be willing to pay in order to avoid the consequences of possible fires in the building during this period of time?”. In analogy to the term “expected loss”, this value will be termed the risk-adjusted expected loss (RAEL). “Risk-adjusted” is added to the expression to emphasise the fact that, in calculating this sum the decision maker’s risk attitude and preferences are taken into account (Cozzolino [6] used the term “Risk adjusted cost” to denote this measure). The definition of risk-adjusted expected loss used here is given in equation (3), where λ is the expected number of fires in the building during the period of a year, t is the number of years the decision maker’s evaluation of the risk involved, $P(s_i)$ is the probability of fire

scenario s_i , given that a fire has occurred, $U(s_i)$ is the utility of the consequences (expressed in monetary terms) associated with fire scenario s_i , and $U^I(\cdot)$ is the inverted utility function.

$$RAEL = \lambda \cdot t \cdot U^{-1} \left(\sum_{i=1}^n P(s_i) \cdot U(s_i) \right) \quad (3)$$

In estimating the risk-adjusted expected loss, one can discount the value of consequences that occur in the future. How that can be done is not dealt with in the present paper but is described in [5]. Note that if the decision maker is risk-neutral, the risk-adjusted expected loss is equal to the expected loss.

3. Using Bayesian networks to measure fire risk

Having identified a measure of risk that can be used to describe the level of risk in a particular building, one needs to consider ways in which it can be obtained. It is also important to consider how it would be used. Seen in isolation, risk-adjusted expected loss has no clear meaning to a decision maker. It represents the monetary value that the decision maker should, according to Bayesian decision theory, regard as being equivalent to the possible losses due to fire during a given period of time. It is not easy, however, to state whether a particular value on such a measure represents a high or a low risk. Doing that requires that one compare the risk measure obtained for the building in question with something else. One could compare it with a risk measure of the same type obtained for some other building so as to be able to compare the relative risks involved. One can also obtain a measure of *change* in the level of risk. Estimating the expected loss and the risk-adjusted expected loss for a number of consecutive time periods allows one to determine whether these measures are increasing or decreasing.

An increase or a decrease in risk has a clear meaning to the decision maker. The risk-adjusted expected loss is the monetary value the decision maker should be willing to pay in order to avoid the possible consequences of fires during a particular time period if adhering to the principle of maximising expected utility. Thus, if the risk-adjusted expected loss changes, the decision maker should evaluate that change in

risk as being equivalent to the monetary difference between the risk-adjusted expected loss before and that after the change has taken place.

Equation (3) indicates how the risk-adjusted expected loss (RAEL) can be calculated. As can be seen, in order for RAEL to change, a change needs to occur either in the expected number of fires per year in the building (λ), the probabilities of different fire scenarios given that a fire has occurred ($P(s_i)$), or the decision maker's preferences as represented by the utility function $U(\cdot)$. In estimating the amount of change that has occurred in the expected number of fires or in the probability of a particular fire scenario, given that a fire has occurred, both the decision maker's risk attitude and the time period considered in the calculations (t in equation 3) are assumed to be constant.

It is suggested here that in measuring changes in the expected number of fires during a specific period of time and in the probability of a particular fire scenario, given that a fire has occurred, use be made of Bayesian networks. Although these have been used for the probabilistic analysis of fire safety earlier (see [7], [8], and [9]), their use there was not, as is the case here, for the continual revision of fire risk measures by means of the updating of Bayesian networks.

In constructing a Bayesian network, one starts by specifying the network's structure. The network shown in Figure 2 represents a model involving different events that can influence the outcome of a fire in an electronic manufacturing facility belonging to the firm ABB (the network was created using the software Hugin Researcher, [10]). The risk analysis of this manufacturing facility was originally part of a decision analysis concerned with whether to invest in a sprinkler system. Here, however, the same analysis is used to exemplify the use of Bayesian networks for the continual revision of fire risk measures.

As can be seen in Figure 2, the variables Area, Sprinkler and Detection have no parent variables. Thus, they are assumed to not be dependent upon the states of the other variable in the network. The Area variable represents the area in which the fire starts, the Detection variable indicates whether the automatic smoke detection system has detected the fire, and the Sprinkler variable indicates whether the sprinkler system has succeeded in extinguishing the fire. The structure of the network, including the direction of the arcs, provides information on the direct

causal effects of one variable on another. Thus, one can see that the probability that the employees will succeed in extinguishing a fire (the variable Employees) is affected by the area in which the fire occurs and by whether the smoke detection system has detected the fire. The fire department can, as indicated, either extinguish a fire quickly, extinguish it slowly or not succeed in extinguishing it at all. A fire that is extinguished quickly is defined as one extinguished by the first squad arriving at the factory. A fire that is extinguished slowly is one that cannot be extinguished by the first squad, which has to wait for reinforcements before being able to extinguish it. A fire which is not extinguished involves the area in which the fire started being fully destroyed, there also being the possibility that the fire will spread to other areas. The probabilities of the different states “Quickly”, “Slowly”, and “Not extinguish” is affected by the area in which the fire started, what potential the fire has, and whether the detection system detected the fire or not. The probabilities of the different potential of the fire depend upon the area in which the fire has started. The fire potential, as indicated, can be either be small, medium or large. A large fire potential means that if none of the extinguishing operations are successful the fire will spread so as to involve the entire area in which it started. A medium fire potential implies that if none of the extinguishing operations are successful the fire will spread so as to involve a certain part of the area in which it started but will not spread beyond that part, although it might destroy a particular machine, for example. A small fire potential implies that the fire will be of limited scope and will not cause the company any serious losses.




Figure 2 Bayesian network representing the different possible fire scenarios in the ABB building.

The strength of the causal effect that the arrows in the network indicate can be determined by use of the conditional probability tables presented in Table 2. Note that those tables pertain to a specific area in the factory, the A workshop, which is an area in which ABB produces equipment for the measurement of different forces. Thus, the applicability of the probability tables presented in Table 2 is conditional upon the Area-variable being in the “A workshop” state.

Table 2 Conditional probability tables pertaining to the Bayesian network shown in Figure 2.

Detection	
Yes	0.98
No	0.02
Potential	
Small	0.8
Medium	0.1
Large	0.1
Sprinkler	
Yes	0.96
No	0.04

Table 2 (continued)

Detection	Yes		No	
Employees				
Yes	0.6		0.5	
No	0.4		0.5	

	Yes		No	
Potential	Medium	Large	Medium	Large
Fire department				
Quickly	0.8	0.3	0.7	0.1
Slowly	0.2	0.05	0.3	0.01
Not extinguishing	0	0.65	0	0.89

The probability tables pertaining to a Bayesian network can be used to calculate the probabilities of the different fire scenarios ($P(s_i)$) that are needed for calculating the risk-adjusted expected loss in equation (3). Since a given fire scenario can be defined over the states of the different variables in the Bayesian network, the probability of a fire scenario can be calculated by use of equation (4), in which the names of the variables have been abbreviated.

$$P(De, Po, Sp, Em, Fd) = P(De)P(Po)P(Sp)P(Em|De)P(Fd|De, Po) \quad (4)$$

The consequences, given that a fire has occurred in the A workshop, are summarised in Table 3. Since for many of the combinations of the different states of the variables the consequences are the same, only that part of the table of consequences in which rather high losses are involved is presented. For example, regardless of whether the employees succeed in extinguishing the fire or the sprinkler system succeeds at it, the consequence is assumed to be equivalent to a loss of 0.01 \$ million¹. Also, if the potential of the fire is small, the losses are estimated to be 0.002 \$ million.

¹ The monetary sums given in the paper, originally in Swedish crowns (SEK), were converted to US dollars at the rate of \$1 to 10 SEK.

Table 3 Consequences and utilities associated with different fire scenarios.

Employees	Not extinguishing				
Sprinkler	Not extinguishing				
Potential	Large			Medium	
Fire department	Quick	Slow	Not	Quick	Slow
Consequences (\$ million)	-0.01	-1.5	-30	-0.01	-1.5
Utility (Risk-neutral)	0.99967	0.95000	0	0.99967	0.95000
Utility (Risk-seeking)	0.99933	0.90250	0	0.99933	0.90250
Utility (Risk-averse)	0.99993	0.97254	0	0.99993	0.97254

Using the table of consequences together with tables of conditional probabilities enables one to calculate expected losses. Alternatively, if the consequences are translated into utility values, the risk-adjusted expected losses can be calculated.

The last three lines in Table 3 show the utilities associated with the different consequences, those which apply depending on which of the three utility functions in Figure 1 best characterises the decision maker. Using the Bayesian network shown in Figure 2 together with the probability tables in Table 2 and Table 3 allows one to calculate the certainty equivalent of a *single* fire. If the decision maker is risk-seeking, the result is -\$20,000, whereas if he/she is risk-averse the result is -\$101,780. For a risk-neutral decision maker the result is -\$36,910.

Note that nothing has been said thus far about the uncertainty regarding the probability estimates and estimates of the consequences. One could expect, for example, the estimates pertaining to the variables Potential, Employees and Fire Department to be considered uncertain due to statistical information regarding such variables generally being scarce. Modelling uncertainty of this type is the key to updating the network by use of assessments by experts.

4. Updating Bayesian networks

An important aspect of Bayesian networks that makes them particularly useful in the present context when a continual revision of risk

measurements is sought is the fact that they are very easy to update when new information is received. New information can be that concerning a fire that has occurred in the building, for example. In such a case, updating of the conditional probability distributions represented by the variables in the Bayesian network can readily be carried out by use of fractional updating, described later in the paper (see also [11]). In some buildings, however, fires occur so seldom that one may want to use other sources of information, such as annual inspections of the fire protection available. The present section deals with how Bayesian networks can be updated on the basis of judgements by experts.

Authors have approached the problem of incorporating expert judgement into a decision maker's body of knowledge in various ways. Genest and Zidek [12] review various procedures for aggregating assessments by experts. Apostolakis and Mosleh [13] suggest a method in which a point estimate by an expert is treated as being evidence and the expert's credibility is modelled explicitly. Apostolakis and Mosleh's method is intended for use in adjusting estimates of the frequency of reactor core meltdowns. The method suggested for use in the present context involves a technique similar to that the two authors just referred to employed. It will be shown how the model can be used in conjunction with Fractional updating and Fading together with Bayesian networks (see [11], pp. 87-91).

In the Bayesian network shown in Figure 2, no modelling of uncertainty regarding probability values is performed. In updating the Bayesian network by use of experts, however, consideration of such uncertainty is necessary for modelling. The procedure involved in modelling the uncertainty regarding the probabilities for the network as a whole as presented in Figure 2 would be difficult to show here, however.

For this reason, the major emphasis will be placed on showing how the modelling of uncertainty can be performed for a single variable. The procedure is the same when modelling uncertainty for the other variables as well. The variable of interest is denoted as I and its parent variables as A and B . Note that more than two parent variables can be employed in any given case.

The fractional updating method allows one to model the uncertainty regarding the values of $P(I|A,B)$ by assuming $P(I|a_i,b_j)$ to be related to frequencies obtained from a fictitious sample of s cases. The probability of each state of the variable I , given that A and B are known, is

determined in accordance with Table 4, where n_1, n_2, \dots, n_n are fictitious cases ($s = n_1 + n_2 + \dots + n_n$). A small sample size, s , implies the uncertainty regarding the $P(I|a_i, b_j)$ values to be high.

Table 4 *Probability table for a variable termed I , given that its parent variables A and B are in the states a_i and b_j , respectively.*

I	P(I a_i, b_j)
i_1	n_1/s
i_2	n_2/s
\vdots	\vdots
i_n	n_n/s

If a fire occurs in the building in question and the state of any variable of interest and of its parent variables are recorded, one can use this information in updating the probability table of the I -variable. If the states of the parent variables are $A = a_i$, $B = b_j$, and $I = i_l$, both the sample size and the cases themselves are updated in the following way: $s' = s + 1$, $n'_1 = n_1 + 1$, $n'_2 = n_2, \dots, n'_n = n_n$, where s' is the updated sample size and n'_i the updated cases.

If no fires in the building in question have occurred, however, and the decision maker wants to update the information concerning $P(I|A, B)$ by use of an annual inspection by a fire safety expert, for example, the problem becomes more complicated.

The task then becomes that of incorporating the expert's estimate of $P(I|A, B)$ into one's previous body of knowledge. As noted above, authors have differed in their manner of approaching this problem. Here a method based on Apostolakis and Mosleh's approach [13] to the incorporation of new evidence into one's previous body of knowledge will be employed. A basic idea behind their method is to have an expert provide a point estimate of the parameter in question and to interpret this estimate as the observed result of an experiment. In the present context, a point estimate provided by an expert would be interpreted as if the expert had actually observed one or more fires and noted the state there of the parameters of interest.

Assume that the expert provides $P(I|a_i, b_j)^*$ as an estimate of the probabilities of concern. That estimate can be interpreted as though the expert had observed s^* fires in which the parent variables were in the states $A = a_i$ and $B = b_j$, respectively, in n_k^* of the fires the variable of interest I being in the state i_k . The relationship between $P(i_k|a_i, b_j)^*$, n_k^* and s^* is given in equation (5).

$$P(i_k|a_i, b_j)^* = \frac{n_k^*}{s^*} \quad (5)$$

The decision maker needs to assign values to s^* for expressing his/her confidence in the expert. For example, if the decision maker has strong confidence in the expert, the s^* -value should be high. Viewing the expert's point estimate in this way allows it to be incorporated into the decision maker's previous body of knowledge by use of fractional updating. Note that since the sample size s^* is fictitious, it can have non-integer values.

After the expert's estimate $P(I|a_i, b_j)^*$ has been obtained, the new sample size is $s' = s + s^*$, and $n'_1 = n_1 + n_1^*$, $n'_2 = n_2 + n_2^*$, ..., $n'_n = n_n + n_n^*$.

A potential problem in using the fractional updating method is that, when the conditional probabilities of the system being modelled change over time, previous counts of n_1 , n_2 , etc. may prevent the model from taking adequate account of the changes that have occurred. Suppose an expert provides on repeated occasions his/her assessment of the reliability of a particular fire-rated wall in stopping the spread of a fire. If, for some reason, the wall's reliability in this respect changes significantly and the expert takes note of this, the effect this has on the overall assessment of the wall may be only marginal, due to the numerous past assessments of the wall the expert has made having pointed in the opposite direction. To avoid this problem, one can employ fading. This involves introducing a factor (a "fading factor") that reduces the effect of earlier assessments or of earlier experience, which might be said to "fade away", so that the overall assessment arrived at is largely based, or based to a considerable extent, on recent experience [11]. More precisely, when new cases are observed, either through actual observation or through obtaining expert judgements of them, the size of the updated sample is treated not as being equal to that of the old sample

plus one, but instead as being equal to that of the old sample, times the fading factor, plus one. Denote the fading factor as q and assign it a value of between 0 and 1. Note that $s' = qs + s^*$, and that $n'_1 = qn_1 + n_1^*$, $n'_2 = qn_2 + n_2^*, \dots, n'_n = qn_n + n_n^*$. Use of a fading factor of less than 1 results in past experience "fading away" exponentially. The lower the fading factor is, the more rapidly past experience fades away

In using the fractional updating method described above, both *global* and *local* independence are assumed [11]. Global independence involves the uncertainty regarding the probabilities of a variable $P(I|A,B)$ being independent of the uncertainty regarding the probabilities of other variables. This means that changing the probability table of a variable has no effect on the probability tables of the other variables. Local independence, in turn, involves the probability distributions of a variable for different parent configurations being independent of each other. Let (a_i, b_j) and (a'_i, b'_j) represent two different configurations for the parents of the variable in question. The uncertainty regarding $P(I|a_i, b_j)$ is independent then of the uncertainty regarding $P(I|a'_i, b'_j)$.

The achievement of global independence in the present context appears very likely. For example, the uncertainty regarding the probability that a water sprinkler system will operate in case of fire can surely be modelled independent of the uncertainty regarding the probability that a fire will occur in a particular area of a building. Local independence, however, can be more difficult to verify. Attempts to do so might involve, for example, investigating whether the uncertainty regarding the probability that the employees will succeed in extinguishing a fire in one area is affected by the uncertainty regarding the probability that they will succeed in extinguishing a fire in another area. As far as the author is aware, no such investigations, which are outside the scope of the present paper, have been reported.

5. An example of continual updating of fire risk measurements

Consider the Bayesian network shown in Figure 2. Assume that one wants to use the method presented above for measuring changes in risk during a particular period of time, five years, for example. To do this, one needs to determine which of the different variables should be updated during that time. Assume that it is only the variable Potential which is updated. Updating it would require that one establish the credibility of the values in the tables presented in Table 2, which involves assigning the probability table the fictitious sample sizes s . Assume that the initial sample size is set to consisting of 5 fires and that the relative numbers of fires of small, medium and large fire potential are those shown in Table 5. Assume too that an expert performs 5 annual assessments of the probabilities of the different fire potentials in the area in question and that this results in the assessments given in Table 6.

Table 5 *The initial counts pertaining to the probability table for the variable Potential.*

Potential	n
Small	4
Medium	0.5
Large	0.5

Table 6 *The expert's assessments of the probabilities of the different levels of fire potential in the A workshop for different years.*

Potential	Year				
	1	2	3	4	5
Small	0.7	0.7	0.75	0.75	0.75
Medium	0.2	0.2	0.2	0.2	0.2
Large	0.1	0.1	0.05	0.05	0.05

Assume that the decision maker has chosen to consider each of the annual assessments provided by the expert as being equivalent to a fictitious sample of 1 fire, so that $s^* = 1$. Assume as well that the decision maker decides to use a fading factor of 0.8. The initial probability table for the variable Potential provided in Table 2 can then be updated, using the expert's estimates.

Considering the expert's annual estimates as being equal to 1 fire implies the initial sample size $s = 5$ being updated to $s' = 5 \cdot 0.8 + 1 = 5$. Thus, the sample size continues to be 5 fires, due to the use of fading. The initial counts given in Table 5 are updated as follows:

$$\begin{aligned} n'_{Small} &= qn_{Small} + n_{Small}^* = 0.8 \cdot 4 + (1 \cdot 0.7) = 3.9 \\ n'_{Medium} &= qn_{Medium} + n_{Medium}^* = 0.8 \cdot 0.5 + (1 \cdot 0.2) = 0.6 \\ n'_{Large} &= qn_{Large} + n_{Large}^* = 0.8 \cdot 0.5 + (1 \cdot 0.1) = 0.5 \end{aligned}$$

These updated counts are used to provide an updated estimate of the probability of a fire of Small, Medium and Large potential. The result is shown in the column in Table 7 representing year 1. Making use of the expert's estimates during the 4 years following that (Table 6) allows the updated estimate of the probabilities of the different potentials to be calculated for each year (see Table 7).

Table 7 Updated estimates of the probability of a Small, a Medium and a Large fire potential in the A workshop.

Potential	Year				
	1	2	3	4	5
Small	0.780	0.764	0.761	0.759	0.757
Medium	0.120	0.136	0.149	0.159	0.167
Large	0.100	0.100	0.090	0.082	0.076

Since the variable Potential is part of the Bayesian network representing the possible development of a fire in the A workshop one can utilise the updated probability tables for the variables so as to also update the risk-adjusted losses. Assume that the risk-adjusted expected loss is calculated for a ten-year period (t in equation (3) is 10) and that the expected number of fires during any given year is 0.1 (λ in equation (3) is 0.1). On the basis of these assumptions one can calculate the risk-adjusted expected loss for each of the five years the expert provided the estimates for shown in Table 6. Figure 3 shows risk-adjusted expected loss as a function of time. The three graphs presented there correspond to the different risk attitudes, that the utility functions presented in Figure 1 represents.

Figure 3 shows the fire risk in the A workshop to have decreased during the five years for which measurements were performed. The magnitude

of the risk reduction differs depending upon whether the decision maker's risk-attitude can be characterised as being risk-averse, risk-seeking or risk-neutral. If the decision maker is risk-neutral, the reduction in risk is worth approximately \$7,000 to him/her. If, on the other hand, his/her risk-attitude is characterised by the risk-seeking utility function shown in Figure 1, the risk reduction is worth approximately \$3,000. If it is characterised by the risk-averse utility function shown there, the reduction is worth approximately \$19,600.




Figure 3 The risk-adjusted expected loss that results.

6. Summary, discussion and future work

The paper discusses the measurement of fire risk in a specific industrial facility. Use of a measure of risk based on decision theory is suggested, a risk measure termed risk-adjusted expected loss. It represents the monetary value that the building owner, or the decision maker, should be willing to pay in order to avoid the negative consequences due to fire in the building during a given period of time. It is shown how this measure can be used in conjunction with a Bayesian network for measuring *changes* in risk over time through utilising fractional updating and fading. The changes are measured by updating the Bayesian network on the basis of information concerning fires that have occurred in the building in question.

Since in most buildings fires do not occur very often, information concerning fires in the building considered is likely to be scarce. It is suggested that, due to this, annual inspections by fire experts be used for updating the Bayesian networks involved in this way updating the measure of fire risk as well. A model in which expert estimates are treated as evidence and used to update previously made estimates is presented in the paper.

Although the paper illustrates how expert assessments can be used so as to update a risk measure continually, it does not address issues concerning bias in expert assessments and the quality they possess. One might argue, with support of empirical results (see [14], pp. 533-544, for example), that humans are poor probability assessors and that we are subject to a very definite set of biases in estimating probabilities. The decision maker can endeavour to take account of such biases. Developing a model that facilitates the decision maker's doing this could be a goal for future research.

Another aspect of such a model that would be useful to develop is to provide the expert means of readily expressing his/her uncertainty regarding estimates that are made. In the model presented here, doing this is not possible, no distinction is being made of whether an expert is very confident in his/her assessments or is very uncertain about them. It should be possible to extend the model in such a way that the fictitious sample size can be adjusted in a manner appropriate to the confidence the expert has in his/her assessments.

7. Acknowledgements

The author would like to thank The Swedish Fire Research Board (BRANDFORSK) for funding the research on which the present paper is based.

8. References

1. Johansson, H., "Decision analysis concerned with investments in fire safety", *Seventh International Symposium on Fire Safety Science*, Worcester, 2002.
2. Hall, J.R., Sekizawa, A., "Fire Risk Analysis: General Conceptual Framework For Describing Models", *Fire Technology* 27, pp. 33-53, 1991.

3. Luce, R.D., Raiffa, H., *Games and decisions Introduction and critical survey*, John Wiley & Sons, New York, 1957.
4. Farquhar, P. H., "Utility assessment methods", *Management Science*, Vol. 30, No. 11, pp. 1283-1300, 1984.
5. Johansson, H., *Decision Analysis in Fire Safety Engineering – Analysing Investments in Fire Safety*, Doctoral thesis, Lund University, Department of fire safety engineering, Lund, 2003.
6. Cozzolino, J. M., "A Method for the Evaluation of Retained Risk", *The Journal of Risk and Insurance*, Vol. XLV, No. 3, pp. 449-471, 1978.
7. Gulvanessian, H., Holický, M., Cajot, L.-G. and Schleich, J.-B., "Probabilistic analysis of fire safety using Bayesian causal network", *Safety and Reliability*, Schueller & Kafka (eds), Balkema, Rotterdam, 1999.
8. Holicky, M., Schleich, J.-B., "Modelling of a Structure under Permanent and Fire Design Situation", *Proc. of Risk and Reliability – Trends in Engineering*, Safety and Risk in Engineering, Zürich, pp. 789-794, 2001.
9. Holicky, M., Schleich, J.-B., "Fire Safety Assessment using Bayesian Casual Network", *Proc. of Foresight and Precaution Conf., ESREL2000*, Edingburgh, A.A. Balkema, Rotterdam, pp. 1301-1306, 2000.
10. Hugin Researcher Version 6.0, Hugin Expert A/S, Aalborg, Denmark.
11. Jensen, F. V., *Bayesian Networks and Decision Graphs*, Springer-Verlag, New York, 2001.
12. Genest, C., Zidek, J.V., "Combining Probability Distributions: A Critique and an Annotated Bibliography", *Statistical Science*, Vol. 1, No. 1, pp. 114-135, 1986.
13. Apostolakis, G., Mosleh, A., "Expert Opinion and Statistical Evidence: An Application to Reactor Core Melt Frequency", *Nuclear science and engineering*: 70, pp. 135-149, 1979.
14. Von Winterfeldt, D., *Decision Analysis and Behavioral Research*, Cambridge University Press, Cambridge, 1986.

Appendix

Appendix A: Bayesian updating of a Dirichlet distribution

From equation (A.1), which provides a general expression of Bayes' theorem for the case of more than two possible alternative events, one can see that the Dirichlet distribution is the conjugate form of the Multinomial-distribution. The posterior distribution $f''(p_1, \dots, p_k)$ is derived by multiplying the likelihood $L(E|p_1, \dots, p_k)$ by the prior distribution $f'(p_1, \dots, p_k)$. The result is then divided by the normalisation factor c so as to ensure that the posterior distribution is a true probability distribution.

$$f''(p_1, \dots, p_k) = \frac{1}{c} \cdot L(E|p_1, \dots, p_k) \cdot f'(p_1, \dots, p_k) \quad (\text{A.1})$$

Use of distributions corresponding to the prior distribution and the likelihood, respectively, for the case in which the likelihood can be represented by a multinomial distribution and the prior by a Dirichlet distribution allows equation (A.1) to be written in a different form, as equation (A.2). There, p_1, \dots, p_k , and $(1-p_1-\dots-p_k)$ are the probabilities of the uncertain events of interest, for example, the probability that a fire has started in a specific area, given that a fire has occurred; r_1, \dots, r_k , and $(n-r_1-\dots-r_k)$ are the number of occurrences one has observed for a specific event; n is the total number of events observed; and $\nu_1, \dots, \nu_k, \nu_{k+1}$, are the parameters of the Dirichlet distribution, which represents the prior information regarding the probabilities in question.

$$\begin{aligned} f''(p_1, \dots, p_k) &= \\ &= \frac{1}{c} \cdot \frac{n!}{r_1! \cdot \dots \cdot r_k! (n-r_1-\dots-r_k)!} p_1^{r_1} \cdot \dots \cdot p_k^{r_k} (1-p_1-\dots-p_k)^{(n-r_1-\dots-r_k)} \cdot \\ &\cdot \frac{\Gamma(\nu_1 + \dots + \nu_{k+1})}{\Gamma(\nu_1) \cdot \dots \cdot \Gamma(\nu_{k+1})} p_1^{\nu_1-1} \cdot \dots \cdot p_k^{\nu_k-1} (1-p_1-\dots-p_k)^{\nu_{k+1}-1} \end{aligned} \quad (\text{A.2})$$

Rearranging the terms in equation (A.2) yields equation (A.3).

$$\begin{aligned}
 f''(p_1, \dots, p_k) &= \\
 &= \frac{1}{c} \cdot \frac{n!}{r_1! \cdots r_k! (n - r_1 - \dots - r_k)!} \cdot \frac{\Gamma(v_1 + \dots + v_{k+1})}{\Gamma(v_1) \cdots \Gamma(v_{k+1})} \cdot \\
 &\cdot p_1^{v_1+r_1-1} \cdots p_k^{v_k+r_k-1} (1 - p_1 - \dots - p_k)^{v_{k+1}+n-r_1-\dots-r_k-1}
 \end{aligned} \tag{A.3}$$

The factor c in equation (A.3) is defined in equation (A.4), where

$$\begin{aligned}
 S &= \left\{ (p_1, \dots, p_k) : p_i \geq 0, i = 1, \dots, k, \sum_{i=1}^k p_i \leq 1 \right\}. \\
 c &= \int_S \frac{n!}{r_1! \cdots r_k! (n - r_1 - \dots - r_k)!} \cdot \frac{\Gamma(v_1 + \dots + v_{k+1})}{\Gamma(v_1) \cdots \Gamma(v_{k+1})} \cdot \\
 &\cdot p_1^{v_1+r_1-1} \cdots p_k^{v_k+r_k-1} (1 - p_1 - \dots - p_k)^{v_{k+1}+n-r_1-\dots-r_k-1} dp_1 \dots dp_k
 \end{aligned} \tag{A.4}$$

The integral in equation (A.4) is called the Dirichlet integral and is equal to the expression shown in equation (A.5) [84].

$$\begin{aligned}
 &\int_S p_1^{v_1+r_1-1} \cdots p_k^{v_k+r_k-1} (1 - p_1 - \dots - p_k)^{v_{k+1}+n-r_1-\dots-r_k-1} dp_1 \dots dp_k = \\
 &= \frac{\Gamma(v_1 + r_1) \cdots \Gamma(v_k + r_k) \cdot \Gamma(v_{k+1} + n - r_1 - \dots - r_k)}{\Gamma(v_1 + \dots + v_k + v_{k+1} + n)}
 \end{aligned} \tag{A.5}$$

If equation (A.5) is combined with equation (A.4) and this, in turn, is inserted into equation (A.3), this results in equation (A.6).

$$\begin{aligned}
 f''(p_1, \dots, p_k) &= \frac{\Gamma(v_1 + \dots + v_k + v_{k+1} + n)}{\Gamma(v_1 + r_1) \cdots \Gamma(v_k + r_k) \Gamma(v_{k+1} + n - r_1 - \dots - r_k)} \cdot \\
 &\cdot p_1^{r_1+v_1-1} \cdots p_k^{r_k+v_k-1} (1 - p_1 - \dots - p_k)^{v_{k+1}+n-r_1-\dots-r_k-1}
 \end{aligned} \tag{A.6}$$

Equation (A.6) is a Dirichlet distribution with the parameters $(\nu_1 + r_1)$, $(\nu_2 + r_2), \dots, (\nu_{k+1} + n - r_1 - \dots - r_k)$, which implies that the Dirichlet distribution is the natural conjugate obtained when the likelihood has a multinomial distribution. The prior distribution is updated to the posterior distribution by use of the following equations:

$$\begin{aligned} \nu_1'' &= \nu_1' + r_1 \\ &\vdots \\ \nu_k'' &= \nu_k' + r_k \\ \nu_{k+1}'' &= \nu_{k+1}' + n - r_1 - \dots - r_k \end{aligned} \tag{A.7}$$

Appendix B: Estimates pertaining to the ABB analysis

Many different parameters (costs and probabilities) are involved in analysing the investment in a sprinkler system for building 358. Each parameter in the risk-model is estimated using a minimum, a most likely and a maximum value. These estimates are presented here in appendix B. The parameters considered are shown in the event tree presented in Appendix D. In which tables estimates of the different parameters appear depends on the area of the building to which they pertain, on whether they concern costs, and on whether they concern the reliability of the barriers in the building. A brief account of the circumstances under which the estimates were made is presented prior to each table. Each of the parameters concerned is assigned an abbreviation that can be used to help find the event trees in Appendix D in which it is located.

Estimation of Costs

The costs presented are ones associated with fire scenarios that are not particularly serious. The costs associated with more serious fire scenarios are dealt with in chapter 5. The event trees to which the costs pertain are presented in Appendix D. There, one can see the costs corresponding to each of the fire scenarios. In general terms, one can say that if the fire department fails to succeed in extinguishing a fire of medium potential quickly or if it slowly manages to extinguish a fire of large potential, such that the first group of fire men who arrive fail to extinguish it but by the supporting forces arriving soon after succeed in doing so, the costs of the fire are of medium size.

Table 13 Minimum, maximum and most likely values for the costs associated with fire scenarios that are not very severe. The costs are given in millions of SEK.

	Min	Most likely	Max	Abbreviation
Costs of a small fire regardless of where it started.	0.005	0.02	0.05	C9ABB
Area 1				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.05	0.1	0.15	C10ABB
Costs of a medium-sized fire.	5	11	19	C11ABB

Table 13 (continued)

Area 2				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.05	0.1	0.15	C12ABB
Costs of a medium-sized fire.	6	15	21	C13ABB
Area 3				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.025	0.05	0.075	C14ABB
Costs of a medium-sized fire.	0.05	0.1	0.2	C15ABB
Area 4	0.05	0.1	0.15	
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.1	0.2	0.5	C16ABB
Costs of a medium-sized fire.	0.05	0.1	0.15	C17ABB
Area 5				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.05	0.1	0.15	C18ABB
Costs of a medium-sized fire.	0.2	0.4	0.6	C19ABB
Area 6				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.05	0.1	0.15	C20ABB
Costs of a medium-sized fire.	2	5	18	C21ABB
Area 7				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.025	0.05	0.075	C22ABB
Costs of a medium-sized fire.	0.2	0.4	0.6	C23ABB
Area 8				
Costs of a fire extinguished by the sprinkler system or the employees, or extinguished quickly by the fire department.	0.05	0.1	0.15	C24ABB
Costs of a medium-sized fire.	8	18	38	C25ABB

Estimates of probabilities

The model used to represent different fire scenarios in the ABB building (Appendix D) contains many probabilities. Some of these concern events for which very little “objective” information is available. Objective information refers here to the results of statistical investigations of past fires. For example, no statistical information relevant to the present context is available for the probability that the employees will succeed in extinguishing a fire in a specific area.

This means that in such cases the estimates of probabilities are based solely on the judgements of the analyst and of people from ABB. This poses no difficulties from a decision analytical standpoint since the definition of probability employed there is subjective and thus represents the decision maker’s degree of belief (see chapter 5 for a brief discussion of this). In presenting estimates of the probabilities, it is useful to also present insofar as possible the information on which the estimates are based. Accordingly, each of the tables containing probability estimates is preceded by a short account of the circumstances relevant to the estimates in question.

The first two probabilities taken up are the probability that the sprinkler system will succeed in extinguishing a fire and the probability that the smoke detection system will sound the alarm, given that a fire has occurred. The estimated probability that the sprinkler will succeed in extinguishing a fire is somewhat lower than what is indicated in the general statistics presented in chapter 5. This is because the present building has a large amount of storage rack in which a fire might be difficult to extinguish. Although the sprinkler system is designed to be able to deal with storage rack fires, there are places in the building in which it is considered difficult for the sprinkler to extinguish a fire, storage racks being one of them. The probability that the smoke detection system will sound the alarm, given that a fire in the building occurs, has been estimated to be a bit higher than what is reported in the general investigations summarised in chapter 5. This is because the smoke detection system is of a modern type and is carefully maintained.

Table 14 The probability that the sprinkler system will succeed in extinguishing a fire and the probability that the smoke detection system will sound the alarm, given that a fire has occurred.

Probability	Min	Most likely	Max	Abbreviation
The probability that the sprinkler system will succeed in extinguishing a fire if one starts.	0.94	0.96	0.98	P1ABB
The probability that the smoke detection system will sound the alarm, given that a fire has occurred.	0.9	0.98	0.99	P2ABB

Distribution of fires

If a fire in the building occurs, it may be at any one of 9 different areas there. The probability of a fire occurring in a specific area, given that a fire in the building has occurred, can be assessed by considering the size of the respective area and the activities that are performed there. It would be reasonable to expect that an area being larger would involve the probability of a fire occurring there being greater (provided the activities in all the areas are the same). It would also be reasonable to expect some activities, such as those of the production of components, to involve a greater number of potential ignition sources than storage spaces and office spaces would, for example. A brief account of the activities carried on in the different areas can be found in chapter 5, as well as later in this appendix.

Table 15 The probability of a fire occurring in each of the areas.

Area	Min	Most likely	Max	Abbreviation
The new PK workshop	0.04	0.061	0.08	P4ABB
The A workshop	0.15	0.24	0.30	P5ABB
Storage area	0.01	0.017	0.03	P6ABB
ABB Training Center	0.001	0.005	0.01	P7ABB
EMC	0.001	0.009	0.02	P8ABB
The PS Workshop	0.15	0.197	0.25	P9ABB
The office	0.002	0.004	0.008	P10ABB
The PK Workshop	0.10	0.161	0.20	P11ABB
Other sections	0.20	0.306	0.35	P12ABB

Barriers

There are barriers between the different areas (illustrated in Figure 22) of the building that can stop a fire from spreading. In the model of fire spread employed here, the performance of each barrier is represented by the probability that it will stop the spread of fire from the one area to another. This model of fire spread is a very simple one. A much more complicated one could be employed instead, such as a model taking account of the length of time a barrier has been exposed to fire or to the total amount of energy it has absorbed, so as to be able to express the probability of failure as a function of time (as described in [100], for example). Such a complicated analysis is outside the scope of the case study presented here however. Employing such a model would do little to help in exemplifying how decision analysis can be applied to the analysis of investments in fire safety. Accordingly, the probability that a fire will spread from one side of the barrier to the other is simply estimated for the final state of the barrier, i.e. the probability of spread of fire through the barrier at some point during the fire scenario.

The barriers in the building are considered here to be of two major types, those that separate the building into two halves, e.g. the barriers between areas 1 and 8, 2 and 7, etc., and those that separate areas within each of the two halves. The barriers separating the building into two halves are judged generally to have a higher probability of succeeding in limiting spread of fire than the others. One can see in Table 16 that the estimates of the probability that a specific barrier will succeed in limiting the spread of fire is either close to 0.9, which is the value recommended in the BSI guide [96] for a fire-rated barrier that contains openings, or close to 0.5, which is the recommended value for walls without doors for which there is no documented fire rating. Note that the barrier separating area 9 from the rest of the building is of particularly high quality. Since the fire load in area 9 is low and ABB Automation Products has no activities in that area, fires occurring there are disregarded in the analysis.

Table 16 Minimum, maximum and most likely values for the probability that a barrier between two fire compartments can withstand a fully developed fire.

Between areas		Minimum	Most likely	Maximum	Abbreviation
1	2	0.45	0.6	0.75	P13ABB
1	8	0.7	0.9	0.95	P14ABB
2	3	0.3	0.5	0.6	P15ABB
2	4	0.3	0.5	0.6	P16ABB
2	6	0.7	0.9	0.95	P17ABB
2	7	0.7	0.9	0.95	P18ABB
2	8	0.7	0.9	0.95	P19ABB
3	4	0.3	0.5	0.6	P20ABB
3	5	0.3	0.5	0.6	P21ABB
3	6	0.7	0.9	0.95	P22ABB
4	5	0.3	0.5	0.6	P23ABB
6	7	0.3	0.5	0.6	P24ABB
7	8	0.3	0.5	0.6	P25ABB

Area 1: The New PK workshop

In the New PK workshop, ABB is assembling circuit cards in seven production lines. The area contains a large amount of electronic equipment and some storage rack space. The most serious fire scenarios would be those of a fire spreading to the storage rack space, since this will very likely result in the entire area becoming involved (unless the fire was extinguished by the employees or by the sprinkler system). Some 80 employees are working in this area.

Table 17 The minimum (“Min”), maximum (“Max”) and most likely values for the different probabilities in the model for fire spread in the new PK workshop.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.55	0.78	0.85	P26ABB
The probability of a medium fire	0.05	0.15	0.3	P27ABB
The probability of a large fire	0.03	0.07	0.15	P28ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.7	0.9	0.95	P29ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.5	0.8	0.9	P30ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.4	0.7	0.9	P31ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.1	0.4	0.6	P32ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.05	0.2	0.4	P33ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.005	0.01	0.1	P34ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.01	0.07	0.15	P35ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.001	0.005	0.1	P36ABB

Area 2: The A workshop

In the A workshop, ABB manufactures force measurement equipment that is sold primarily to the forest and steel industries. The area contains workstations for both production and testing. There are also offices within the area and storage rack spaces. Some 75 employees work in the area. There are high storage racks here.

Table 18 The minimum (“Min”), maximum (“Max”) and most likely values of the different probabilities in the model for fire spread in the A workshop.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.7	0.8	0.9	P37 ABB
The probability of a medium fire	0.05	0.1	0.4	P38 ABB
The probability of a large fire	0.05	0.1	0.15	P39 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.5	0.6	0.8	P40 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.3	0.5	0.6	P41 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.7	0.8	0.9	P42 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.5	0.7	0.8	P43 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.2	0.3	0.6	P44 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.01	0.05	0.1	P45 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.05	0.1	0.15	P46 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.001	0.01	0.05	P47 ABB

Area 3: The Storage Area

In the storage area there are large amounts of combustibles in storage racks. There are no employees working in this area.

Table 19 The minimum (“Min”), maximum (“Max”) and most likely values for the different probabilities in the model for fire spread in the Storage area.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.4	0.6	0.7	P48 ABB
The probability of a medium fire	0.15	0.3	0.4	P49 ABB
The probability of a large fire	0.05	0.1	0.2	P50 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.2	0.4	0.6	P51 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.1	0.2	0.3	P52 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.5	0.6	0.7	P53 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.2	0.4	0.5	P54 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.3	0.5	0.6	P55 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.01	0.02	0.1	P56 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.05	0.1	0.15	P57 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.005	0.01	0.1	P58 ABB

Area 4: ABB Training Center

This area consists of classrooms and offices. There are usually around 40 employees or students in this area.

Table 20 The minimum (“Min”), maximum (“Max”) and most likely values for the different probabilities in the model for fire spread in the ABB Training center.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.6	0.75	0.85	P59 ABB
The probability of a medium fire	0.1	0.2	0.3	P60 ABB
The probability of a large fire	0.01	0.05	0.1	P61 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.5	0.6	0.8	P62 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.4	0.5	0.6	P63 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.3	0.4	0.5	P64 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.2	0.3	0.4	P65 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.3	0.4	0.5	P66 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.01	0.05	0.1	P67 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.01	0.02	0.05	P68 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.005	0.01	0.1	P69 ABB

Area 5: EMC

In this area ABB tests their products and measures electromagnetic emissions. There are 6 employees working in this area.

Table 21 The minimum (“Min”), maximum (“Max”) and most likely values for the different probabilities in the model for fire spread in the EMC.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.6	0.8	0.9	P70 ABB
The probability of a medium fire	0.09	0.15	0.3	P71 ABB
The probability of a large fire	0.01	0.05	0.1	P72 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.5	0.6	0.7	P73 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.4	0.5	0.6	P74 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.2	0.25	0.4	P75 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.05	0.15	0.2	P76 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.1	0.2	0.25	P77 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.1	0.15	0.2	P78 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.01	0.05	0.1	P79 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.005	0.01	0.05	P80 ABB

Area 6: The PS workshop

In this area, there are automation systems located in metal cabinets. The assembly of the automation systems requires a large amount of electronics equipment and also a large amount of cardboard boxes and paper in which the equipment is packed. About 100 employees work in this area. There are storage racks in the area.

Table 22 The minimum (“Min”), maximum (“Max”) and most likely values for the different probabilities in the model for fire spread in the PS Workshop.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.6	0.7	0.8	P81 ABB
The probability of a medium fire	0.1	0.2	0.3	P82 ABB
The probability of a large fire	0.01	0.1	0.15	P83 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.4	0.6	0.8	P84 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.3	0.5	0.7	P85 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.6	0.8	0.85	P86 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.4	0.7	0.8	P87 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.3	0.4	0.5	P88 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.01	0.02	0.1	P89 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.01	0.1	0.15	P90 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.001	0.01	0.1	P91 ABB

Area 7: The office

The administrative staff is located in this area. The area consists primarily of a large open space divided into cubicles. There are also a number of office rooms and conference rooms. About 100 employees work in this area.

Table 23 The minimum (“Min”), maximum (“Max”) and most likely values for the different probabilities in the model for fire spread in the Office Area

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.6	0.85	0.9	P92 ABB
The probability of a medium fire	0.05	0.1	0.25	P93 ABB
The probability of a large fire	0.1	0.05	0.15	P94 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.7	0.8	0.9	P95 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.65	0.75	0.8	P96 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.3	0.4	0.5	P97 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.15	0.3	0.4	P98 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.25	0.35	0.5	P99 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.005	0.01	0.1	P100 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.05	0.1	0.2	P101 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.005	0.01	0.1	P102 ABB

Area 8: The old PK workshop

This area is very similar to the new PK workshop. There are 4 assembly lines for the production of circuit cards here. Some 360 employees work in the area.

Table 24 The minimum ("Min"), maximum ("Max") and most likely values for the different probabilities in the model for fire spread in the old PK workshop.

	Min	Most likely	Max	Abbreviation
The probability of a small fire	0.7	0.8	0.9	P103 ABB
The probability of a medium fire	0.09	0.15	0.25	P104 ABB
The probability of a large fire	0.01	0.05	0.1	P105 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>works</i>	0.7	0.8	0.9	P106 ABB
The probability that the employees succeed in extinguishing a fire, given that the smoke alarm <i>fails to work</i>	0.5	0.7	0.8	P107 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> given a <i>medium</i> fire-growth potential and that the smoke alarm <i>works</i>	0.5	0.8	0.85	P108 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>medium</i> fire-growth potential and that the smoke alarm <i>fails to work</i>	0.3	0.4	0.6	P109 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.1	0.2	0.3	P110 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>works</i>	0.005	0.01	0.1	P111 ABB
The probability that the fire department extinguishes the fire <i>quickly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.01	0.07	0.1	P112 ABB
The probability that the fire department extinguishes the fire <i>slowly</i> , given a <i>large</i> fire potential and that the smoke alarm <i>fails to work</i>	0.005	0.01	0.1	P113 ABB

Comments on some of the probability estimates

The estimates of the probabilities pertaining to the performance of the employees and of the fire department were made without access to any “objective” information. Thus, in performing the estimates presented here the analyst and the persons involved from ABB and from the fire department had no access to any general information about how often employees can be expected to extinguish a fire, given a particular fire potential. This of course results in uncertainty regarding the estimates. The reasoning on which the estimates were based will be summarised here.

As can be seen in Table 25, the estimates of the most likely value for the probability that the employees succeed in extinguishing a fire is highest for the new PK workshop, where it is 0.9 (given that the automatic detection system works). The reason for this high estimate is that the equipment for the manual suppression of fires in that area is very modern, that the number of employees per square meter is fairly high (0.014), and that the employees in this area have a better education in manual suppression than those in many of the other areas. The second highest estimates for the probability of interest is for the ABB training center, the office and the old PK workshop. The ABB training center is given a high probability because of the quality of the manual suppression equipment there, which is of a character similar to that in the new PK workshop. In addition, a large part of the area consists of classrooms that contain very little combustible material and are thus unlikely to serve as a starting point for a large or a medium-sized fire. Instead, the highest concentration of combustibles is in the part of the ABB training center in which the offices are located. In that part, there is a high concentration of persons who are well educated in manual firefighting. The office area has a higher concentration of persons per square meter, which indicates the probability of their succeeding in extinguishing a fire to be high. Since the employees there do not have any education at all in manual firefighting, however, the probability is estimated to be about same as for the ABB training center. The old PK workshop has the highest concentration of employees per square meter. Since this area, however, does not have as good manual firefighting equipment as the new PK workshop, for example, the probability estimate is lower than for that area.

In three of the areas – the A workshop, EMC and the PS Workshop – the estimate of the probability that the employees succeed in extinguishing a

fire is 0.6. The reason for the estimate there being lower than in the areas described above is that fire spread in the A workshop and in the PS workshop is estimated to be faster than in the other workshops. This is due to the presence of large amounts of combustibles in the form of cardboard and the like, and to the large number of high storage racks. Since the EMC is the area in which the number of persons per square meter is lowest, the probability that the employees there would succeed in extinguishing a fire is also low.

The lowest estimate for the probability of the employees succeeding in extinguishing a fire is for the storage area. This is because no one works in that area, making it likely that the time from the start of a fire until the manual extinguishing operation can begin would be greater than in the other areas.

Table 25 The most likely values of the probability that the employees succeed in extinguishing a fire.

Area	The probability that the employees succeed in extinguishing a fire
The new PK workshop	0.9
The A workshop	0.6
Storage space	0.4
ABB Training Center	0.8
EMC	0.6
The PS workshop	0.6
The office	0.8
The old PK workshop	0.8

In estimating probabilities related to the work of the fire department, different matters have been taken into account, for example whether the automatic detection system works, whether the areas in question are easily accessible, and the amounts of combustibles there. If the automatic detection system works, the fire department gets the alarm early and can start dealing with the fire as early as possible. Some areas of the building are difficult, because of their large size, for the fire department to reach all parts of, making it more difficult to extinguish a fire. The amount of combustibles in an area is assumed to affect the probability that the fire department will succeed in extinguishing a fire. Also, the probability is assumed to be affected by the configuration of the fuel.

Appendix C: Estimates pertaining to the Avesta Sheffield analysis

There are many uncertain parameters in analysing the investment in a sprinkler system for Avesta Sheffield's cold-rolling mill. The uncertainties regarding these parameters are presented here in terms of a maximum, a most likely and a minimum value estimated for each of the parameters. The estimates were performed in discussion between the analyst (the author) and members of the personnel of Avesta Sheffield and of the fire department in Eskilstuna.

Estimation of the costs

The costs associated with the more serious fire scenarios are presented in chapter 5. Table 26, in contrast, presents the costs associated with the fire scenarios that are much less serious. A fire scenario of medium seriousness in a machine located somewhere in the cold-rolling mill is assumed to destroy parts of the machine, although it is estimated that a large part of the machine would remain intact. Thus, the costs associated with such a scenario are not considered to be as high as those associated with the complete destruction of a machine with which the costs in Table 8 are associated. Instead, the costs of a fire of medium seriousness are assumed to amount to a certain part of the costs of the total destruction of the machine in question (given in Table 8). The ratio of the cost of a medium-sized fire in a machine to the costs of complete destruction of the machine is estimated to be somewhere between 0.05 and 0.25, the most likely appearing value being 0.1. This parameter is called the "Ratio".

Table 26 Maximum, minimum and the most likely values for the costs associated with the fire scenarios that are not particularly severe.

Fire scenario	Min	Most likely	Max	Abbreviation
A small fire in a machine.	0.025	0.05	0.075	C30Av
A medium-sized or large fire in a machine that is extinguished by the employees.	0.05	0.1	0.15	C31Av
Small fires that do not occur in the machines.	0.005	0.01	0.025	C32Av
A medium-sized fire that does not occur in any of the machines.	0.025	0.05	0.075	C33Av
A large fire that does not occur in a machine and is extinguished by the fire department.	0.05	0.1	0.15	C34Av

Estimation of probabilities

Distribution of fires

A fire in the building may occur in any one of 9 different areas. The probability of a fire occurring in a specific area, given that a fire somewhere in the building has occurred, is determined on the basis of the size of the respective area and the nature of the activities performed there. Generally speaking, one should expect a larger area to involve a higher probability of a fire (provided that the activities in the areas compared are about the same). Also, areas that contain large amounts of combustibles and many potential ignition sources, such as machines operating at high temperatures, can be expected to have a higher probability of fire. A brief account of the activities that take place in the different areas can be found in chapter 5, as well as later in this appendix.

Although areas 1 and 4 are of approximately the same size as area 3, the latter area is judged to have a higher probability of fire than the others. The reason is that two cold-rolling mills are located in this area and that these machines are known to cause many fires. Area 5 also has a high probability of fire. The reason for this is that the abrasive-belt grinder, which is a known source of fire, is located there. Note that the smaller areas in the building are treated as constituting a single area (area 6). In

the event tree that pertains to area 6, however, the separate areas of which it consists are modelled separately.

Table 27 *Estimates of the probability of a fire occurring in a specific area in the Avesta Sheffield factory.*

	Min	Most likely	Max	Abbreviation
Area 1	0.03	0.07	0.15	P3Av
Area 2	0.01	0.05	0.15	P4Av
Area 3	0.25	0.37	0.55	P5Av
Area 4	0.1	0.12	0.2	P6Av
Area 5	0.15	0.28	0.35	P7Av
Area 6	0.05	0.11	0.15	P8Av

Barriers

The barriers in the Avesta Sheffield factory are judged to be lower in quality than those in the ABB factory. This is because none of the barriers there are fire rated, although some of them appear to have a fire resistance about the same as that of a fire-rated wall. Note that the BSI Guide [95] recommends that the reliability of fire rated walls that have doors is 0.9 and that the reliability of walls without any documented fire rating is 0.5.

Table 28 *The probability that a particular barrier will succeed in limiting the spread of fire.*

Barrier areas	between	Min	Most likely	Max	Abbreviation
1	3	0.2	0.4	0.5	P143Av
1	4	0.3	0.5	0.6	P144Av
2	3	0.6	0.7	0.8	P145Av
3	4	0.4	0.5	0.6	P146Av
3	5	0.4	0.5	0.6	P147Av
3	6	0.6	0.8	0.9	P148Av
3	7	0.6	0.8	0.9	P149Av
3	8	0.6	0.8	0.9	P150Av
3	9	0.6	0.8	0.9	P151Av
4	5	0.6	0.8	0.9	P152Av
3	Oil room	0.6	0.8	0.9	P153Av
2	Pallet storage	0.4	0.5	0.6	P154Av

The development of a fire in each of the different areas.

The event trees describing the different fire scenarios are presented in Appendix E. Each of the areas has a particular event tree associated with it, each tree having various probabilities associated with it. In this section, all the probability estimates pertaining to a particular area are presented. A brief account of each of the areas is provided first.

Probability estimates pertaining to Area 1

This is an area containing machines for cutting steel. There is also a smoothing roller there. Besides these machines, there are finished products of steel and large amount of packaging material such as paper and wooden-pallets located there.

Table 29 *Probability estimations related to fire development in area 1.*

	Min	Most likely	Max	Abbreviation
Probability of a fire in one of the machines	0.5	0.6	0.9	P9Av
Probability of a fire in the smoothing roller, given a fire in a machine	0.3	0.4	0.6	P10Av
Probability of a fire in Cutter 1, given a fire in a machine	0.2	0.4	0.45	P11Av
Probability of a fire in Cutter 2, given a fire in a machine	0.1	0.2	0.25	P12Av
<i>Smoothing roller</i>				
Small fire	0.7	0.8	0.85	P13Av
Medium-sized fire	0.14	0.19	0.25	P14Av
Large fire	0.001	0.01	0.05	P15Av
<i>Cutter 1</i>				
Small fire	0.79	0.83	0.9	P16Av
Medium-sized fire	0.08	0.16	0.2	P17Av
Large fire	0.001	0.01	0.02	P18Av

Table 29 (continued)

<i>Cutter 2</i>				
Small fire	0.85	0.93	0.95	P19Av
Medium-sized fire	0.03	0.06	0.13	P20Av
Large fire	0.001	0.01	0.02	P21Av
<i>Other fires</i>				
Small fire	0.75	0.8	0.9	P22Av
Medium-sized fire	0.09	0.19	0.2	P23Av
Large fire	0.001	0.01	0.05	P24Av
<i>The probability of the employees succeeding in extinguishing a fire</i>				
In the smoothing roller	0.5	0.7	0.8	P25Av
In cutter 1	0.6	0.8	0.85	P26Av
In cutter 2	0.6	0.8	0.85	P27Av
In some other location	0.7	0.8	0.9	P28Av
The probability of the fire department succeeding in extinguishing a large fire	0.6	0.95	0.98	P29Av

Probability estimates pertaining to Area 2

Area 2 is a small area with a single cutting-machine, one which is newer than the machines in area 1. In addition to the machine, the area contains packaging-material and wooden pallets.

Table 30 *Probability estimates related to fire development in area 2.*

	Min	Most likely	Max	Abbreviation
Probability of a fire in a machine	0.7	0.82	0.9	P30Av
<i>Cutter 3</i>				
Small fire	0.8	0.9	0.95	P31Av
Medium-sized fire	0.02	0.08	0.17	P32Av
Large fire	0.001	0.02	0.03	P33Av
<i>Other fires</i>				
Small fire	0.7	0.83	0.9	P34Av
Medium-sized fire	0.09	0.16	0.28	P35Av
Large fire	0.001	0.01	0.02	P36Av
<i>The probability of the employees succeeding in extinguishing a fire</i>				
In cutter 3	0.7	0.8	0.9	P37Av
In other locations	0.5	0.8	0.85	P38Av
The probability of the fire department succeeding in extinguishing a large fire.	0.6	0.95	0.98	P39Av

Probability estimates pertaining to Area 3

Area 3 contains three machines: two cold-rolling mills, and a strip coiling machine. Cold-rolling mill 1, located there, is the oldest of the cold-rolling mills. It contains a water sprinkler system. Cold-rolling mill 2 contains a water sprinkler system and a CO₂-system. Both of these machines have caused fires in the past, although fortunately the automatic extinguishing systems there have succeeded in extinguishing the fires. Besides the machines, cardboard, wood and plastics are stored in the area.

Table 31 Probability estimates concerning fire development in area 3.

	Min	Most likely	Max	Abbreviation
The probability of a fire in a machine, given that a fire has occurred	0.8	0.96	0.98	P40Av
Probability of a fire in cold-rolling mill 1, given that a fire has occurred in a machine	0.1	0.25	0.3	P41Av
Probability of a fire in cold-rolling mill 2, given that a fire has occurred in a machine	0.6	0.7	0.85	P42Av
Probability of a fire in the strip coiling machine, given that a fire has occurred in a machine	0.01	0.05	0.1	P43Av
<i>Fire potential in the cold-rolling mill 1</i>				
Small fire	0.5	0.66	0.8	P44Av
Medium-sized fire	0.2	0.31	0.5	P45Av
Large fire	0.01	0.03	0.05	P46Av
<i>Fire potential in the cold-rolling mill 2</i>				
Small fire	0.4	0.52	0.7	P47Av
Medium-sized fire	0.29	0.44	0.53	P48Av
Large fire	0.01	0.04	0.07	P49Av
<i>Fire potential in the strip coiling machine</i>				
Small fire	0.8	0.88	0.95	P50Av
Medium-sized fire	0.04	0.11	0.18	P51Av
Large fire	0.001	0.01	0.02	P52Av
<i>Fire potential in other fires</i>				
Small fire	0.7	0.78	0.85	P53Av
Medium-sized fire	0.1	0.21	0.3	P54Av
Large fire	0.001	0.01	0.04	P55Av
<i>The probability that the employees succeed in extinguishing a fire</i>				
In cold-rolling mill 1	0.6	0.8	0.9	P56Av

Table 31 (continued)

In cold-rolling mill 2	0.6	0.8	0.9	P57Av
In the strip coiling machine	0.7	0.85	0.95	P58Av
A fire at some other location	0.7	0.85	0.95	P59Av
The probability that the fire department succeeds in extinguishing a large fire.	0.5	0.95	0.98	P60Av
The probability that the sprinkler system in cold-rolling mill 1 extinguishes the fire.	0.85	0.94	0.96	P61Av
The probability that the sprinkler system in cold-rolling mill 2 extinguishes the fire.	0.9	0.96	0.98	P62Av
The probability that the CO2 system in cold-rolling mill 2 extinguishes the fire.	0.7	0.95	0.96	P63Av

Probability estimates pertaining to Area 4

Area 4 is the most important area in the building from the standpoint of production. All products that are produced in the cold-rolling mill go through production line 60, which is located in that area. Also, approximately 50% of the products go through production line 55, which is also located there. The area is filled with machines (see Table 8) and combustibles in the form of oil, rubber and plastics.

Table 32 *Probability estimates concerning fire development in area 4.*

	Min	Most likely	Max	Abbreviation
The probability of a fire in a machine, given that a fire has occurred	0.65	0.75	0.9	P64Av
The probability of a fire in the uncoiling capstan, weld, etc. (line 55), given that a fire has occurred in a machine.	0.05	0.1	0.15	P65Av
The probability of a fire in the cold-rolling mill, given that a fire has occurred in a machine.	0.3	0.4	0.6	P66Av
The probability of a fire in the oven (line 55), given that a fire has occurred in a machine.	0.05	0.2	0.3	P67Av
The probability of a fire in the uncoiling capstan, weld, etc. (line 60), given that a fire has occurred in a machine.	0.05	0.1	0.15	P68Av
The probability of a fire in the oven (line 55), given that a fire has occurred in a machine.	0.05	0.2	0.3	P69Av
<i>Line 55</i>				
<i>Fire potential in the weld</i>				
Small fire	0.77	0.87	0.94	P70Av
Medium-sized fire	0.05	0.11	0.2	P71Av
Large fire	0.001	0.02	0.03	P72Av
<i>Fire potential in the cold-rolling mill</i>				
Small fire	0.54	0.67	0.83	P73Av
Medium-sized fire	0.15	0.29	0.4	P74Av
Large fire	0.02	0.04	0.06	P75Av
<i>Fire potential in the oven</i>				
Small fire	0.8	0.88	0.94	P76Av
Medium-sized fire	0.05	0.1	0.17	P77Av
Large fire	0.001	0.02	0.03	P78Av

Table 32 (continued)

<i>Line 60</i>				
<i>Fire potential in the welding machine</i>				
Small fire	0.77	0.87	0.94	P79Av
Medium-sized fire	0.05	0.11	0.2	P80Av
Large fire	0.001	0.02	0.03	P81Av
<i>Fire potential in the oven</i>				
Small fire	0.77	0.88	0.95	P82Av
Medium-sized fire	0.05	0.1	0.2	P83Av
Large fire	0.001	0.02	0.03	P84Av
<i>Fire potential of other fires</i>				
Small fire	0.77	0.82	0.89	P85Av
Medium-sized fire	0.1	0.16	0.2	P86Av
Large fire	0.001	0.02	0.03	P87Av
<i>The probability that the employees succeed in extinguishing a fire in:</i>				
The welding machine (line 55)	0.7	0.85	0.9	P88Av
The cold-rolling mills	0.6	0.7	0.8	P89Av
The oven (Line 55)	0.7	0.8	0.9	P90Av
The weld (Line 60)	0.7	0.85	0.9	P91Av
The oven (Line 60)	0.7	0.8	0.9	P92Av
Other fires	0.7	0.8	0.9	P93Av
The probability of the fire department succeeding in extinguishing a large fire.	0.8	0.95	0.97	P94Av
The probability of the light water system in the cold-rolling mill extinguishing the fire.	0.7	0.94	0.96	P95Av

Probability estimates pertaining to Area 5

Area 5 contains the abrasive-belt grinder, which is used to process the surface of the steel. An oil room in which the oil used in the abrasive-belt grinder is filtered and stored is located near to it. The grinder has caused serious fires in the past.

Table 33 *Probability estimates concerning the development of a fire in area 5.*

	Min	Most likely	Max	Abbreviation
The probability of a fire in a machine, given that a fire has occurred.	0.75	0.97	0.98	P96Av
The probability of a fire in the abrasive-belt grinder, given that a fire in a machine has occurred.	0.4	0.5	0.8	P97Av
The probability of a fire in the oil-room, given that a fire in a machine has occurred.	0.2	0.5	0.6	P98Av
<i>Fire potential in the abrasive-belt grinder</i>				
Small fire	0.7	0.8	0.85	P99Av
Medium-sized fire	0.1	0.15	0.25	P100Av
Large fire	0.03	0.05	0.1	P101Av
<i>Fire potential in the oil room</i>				
Small fire	0.74	0.8	0.88	P102Av
Medium-sized fire	0.1	0.16	0.2	P103Av
Large fire	0.02	0.04	0.06	P104Av
<i>Fire potential of other fires</i>				
Small fire	0.7	0.82	0.9	P105Av
Medium-sized fire	0.09	0.16	0.26	P106Av
Large fire	0.01	0.02	0.04	P107Av
<i>The probability of the employees succeeding in extinguishing a fire</i>				
In the abrasive-belt grinder	0.5	0.75	0.8	P108Av
In the oil room	0.5	0.7	0.8	P109Av
A fire in another location	0.5	0.8	0.9	P110Av
The probability of the fire department succeeding in extinguishing a large fire.	0.7	0.95	0.96	P111Av
The probability of the fire department succeeding in extinguishing a large fire in the abrasive-belt grinder.	0.6	0.7	0.9	P112Av
The probability of the fire department succeeding in extinguishing a large fire in the oil-room.	0.65	0.7	0.95	P113Av
The probability of the CO2 system in the oil room succeeding in extinguishing the fire.	0.8	0.95	0.96	P114Av

Probability estimates pertaining to other areas

The other areas in the building are less important from a production standpoint or are less difficult to model a fire in than the areas taken up above. These other areas are thus dealt with in terms of one large event tree (see Appendix E).

Table 34 *Probability estimates pertaining to the development of fire in the other areas.*

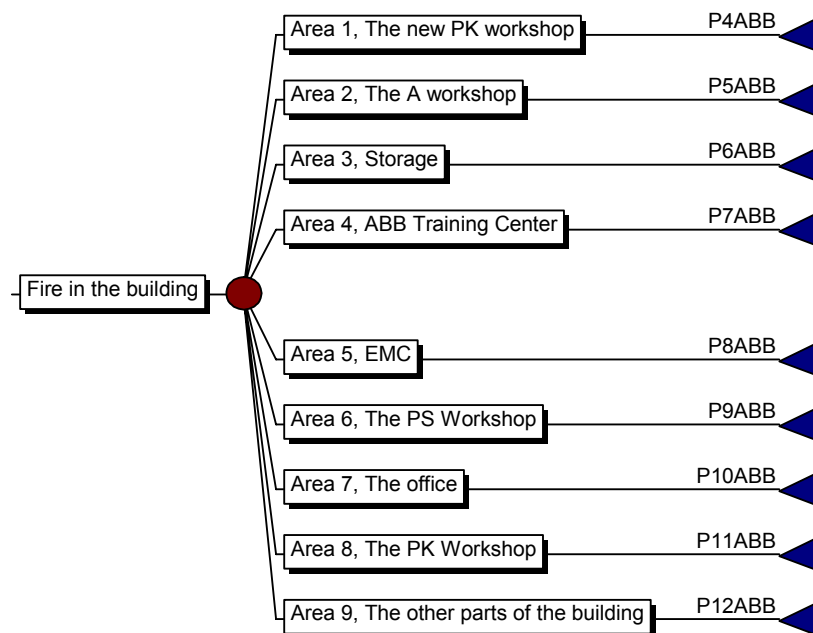
	Min	Most likely	Max	Abbreviation
The probability of a fire in engine room 1	0.2	0.3	0.4	P115Av
The probability of a fire in engine room 2	0.2	0.3	0.4	P116Av
The probability of a fire in machine shop 1	0.05	0.15	0.2	P117Av
The probability of a fire in machine shop 2	0.05	0.15	0.2	P118Av
The probability of a fire in the oil room	0.01	0.05	0.1	P119Av
The probability of a fire in the pallet storage area	0.01	0.05	0.1	P120Av
<i>Fire potential in the engine rooms</i>				
Small fire	0.73	0.8	0.85	P121Av
Medium-sized fire	0.14	0.19	0.25	P122Av
Large fire	0.001	0.01	0.02	P123Av
<i>Fire potential in the machine shops</i>				
Small fire	0.8	0.87	0.9	P124Av
Medium-sized fire	0.05	0.11	0.17	P125Av
Large fire	0.001	0.02	0.03	P126Av
<i>Fire potential in the oil room</i>				
Small fire	0.74	0.8	0.88	P127Av
Medium-sized fire	0.1	0.16	0.2	P128Av
Large fire	0.02	0.04	0.06	P129Av
<i>Fire potential in the pallet storage space</i>				
Small fire	0.6	0.7	0.8	P130Av
Medium-sized fire	0.1	0.15	0.2	P131Av
Large fire	0.05	0.15	0.2	P132Av
<i>The probability that the employees succeed in extinguishing a fire in:</i>				
The engine rooms, given that the alarm works	0.6	0.85	0.9	P133Av
The engine rooms, given that the alarm fails to work	0.4	0.6	0.7	P134Av
In the machine shops	0.6	0.85	0.9	P135Av

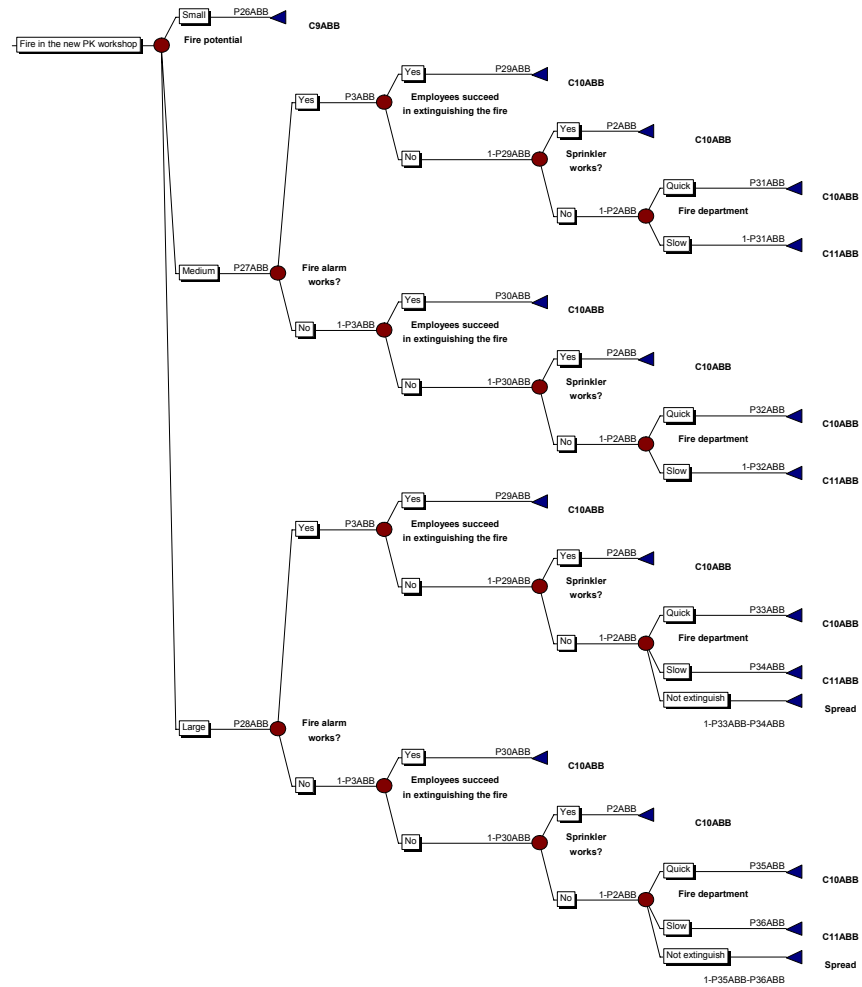
Table 34 (continued)

<i>Fire department</i>				
The probability that the fire department succeeds in extinguishing a fire in an engine room, given that the alarm works.	0.7	0.95	0.97	P136Av
The probability that the fire department succeeds in extinguishing a fire in an engine room, given that the alarm fails to work.	0.5	0.85	0.9	P137Av
The probability that the fire department succeeds in extinguishing a fire in a machine shop.	0.5	0.95	0.97	P138Av
The probability that the fire department succeeds in extinguishing a fire in the oil room.	0.5	0.85	0.9	P139Av
The probability that the fire department succeeds in extinguishing a fire in the pallet storage space.	0.5	0.85	0.9	P140Av
The probability that a fire does not spread from the pallet storage space.	0.6	0.8	0.9	P141Av
The probability that a fire does not spread from the oil room.	0.6	0.96	0.98	P142Av
The probability that the smoke detectors in the motor room detect the fire.	0.8	0.96	0.97	P143Av

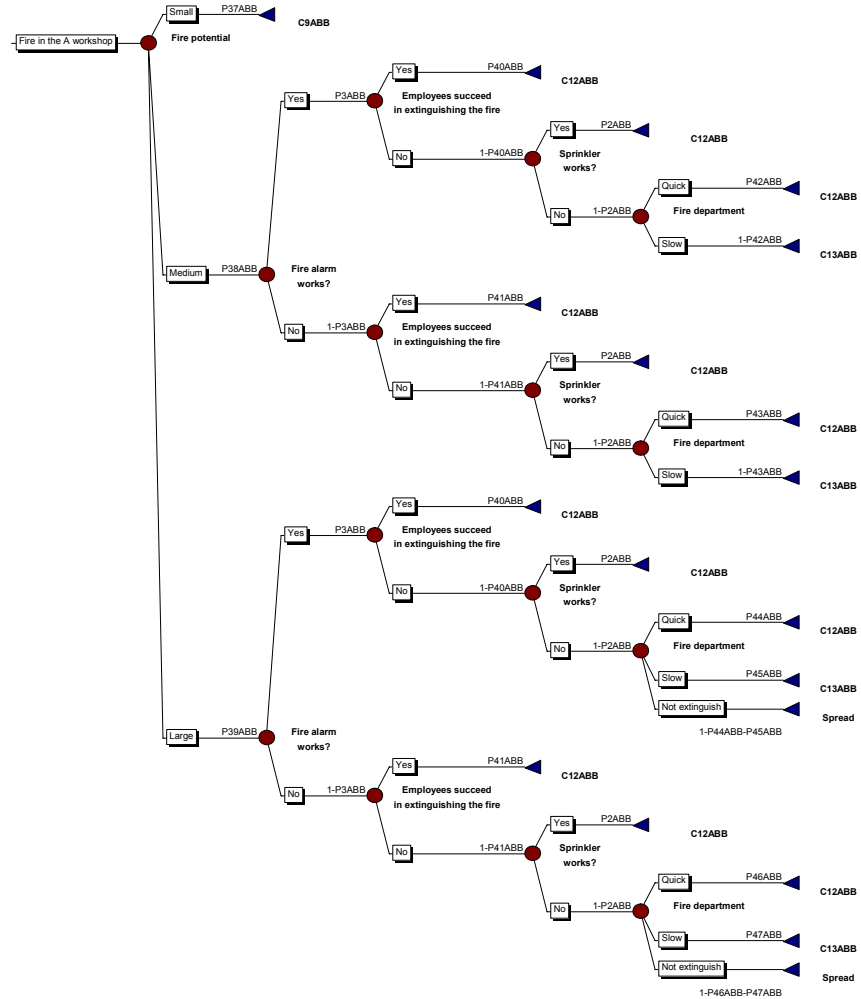
Appendix D: Event trees pertaining to the ABB analysis

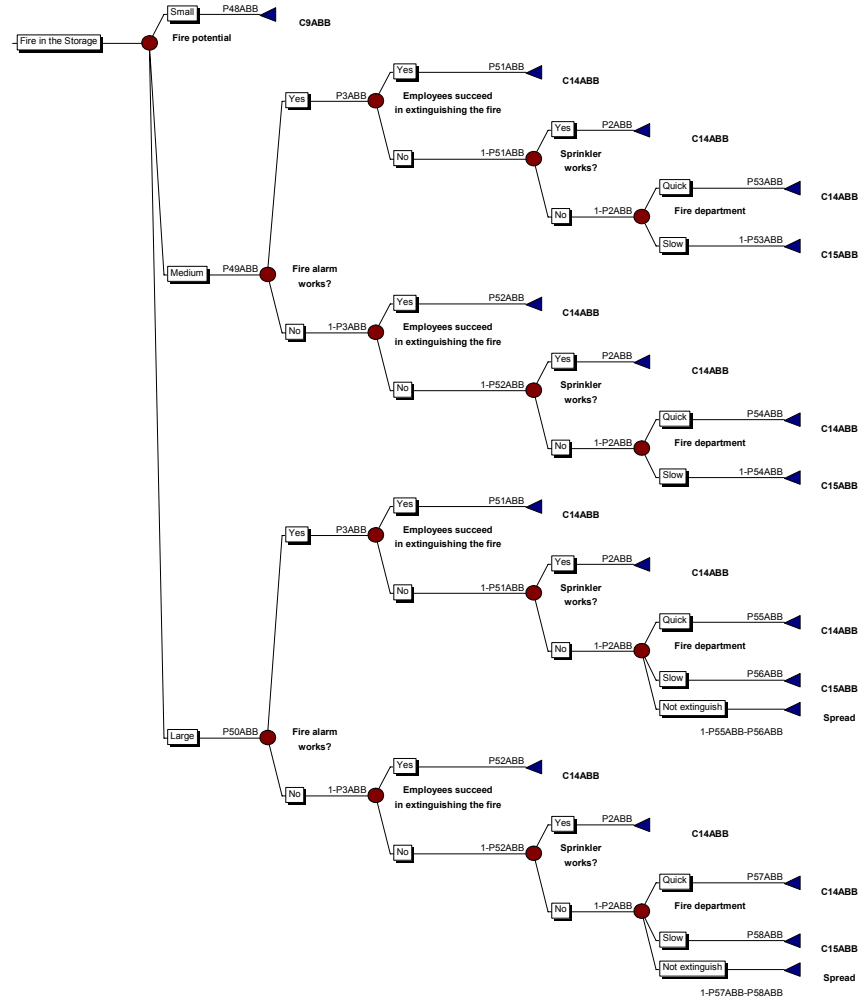
The event tree model used in the ABB-analysis is presented here. Since the tree is very large, it is divided up into several parts. The first part is used to determine the area in which a fire has occurred, each of the parts that follow concerning the development of a fire in a specific area of the building. The estimates of the different probability values can be found in Appendix B.



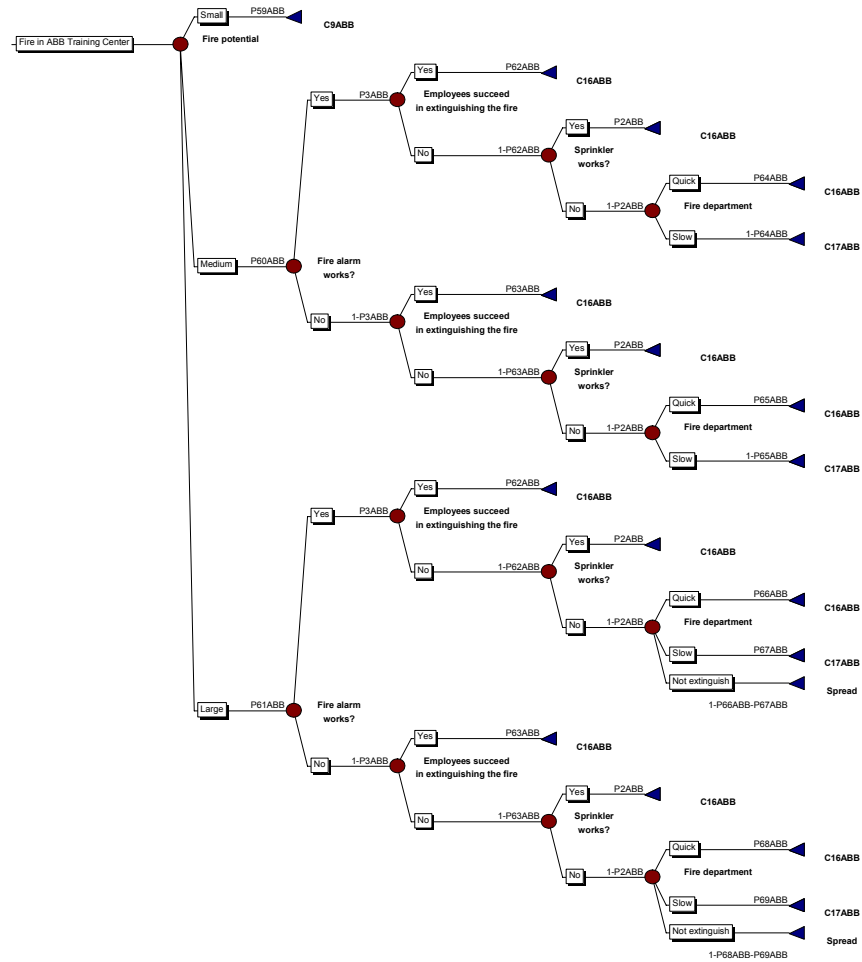



Appendix D: Event trees pertaining to the ABB analysis






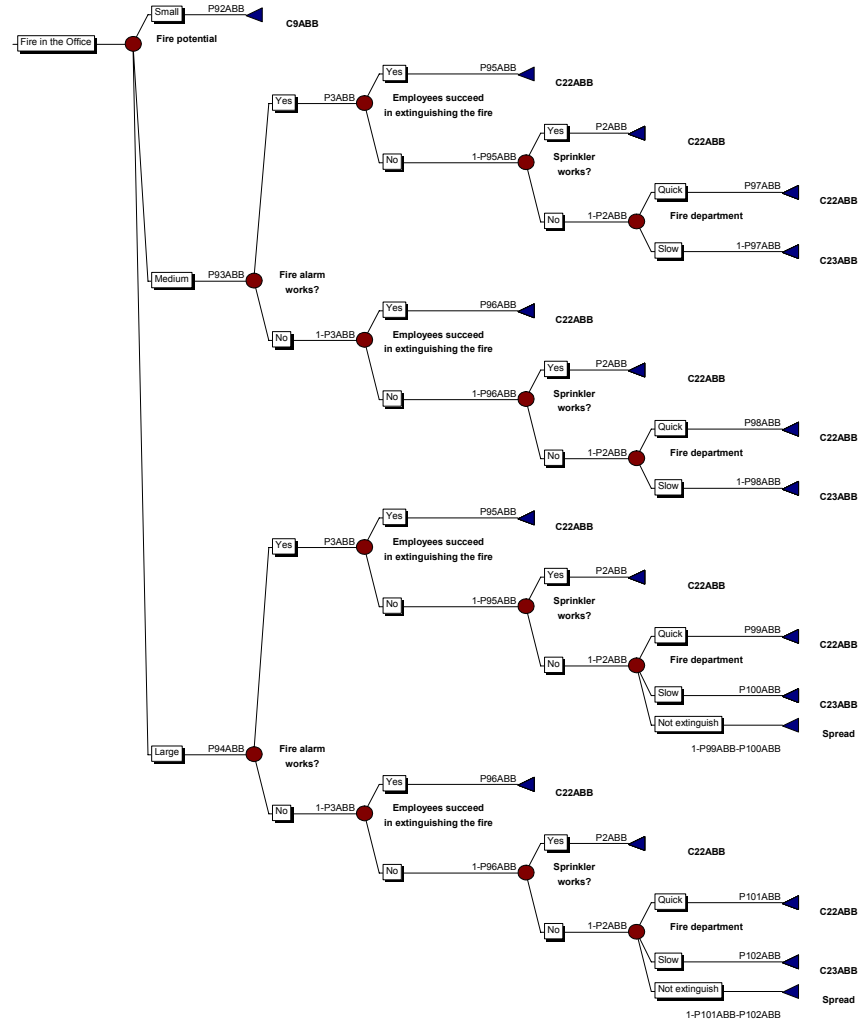
Appendix D: Event trees pertaining to the ABB analysis



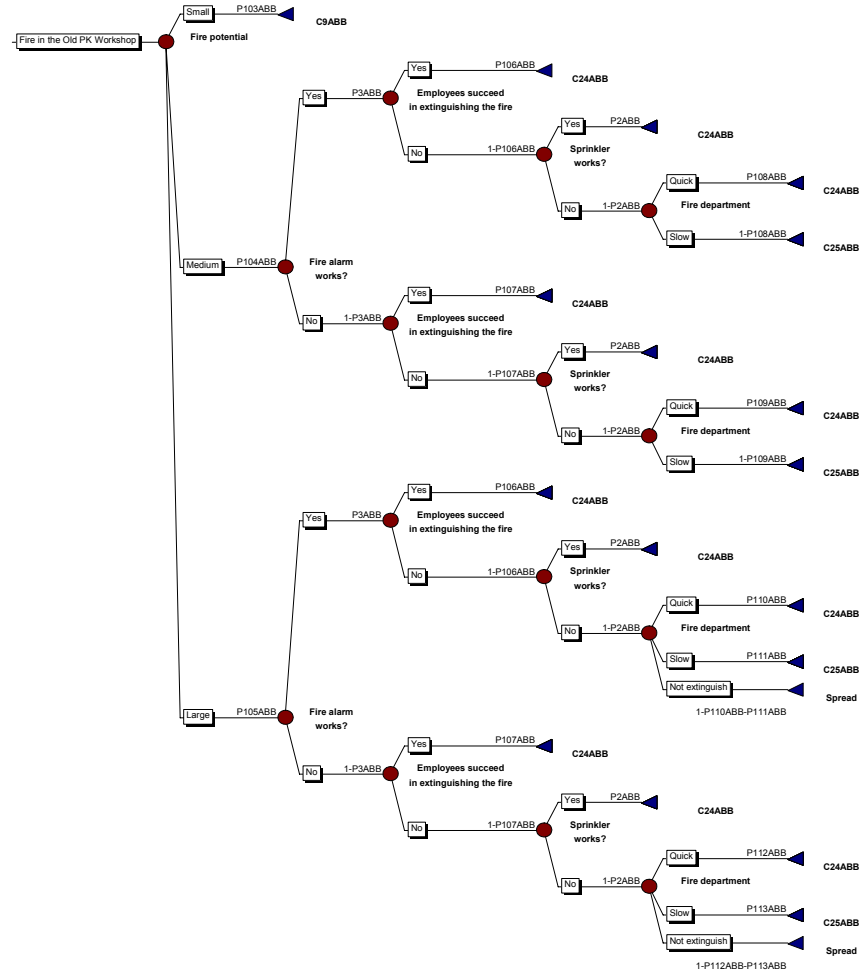


Appendix D: Event trees pertaining to the ABB analysis







Appendix D: Event trees pertaining to the ABB analysis



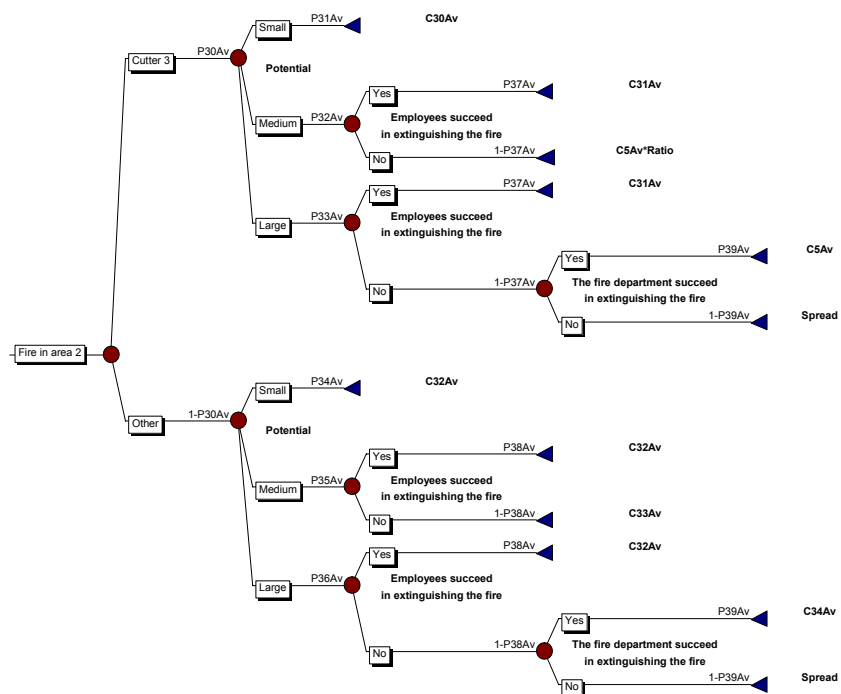
Appendix E: Event trees pertaining to the Avesta Sheffield analysis

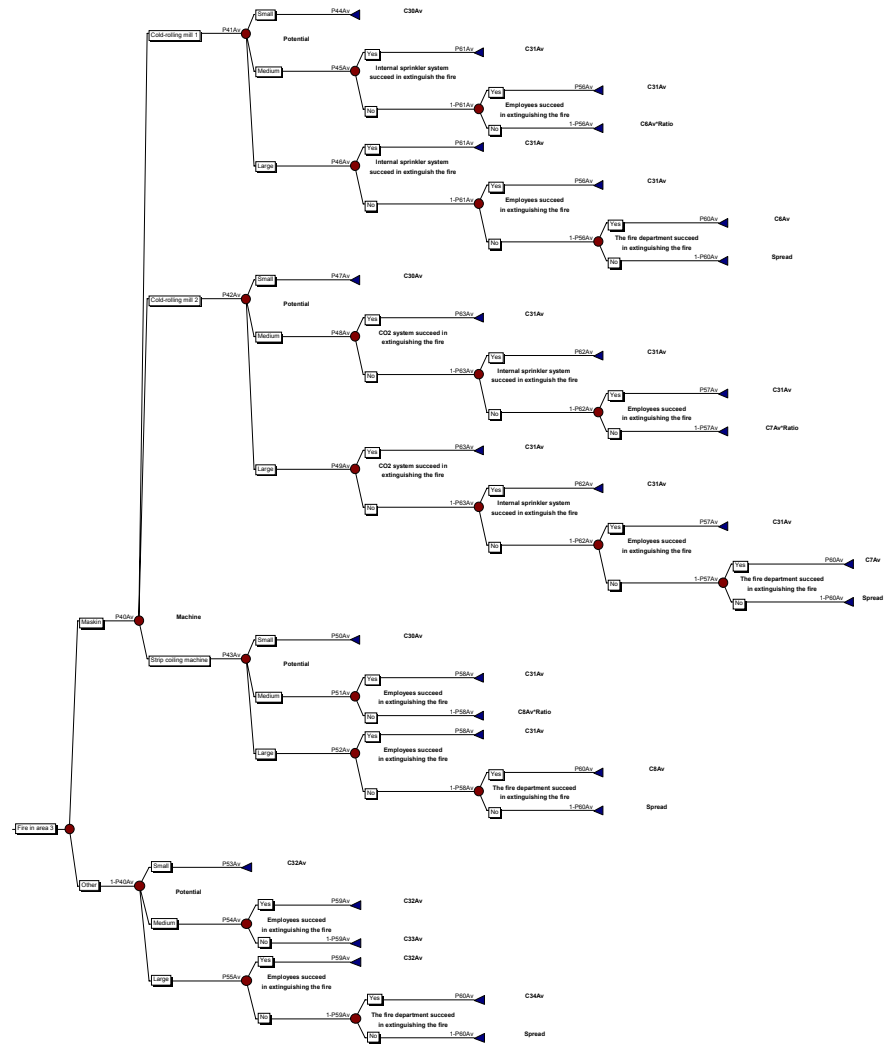
The event tree model used in the Avesta Sheffield analysis is presented here. Since the tree is very large, it is divided up into several parts. The first part is used to determine in which area a fire has occurred. Each of the parts that follow pertain to the development of a fire in a specific area of the building. The estimates of the different probabilities can be found in Appendix C.






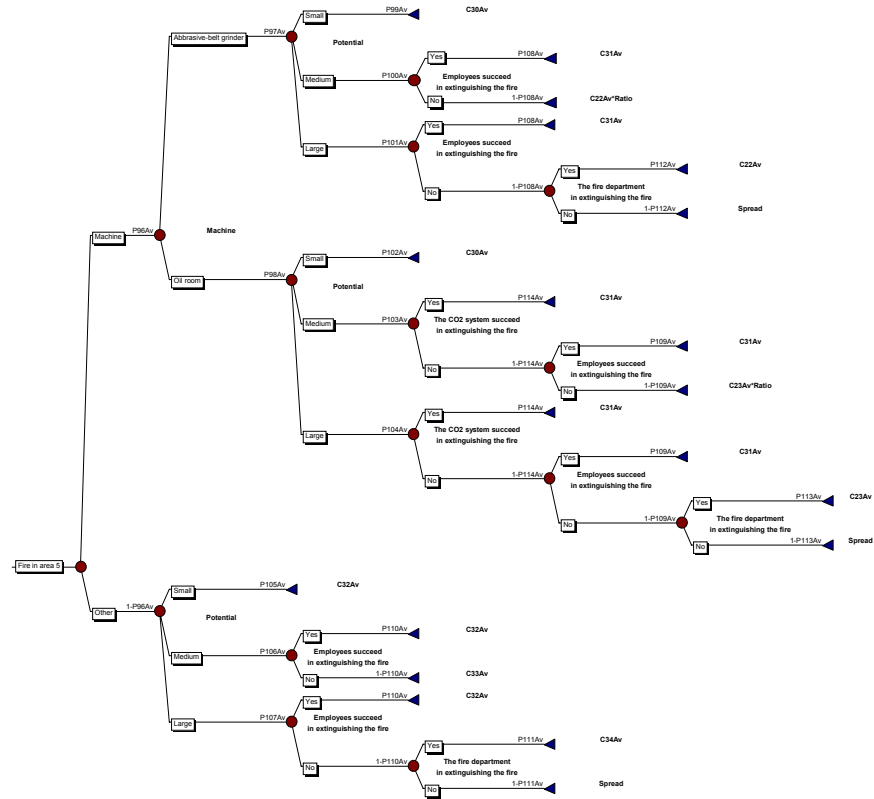
Appendix E: Event trees pertaining to the Avesta Sheffield analysis



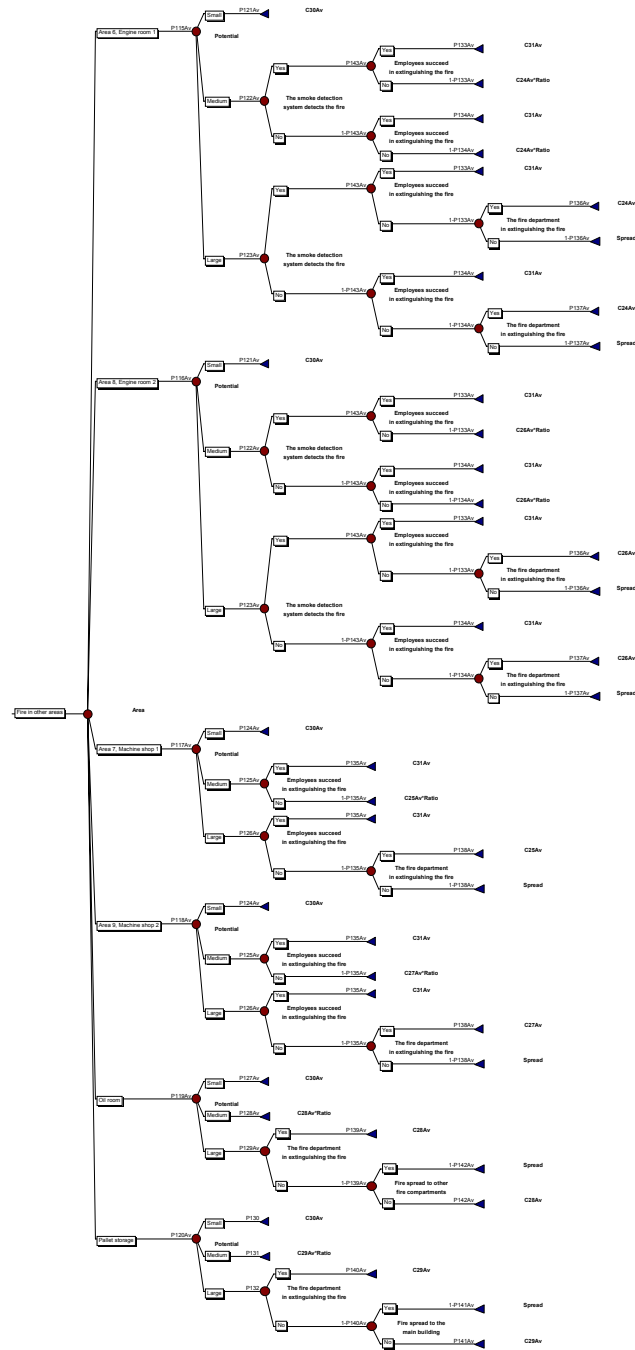


Appendix E: Event trees pertaining to the Avesta Sheffield analysis





Appendix E: Event trees pertaining to the Avesta Sheffield analysis



Appendix F: Computer codes

A computer code written in MATLAB (Version 6.1.0.450 Release 12.1) was used to calculate the expected loss, given that a fire has spread so as to involve a specific area of the building.

It is assumed that the building can be divided into areas that are separated by barriers possibly able to stop the spread of a fire from one area to another. Each of the barriers is assigned a probability value representing the probability that the barrier in question would succeed in limiting the further spread of a fire.

In the ABB-case, one needs to specify the uncompensated losses associated with the destruction of a particular area. In the Avesta Sheffield-case, one needs instead to specify the direct losses that the destruction of a particular area would involve, the consequential losses per month and the estimated time until production can be started again, given that a particular area is destroyed.

The computer program consists of a number of files, which are included here. The files, which are executed in MATLAB, are *Spreadabb.m* and *Spreadavesta.m*. These use other files for calculating the expected uncompensated losses given that a fire has spread to involve a particular area of the building.

A brief description of the different files follows.

Spreadabb.m

Spreadabb.m is the file that is executed in analysing the ABB-case. In that file the user has to create two matrixes, one containing the probabilities pertaining to the different barriers in the building (prob) and the other containing the losses (cost) associated with the destruction of a particular area.

Controlabb.m

This file provides the simulation of fires in different areas of the building, first area 1, then area 2, and so on. It uses the files *Spread.m* and *Expectedabb.m* to perform the calculations.

Spread.m

This file calculates the probabilities of different fire scenarios, i.e. different combinations of effective and non-effective barriers. It also calculates which areas are destroyed in the case of a particular fire scenario. This information is delivered to *Controlabb.m*.

Expectedabb.m

The file *Expectedabb.m* calculates the expected loss, given the probabilities of the different fire scenarios, and provides information on the areas destroyed in each fire scenario and the uncompensated losses associated with the destruction of each area.

Spreadavesta.m

Spreadavesta.m is the file that is executed in analysing fires in the Avesta Sheffield-building. It has the same basic structure as *Spreadabb.m*, except that it provides uncompensated losses in two parts. The one part concerns each of the different areas and the other the business interruption that would follow a serious fire. It also provides estimates of how quickly production can be resumed after a serious fire.

Controlavesta.m

This file controls the simulation of fire in each of the areas, first a serious fire in area 1, then a fire in area 2, and so on.

Expectedavesta.m

Expectedavesta.m is used to calculate expected loss. It differs somewhat from *Expectedabb.m* in that the losses are more difficult to calculate than in the ABB case.

Spreadabb.m

```
clear

%Note that a probability value needs only be given once per
barrier.

%"prob" contains the probability that a barrier will stop a
fire from spreading further.

prob=[0 0.6 0 0 0 0 0 0.9;0 0 0.5 0.5 0 0.9 0.9 0.9;0 0 0
0.5 0.5 0.9 0 0;
0 0 0 0 0.5 0 0 0;0 0 0 0 0 0 0 0;0 0 0 0 0 0 0.5 0;0 0 0 0
0 0 0 0.5;0 0 0 0 0 0 0 0]

%"cost" contains the losses associated with the different
areas. Note that the losses can be divided into three parts
if necessary.

cost=[160000 160000 0;
      120000 190000 0;
      5000 0 0;
      13000 0 0;
      18000 0 0;
      150000 82500 0;
      40000 0 0;
      250000 250000 0];

result=Controlabb(prob,cost);

save C:\Matlabresultat\grund.txt result -ascii
```

Controlabb.m

```
function Controlabb=Controlabb(prob,cost)

%This function simulates possible fire scenarios, given
that a fire has started in a particular area. The result is
a "Control" matrix that contains the expected loss in each
of the areas given that a fire has started there.

clear res

[probres,destroyed]=Spread(1,prob);
costprob=Expectedabb(cost,probres,destroyed);
resegendom=costprob(:,1).*costprob(:,4);
resavb=costprob(:,2).*costprob(:,4);
res(1,1)=sum(resegendom);
res(1,2)=sum(resavb);
save C:\Matlabresultat\costprob1.txt costprob -ascii

%The above code is repeated for all areas of the building,
the results being saved in the matrix "Control".

Control=res;
```

Expectedabb.m

```
function costprob=Expectedabb(cost,probres,destroyed)

info=destroyed;
x=length(destroyed(1,:));
y=length(destroyed(:,1));
for n=1:y
    info(n,x+1)=probres(n);
end

for n=1:y                %Direct losses
    sum=0;
    for i=1:x
        if info(n,i)==1
            sum=sum+cost(i,1);
        end
    end
    costprob(n,1)=sum;
end

for n=1:y                %Consequential losses
    sum=0;
    for i=1:x
        if info(n,i)==1
            sum=sum+cost(i,2);
        end
    end
    costprob(n,2)=sum;
end

for n=1:y                %Hidden losses
    sum=0;
    for i=1:x
        if info(n,i)==1
            sum=sum+cost(i,3);
        end
    end
    costprob(n,3)=sum;
end

sum1=0;
sum2=0;
sum3=0;
for n=1:y                %Expected loss
    sum1=sum1+(costprob(n,1)*probres(n));
    sum2=sum2+(costprob(n,2)*probres(n));
    sum3=sum3+(costprob(n,3)*probres(n));
end
format bank
```

```
sum1
sum2
sum3

%Riskprofile
costprob(:,4)=probres;
costprob=sortrows(costprob,[1]);
for n=1:y
    i=y-n+1;
    if n==1
        costprob(i,5)=costprob(i,4);
    else
        costprob(i,5)=costprob(i,4)+costprob(i+1,5);
    end
end
hold on
plot(costprob(:,1),costprob(:,5),'b')

costprob=sortrows(costprob,[2]);
for n=1:y
    i=y-n+1;
    if n==1
        costprob(i,5)=costprob(i,4);
    else
        costprob(i,5)=costprob(i,4)+costprob(i+1,5);
    end
end
plot(costprob(:,2),costprob(:,5),'g')

costprob=sortrows(costprob,[3]);
for n=1:y
    i=y-n+1;
    if n==1
        costprob(i,5)=costprob(i,4);
    else
        costprob(i,5)=costprob(i,4)+costprob(i+1,5);
    end
end
plot(costprob(:,3),costprob(:,5),'r')
grid
```

Spread.m

```
function [probres,destroyed]=Spread(firepos,prob)

%This function provides the resulting probabilities of the
different fire scenarios as well as a matrix showing which
areas are destroyed in terms of each of the fire scenarios.

clear probres
clear destroyed
clear barrier
clear n
clear numbrcell

[rader,kolumner]=size(prob);

barpos=1;%Pointer in the barrier matrix
numbrcell=size(prob,2);%numbrcell is the number of areas in
the building
for y=1:kolumner %Creates the barrier matrix that
indicates between which areas each of the barriers is
located and the probability of stopping a fire there.
    for x=1:rader
        if prob(x,y)>0
            barrier(barpos,1)=y;
            barrier(barpos,2)=x;
            barrier(barpos,3)=prob(x,y);
            barpos=barpos+1;
        end
    end
end
n=0;
n=size(barrier,1); %n is the number of barriers
works=0;
for x=1:n %Creates the "works"-matrix that
consists of 1 and 0:s indicating which barriers are working
and which are non-working.
    y=0;
    u=1;
    number=2^n;
    while y<(number)
        for y=y+1:(y+(2^(x-1)))
            if u==1
                works(y,x)=1;
            else
                works(y,x)=0;
            end
            u;
        end
        if u==1
```



```
        u=0;
    else
        u=1;
    end
end
end

prob1=barrier(:,3);
prob2=1-prob1;
for y=1:number %Creates the probability-matrix,
where "number" is the number of rows
    for x=1:n %n is the number of barriers
        if works(y,x)==1
            prob3(y,x)=prob1(x);
        else
            prob3(y,x)=prob2(x);
        end
    end
end
%Creates the resulting scenario probabilities
probres=prod(prob3,2);

destroyed=zeros(number,numbrcell); %The "destroyed" matrix
shows which areas are destroyed ("1") and which are not
("0") for all scenarios
pos=1;

for y=1:number
    possible=zeros(1,numbrcell); %"possible"
    indicates the areas that are destroyed, given that a fire
    occurs in a specific area, their being shown in the column
    for x=1:n
        if works(y,x)==0

            k=length(possible(:,barrier(x,1)));
            for i=1:k
                pos=k-i+1;
                possible(pos+1,barrier(x,1))=
=possible(pos,barrier(x,1));
            end

            possible(1,barrier(x,1))=barrier(x,2);

            k=length(possible(:,barrier(x,2)));
            for i=1:k
                pos=k-i+1;
                possible(pos+1,barrier(x,2))=
=possible(pos,barrier(x,2));
            end
        end
    end
end
```

end

Spreadavesta.m

```
clf

prob=[0 0 0.4 0.5 0 0 0 0 0 0;0 0 0.7 0 0 0 0 0 0 0;0 0 0 0.5
0.5 0.8 0.8 0.8 0.8;0 0 0 0 0.8 0 0 0 0 0;0 0 0 0 0 0 0 0 0 0;
0 0 0 0 0 0 0 0 0 0;0 0 0 0 0 0 0 0 0 0;0 0 0 0 0 0 0 0 0 0;0 0 0
0 0 0 0 0 0 0] %The probability that the different barriers
will stop a fire
directloss=[245260 91932 522744 1115740 151400 70200 30400
100200 30400]; %Direct loss when one area is destroyed
losspermonth=90000 %The loss per month for the whole
factory
interruption=[14 14 14 14 14 14 0 14 0]; %Time of
interruption when one area is destroyed
[res,riskprof]=controlavesta(prob,directloss,interruption,
losspermonth);
save C:\Matlabresultat\grund.txt res -ascii;
save C:\Matlabresultat\riskprofgrund.txt riskprof -ascii;
```

Controlavesta.m

```
function [res,riskprof]=  
=controlavesta(spread,directloss,interruption,losspermonth)  
  
clear res  
  
%The following rows are repeated for all areas in the  
building.  
  
[probres,destroyed]=spread(1,spread);  
costs=expectedavesta(directloss,interruption,losspermonth,  
probres,destroyed);  
resdirect=costs(:,1).*costs(:,4);  
resind=costs(:,2).*costs(:,4);  
res(1,1)=sum(resdirect);  
res(1,2)=sum(resind);  
riskprof(:,1)=costs(:,4)  
riskprof(:,2)=costs(:,1)  
riskprof(:,3)=costs(:,2)
```

Expectedavesta.m

```
Function costs=
=expectedavesta(directloss,interruption,losspermonth,
probres,destroyed)

info=destroyed;
x=length(destroyed(1,:));
y=length(destroyed(:,1));
for n=1:y
    info(n,x+1)=probres(n);
end

for n=1:y                %Direct losses
    summa=0;
    for i=1:x
        if info(n,i)==1
            summa=summa+directloss(i);
        end
    end
    costs(n,1)=summa;
end

prodminus=[0 0.5 0 0.5 0 0 0 0 0;
0.5 0 0 0.5 0 0 0 0 0;
0 0.5 1 0.5 0 0 0 0.8 0;
0.5 0 1 0.5 0 0 0 0.8 0;
0 0.5 1 0.5 0 0.2 0 0 0 ;
0.5 0 1 0.5 0 0.2 0 0 0 ]
produktion=[1 0.5 1 0.5 1 1 1 1 1;
0.5 1 1 0.5 1 1 1 1 1;
1 0.5 1 0.5 1 1 1 0.8 1;
0.5 1 1 0.5 1 1 1 0.8 1;
1 0.5 1 0.5 1 0.2 1 1 1 ;
0.5 1 1 0.5 1 0.2 1 1 1]
clear interruptionlength
for n=1:y                %Consequential losses
    summa=0;
    produktion1=produktion
    for i=1:x
        if info(n,i)==1
            produktion1(:,i)=produktion(:,i)-prodminus(:,i)
            interruptionlength(i)=1

        else
            interruptionlength(i)=0
        end
    end
    for i=1:x
        if interruptionlength(i)==1
```

```

        avbrottet(i)=interruption(i);
    else
        avbrottet(i)=0;
    end
end

res=prod(produktion1,2);
faktor=sum(res);
costs(n,2)=losspermonth*(1-faktor)*max(avbrottet);
end

sumdirect=0;
sumconsequential=0;

for n=1:y                %Expected loss
    sumdirect=sumdirect+(costs(n,1)*probres(n));

sumconsequential=sumconsequential+(costs(n,2)*probres(n));

end
format bank
sumdirect
sumconsequential

%Riskprofile
costs(:,4)=probres;
costs=sortrows(costs,[1]);
for n=1:y
    i=y-n+1;
    if n==1
        costs(i,5)=costs(i,4);
    else
        costs(i,5)=costs(i,4)+costs(i+1,5);
    end
end
end
hold on
plot(costs(:,1),costs(:,5),'b')

costs=sortrows(costs,[2]);
for n=1:y
    i=y-n+1;
    if n==1
        costs(i,5)=costs(i,4);
    else
        costs(i,5)=costs(i,4)+costs(i+1,5);
    end
end
end
plot(costs(:,2),costs(:,5),'g')

```

```
costs=sortrows(costs,[3]);  
for n=1:y  
    i=y-n+1;  
    if n==1  
        costs(i,5)=costs(i,4);  
    else  
        costs(i,5)=costs(i,4)+costs(i+1,5);  
    end  
end  
plot(costs(:,3),costs(:,5),'r')  
grid
```