Making the business case for process safety using value-at-risk concepts

Jayming S. Fang, David M. Ford*, M. Sam Mannan

*Mary Kay O'Connor Process Safety Center, Department of Chemical Engineering, Texas
A&M University College Station, TX 77843-3122, USA*

Abstract

An increasing emphasis on chemical process safety over the last two decades has led to the development and application of powerful risk assessment tools. Hazard analysis and risk evaluation techniques have developed to the point where quantitatively meaningful risks can be calculated for processes and plants. However, the results are typically presented in semi-quantitative "ranked list" or "categorical matrix" formats, which are certainly useful but not optimal for making business decisions. A relatively new technique for performing valuation under uncertainty, value at risk (VaR), has been developed in the financial world. VaR is a method of evaluating the probability of a gain or loss by a complex venture, by examining the stochastic behavior of its components. We believe that combining quantitative risk assessment techniques with VaR concepts will bridge the gap between engineers and scientists who determine process risk and business leaders and policy makers who evaluate, manage, or regulate risk. We present a few basic examples of the application of VaR to hazard analysis in the chemical process industry.
© 2004 Published by Elsevier B.V.

Keywords: Chemical process safety; Value at risk; Damage index; Hazard index; Quantitative risk assessment

1. Introduction

1.1. Background

Due to the inherent sensitivity of the chemical process industry (CPI) to the consequences of failure, chemical process safety has been a major concern for some time [1]. Chemical process quantitative risk assessment (CPQRA) identifies areas in operations, engineering, and management systems that might be modified to reduce process risk. CPQRA deals with both aspects of risk, namely likelihood and consequence. Likelihood is typically estimated through some combination of historical data and fault/event tree analysis. Consequence modeling generally consists of two parts; detailed science models predict the parameters of incident-specific events (e.g. gas release, explosion overpressure), and effect/mitigation models predict the final consequences on people and the environment (natural and built). The prod-

uct of likelihood and consequence is a measure of risk. Presently, CPQRA has developed to the point where quantitatively meaningful risks may be calculated for individual processes and entire plants.

Obviously, implementing safety devices and procedures to remove all risks in a chemical plant is not feasible. Thus, an important part of a CPQRA analysis is prioritizing the risks for appropriate action. The results are typically reported in a likelihood-consequence matrix format, or perhaps in a ranked list. While this semi-quantitative approach is useful, we believe that CPQRA has progressed to a point where the results may be presented in more detail and with more quantitative precision. Furthermore, they should be presented in a comprehensive format that is useful to CPI management and other policy makers. This is not an easy task, primarily due to the inherently probabilistic nature of the problem. However, the rewards of such an approach would be substantial; a more quantitative and coherent business case for process safety would certainly result in a better-focused investment by the CPI.

In this paper, we present a new approach for understanding, organizing, and packaging the results of CPQRA

---

* Corresponding author. Tel.: +1-979-862-4850.
E-mail address: D-Ford@chemail.tamu.edu (D.M. Ford).

analyses. The approach is based on a technique, value at risk (VaR), borrowed from the financial industry [2]; it will provide a bridge between the engineers and scientists who *calculate* process risk and the business leaders and policy makers who *evaluate, manage, or regulate* risk in a broader context. VaR is a method of evaluating the probability of a gain or loss by a complex venture, by examining the stochastic behavior of its components. The framework is firmly grounded in the theory of VaR, yet flexible enough so that it may be:

- used at several different organizational levels (process, plant, industry segment);
- integrated with other business risk concerns (operational, market) so that complete and accurate cost-benefit decisions may be made;
- implemented in software targeted for industrial risk professionals;
- extended to other types of risk (environmental, societal) and for use by other stakeholders (governmental agencies, public interest groups).

The primary focus of this paper is to introduce the approach and demonstrate its use on case problems from the CPQRA literature.

We note that VaR concepts have begun to appear in other areas of process design research. For example, Barbaro and Bagajewicz [3] have employed VaR in developing a two-stage stochastic formulation for managing financial risk in planning under uncertainty.

1.2. Organization

Section 2 contains the theoretical development for combining VaR and CPQRA. Section 3 demonstrates the procedure on two different example problems. The first example is based on a single event tree and a simple damage valuation index, with various layers of probabilistic complexity sequentially added in. The second is closer to a real-world example, using a hazard quantification index from the literature. Section 4 contains conclusions and future directions.

2. Theoretical development

2.1. Value at risk

VaR is a method of evaluating the probability of a gain or loss by a complex financial venture, by examining the stochastic behavior of its components [2]. VaR approaches generally involve a combination of likelihood estimation and valuation: how likely is an event to happen, and what is the financial impact on the portfolio? Quantification of both of these aspects may involve sophisticated probabilistic analyses. A major strength of the VaR technique is that it provides a total cost-benefit analysis of an entire portfolio in terms of a single probability distribution function for value. VaR itself is technically defined as the worst loss that is expected in a portfolio, within a given confidence interval, over a specified time period.
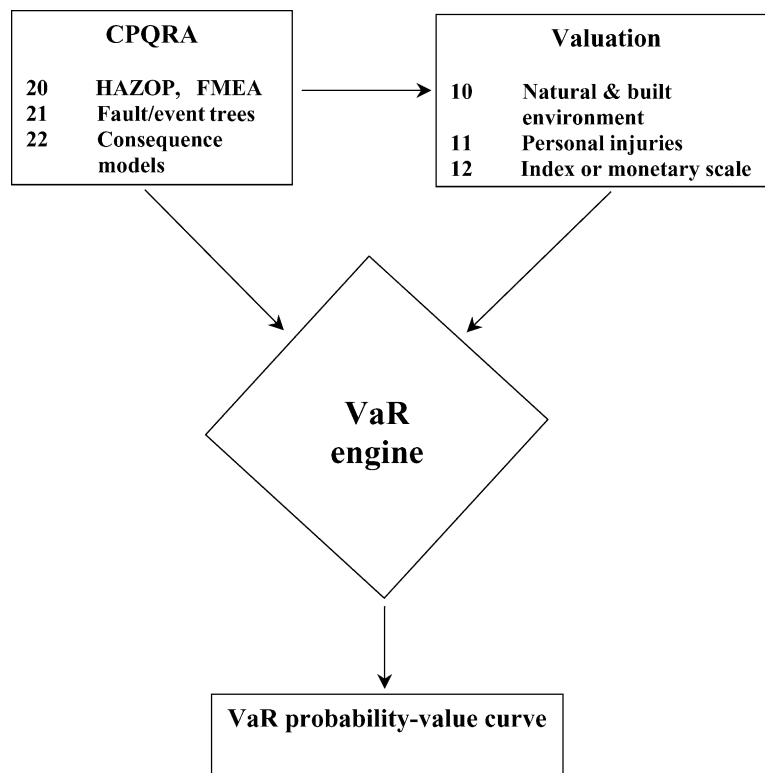

Fig. 1. Flow chart of the integration of VaR and CPQRA.

The flexibility of the VaR approach (i.e., the ability to accept input from different events), combined with the comprehensive, straightforward presentation of results (i.e., the use of a single probabilistic value function), makes it attractive for application to problems in CPQRA.

2.2. Integration of CPQRA and VaR

The diagram in Fig. 1 shows how we envision the procedure. Traditional CPQRA tools are used to determine the probabilities and consequences of undesired events associated with a plant or process. The consequences are passed to a valuation model, where they are assigned values (or distributions of values). The valuation may be done in monetary terms, or with a customized index appropriate to the particular situation or stakeholders. For undesired events, the values will typically be negative by convention. The results of the CPQRA and valuation are sent to the VaR engine, where they are combined to generate a single VaR probability distribution function representing process/plant value.

The VaR approach is capable of handling complex situations in which the fundamental stochastic events are related in a nonlinear fashion within the portfolio; this level of complexity typically requires simulation using Monte Carlo techniques [2]. This level of treatment is not required for the simple example situations described below, but it might be for many real-world problems in the CPI.

We also note that the cumulative versions of our VaR probability curves are somewhat analogous to the frequency-number, or F–N, curves often used to describe societal risk [1]. F–N curves show the cumulative frequency of undesired events with respect to the number of individuals affected

(e.g. killed, injured, exposed). Our cumulative VaR curves represent the cumulative frequency of experiencing a loss with respect to the damage value. In this paper, we consider damage value in an abstract sense and do not relate it to human life.

3. Application examples

3.1. First example problem: leak from LPG storage tank

This example problem applies a VaR analysis to a problem illustrated in chapter 3 of ref. [1]. The possible events and outcomes, and their frequencies, are taken directly from that example. We created the damage index described below, specifically for illustrative purposes related to this example. The values of the damage index for the different possible outcomes were assigned based on our judgment.

3.1.1. Scenario description
In this example, we assume that a fault tree analysis has identified the potential problem of a large leakage from an isolated LPG storage tank and estimated the frequency with which this problem is expected to occur. A further event tree analysis, as shown in Fig. 2, yields 10 possible scenarios comprised of six distinct outcomes. The six outcomes and their associated frequencies are summarized in Table 1 (note that the frequency for UVCE has been reduced to bring the total frequency to 100×10^{-6}, for convenience).

Detailed descriptions of the possible outcomes may be found in ref. [1], but we briefly outline them here. A boiling liquid expanding vapor explosion (BLEVE) occurs when a
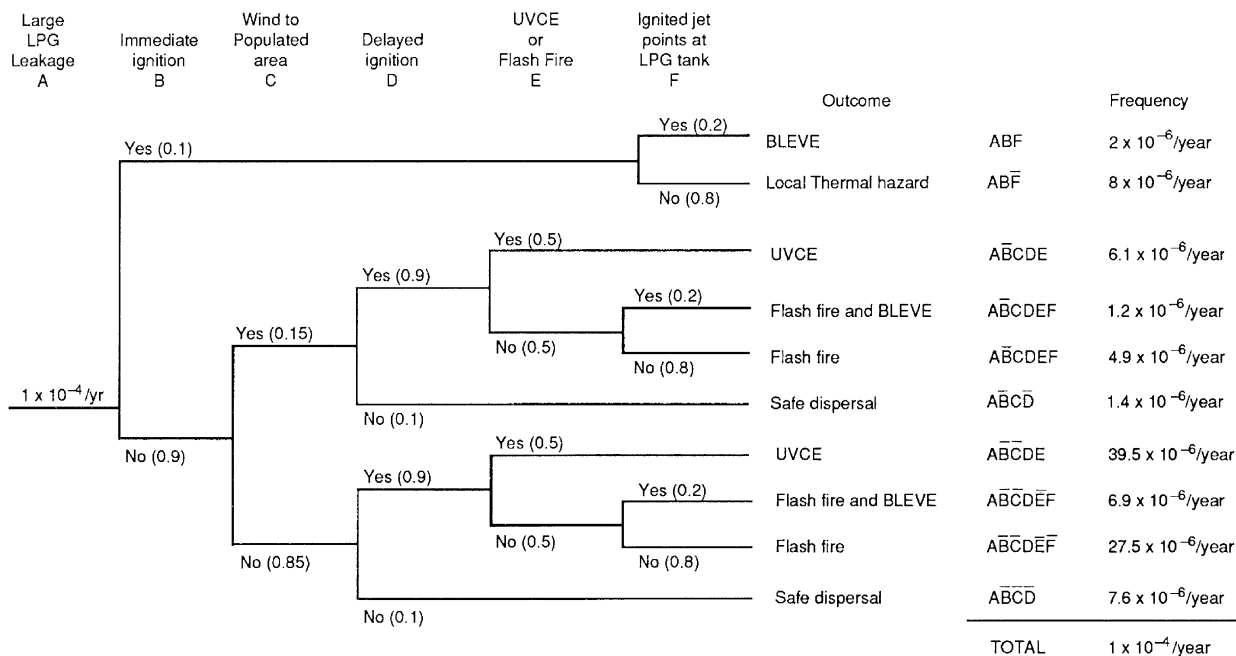
Fig. 2. Event tree for LPG leak problem.

Table 1
Data for the LPG leak problem

Incident	Damage index	Uncertainty	Frequency (10^{-6} per year)
BLEVE	25	7.5	2
Flash fire	20	7.5	32.4
Flash fire and bleve	30	10	8.1
UVCE	40	10	40.5
Local thermal hazard	10	5	8
Safe dispersal	3	1	9

pressurized vessel suddenly fails and its contents flash to the atmosphere, producing a pressure wave. If the expanding substance is also flammable, there is the additional danger of a flash fire. An unconfined vapor cloud explosion (UVCE) occurs when a drifting cloud of flammable vapor ignites and explodes, producing a shock wave. Such a cloud may also ignite but not produce an overpressure wave, thus generating a flash fire. A local thermal hazard will occur if the release burns locally, without flashing back into the tank to cause an explosion. Of course, safe dispersal is the most desirable of these undesirable events, but even this outcome has a negative value associated with a shutdown of the facility.

The event tree supplies the possible outcomes and frequencies. In order to apply the VaR analysis, we also need values for these outcomes. We have done this using a damage index that we created, somewhat arbitrarily, for this example.

3.1.2. Point system for event damage

We perform our valuation based on the following damage index scale.

0–10 points	Minor damage to the local built environment; rare minor injuries
10–20 points	Significant damage to the local built environment; common minor injuries; rare major injuries
20–30 points	Severe damage to the local built environment; significant damage to the surrounding built environment; common minor and major injuries; at least one fatality is likely
30–40 points	Severe damage to the local and surrounding built environment; significant damage to the natural environment; many minor and major injuries; several fatalities
40–50 points	Catastrophic damage to the local and surrounding built environment; permanent damage to the natural environment; many minor and major injuries; dozens of fatalities

Based on this scale and our judgment of the damage potentials of the various outcomes, we have assigned damage points to the outcomes, as shown in Table 1. We have also assigned an "uncertainty" to the damage points, which will

be used and described later (Sections 3.1.4–3.1.6); generally, the uncertainties represent underlying stochastic processes specific to the events but beyond the desired level of model detail.

Note that we will report the negative of the point value when referring to the damage index, so that negative numbers with higher absolute values indicate worse damage.

3.1.3. VaR for the case of no uncertainty in event damage

If there is no uncertainty in the damage associated with any outcome, then the VaR curve is actually a discrete probability *mass* function as opposed to a continuous probability *density* function. This function is shown in the simple bar graph of Fig. 3. Each event contributes to the VaR at exactly one value of the damage index, with a frequency determined by the event tree. We do not show the bar for the outcome of zero damage, which has a frequency of 0.9999 per year (assuming that our other outcomes cover all other possibilities), because it would be well off the scale of the chart. The cumulative probability mass function is shown in Fig. 4.

In the financial world, the actual "value at risk" is defined as the value that sets some lower confidence limit on the VaR probability function. For example, say that the value v represents a lower limit where 95% of the probability lies above it. Then we can state that we are 95% certain that we will lose no more than v over the time horizon used to construct the probability curve, or equivalently, "the value at risk is v." Based on the data in Fig. 4, we may make statements such as the following.

- We are 99.99% certain that we will suffer no damage from an LPG storage tank leak over the next year.
- We are 99.995% certain that we will suffer a damage value of no more than 30 points from an LPG storage tank leak over the next year.
- Over a 1-year time horizon, to a 99.995% confidence level, our value at risk from an LPG storage tank leak is 30 points.

The last two statements are equivalent.

The main assumptions used to generate Figs. 3 and 4 are that (1) the fault tree prediction of 10^{-4} LPG storage tank failures per year is accurate, (2) the event tree captures all possible failure outcomes and their associated probabilities, and (3) a single number is sufficient to capture the damage effects of each outcome. The next few sections address the relaxation of the third assumption.

3.1.4. VaR for the case of uniform uncertainty in event damage

In reality, many failure outcomes will result in a distribution of possible damage effects, due to stochastic variables such as atmospheric conditions or human factors. To capture the random nature of these processes, damage effects are often modeled as probabilistic functions instead of single values.
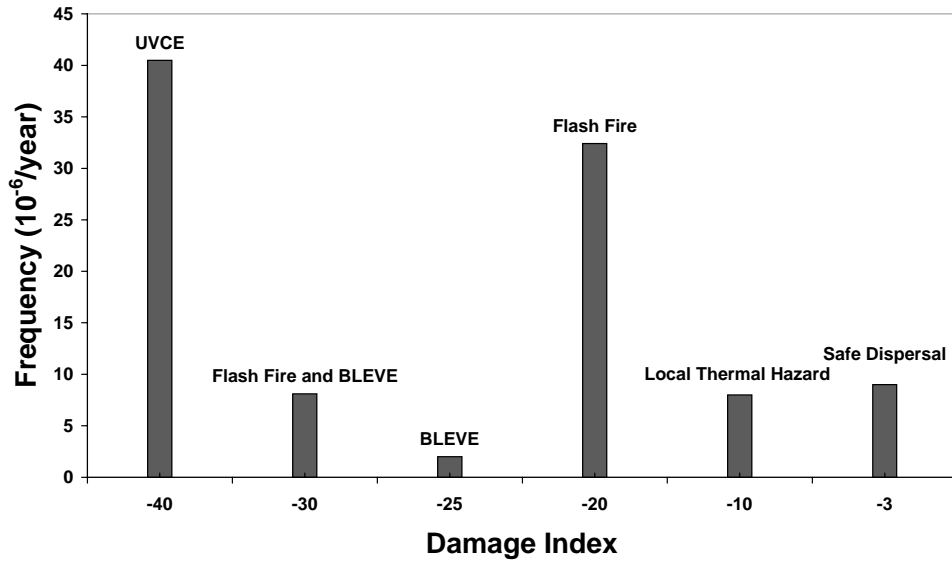
Fig. 3. Probability mass function for the LPG leak (no uncertainty).

The simplest approach is to assume a uniform distribution of frequency (or equivalently, probability) across some damage range, for each outcome. We demonstrate this approach using the numbers given in Table 1 for the LPG storage tank scenario. The uncertainties of the damage events in the Table are used as upper and lower bounds on the distributions, with the frequency being constant between them and zero elsewhere, and the total frequency (area under the curve) being equal to the frequency given in the Table. For example, the damage index associated with the "safe dispersal" outcome ranges from 2 to 4, with a uniform probability density of 4.5×10^{-6} events per year per damage point, yielding a total (integrated) frequency of 9.0×10^{-6} events per

year. Fig. 5 shows the resulting VaR curve. With the use of probability distributions to describe the damage effects, the curve becomes a probability density function, instead of a probability mass function. The curves for different individual outcomes now overlap in certain regions of damage index value and are combined additively in those regions. This additivity is justified because the event tree produces the outcomes as a set of complementary events, in a probabilistic sense. The corresponding cumulative curve is shown in Fig. 6. The effects of the sharp discontinuities in probability that exist at the edges of the uniform distributions are evident in the discontinuities of the slope at several locations in Fig. 6.
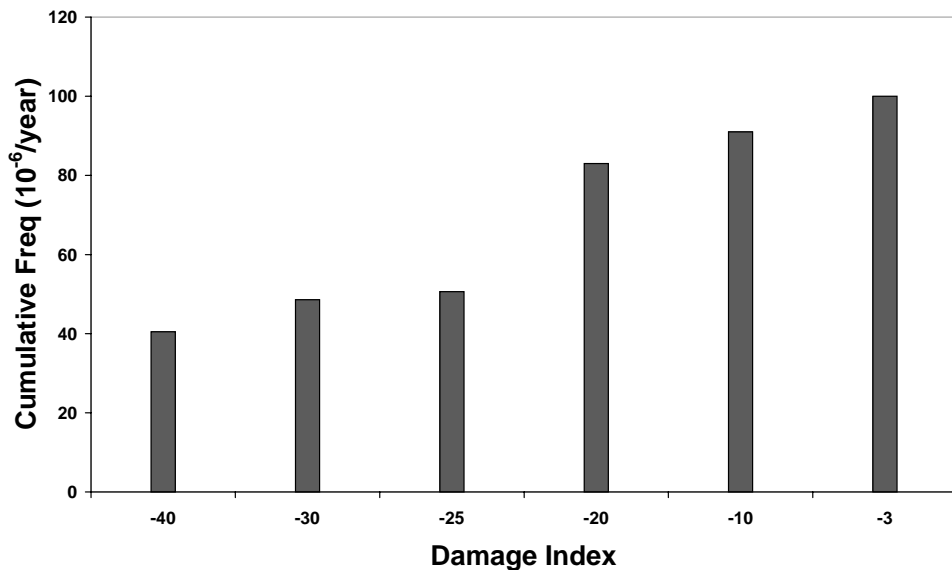

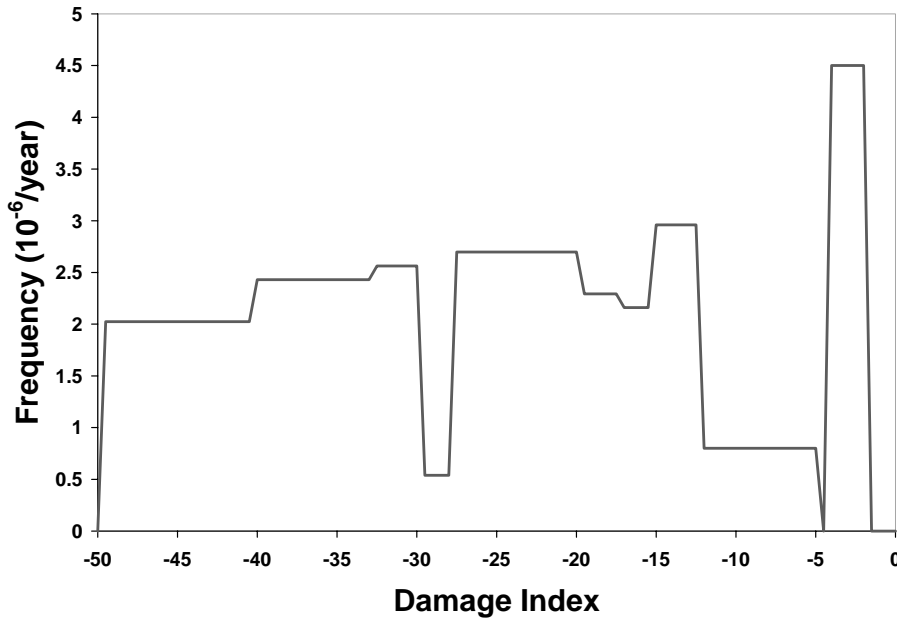Fig. 4. Cumulative mass function for the LPG leak (no uncertainty).

Fig. 5. Probability density function for the LPG leak (uniform uncertainty).

3.1.5. VaR for the case of Gaussian uncertainty in event damage

In this case, we assume that the damage effects are distributed normally. The uncertainties listed in Table 1 are now assumed to be the standard deviations in the Gaussian distributions. For clarity, the entire point scale for damage (Section 3.1.2) has been increased by a factor of 10 with new damage scores for each event. These new scores are reflected in Table 2. As with the uniform distributions, each Gaussian is normalized so that the total area under the curve equals the frequency given in Table 1. The probability density function is shown in Fig. 7 and the corresponding cumulative function is shown in Fig. 8.

With the Gaussian curves, both probability functions are now smoother. One problem with the normal distribution is that it has infinite range, which may have two undesirable side effects in the present analysis. First, all damage events make some contribution (albeit small) to the positive side of the value curve, which is not sensible. Furthermore, even minor damage events make some contribution (albeit small) to extreme damage values, which is also not sensible.
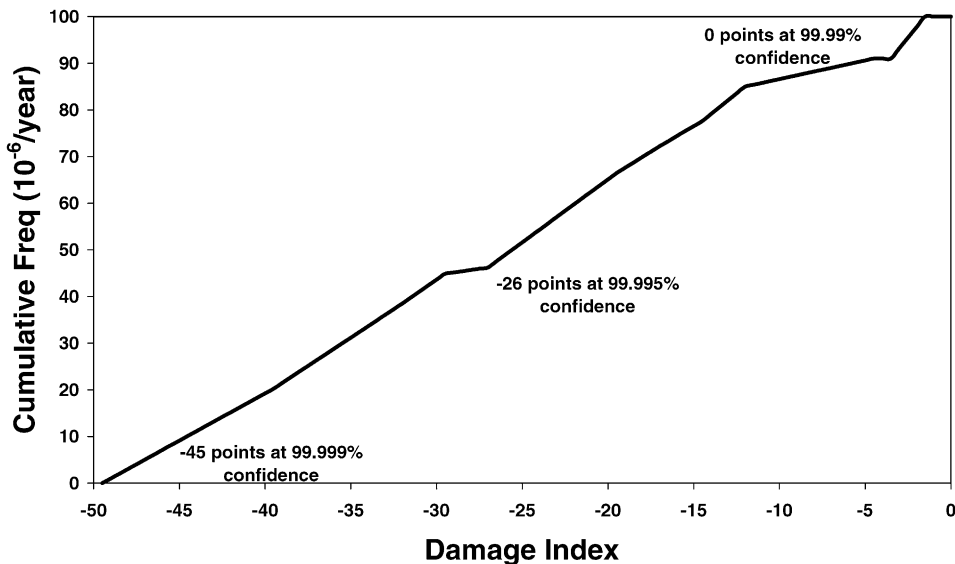

Fig. 6. Cumulative distribution function for the LPG leak (uniform uncertainty).

Table 2
Revised data for the LPG leak problem

Incident	Damage index	Uncertainty	Frequency (10^{-6} per year)
BLEVE	−200	25	2
Flash fire	−150	15	32.4
Flash fire and bleve	−275	20	8.1
UVCE	−425	20	40.5
Local thermal hazard	−30	5	8
Safe dispersal	−3	1	9

3.1.6. VaR for the case of beta uncertainty in event damage

An obvious fix to the problem mentioned above is to use a probability function with limited range. For this purpose, we employed the beta distribution, which has both lower and upper bounds. The parameters α and β for the beta distribution were chosen to match the averages and standard deviations (uncertainties) given in Table 1.

The density function is shown in Fig. 9, while the cumulative function is shown in Fig. 10. In theory, this is probably the best representation of the results, in that the individual damage events are bounded appropriately. In practice, it does not appear to be much different from the Gaussian results, on this scale.

3.2. Second example problem: loading of chlorine rail tank car

This example problem applies VaR analysis to a problem illustrated in chapter 8 of ref. [1]. The representative outcomes and their frequencies are taken directly from that ref-erence. The damage index used for this example was created by Khan and Abbasi [4,5].

3.2.1. Scenario description

In this example, we assume that an incident identification analysis has generated a set of representative events associated with a chlorine tank car loading facility, and we further assume that a combination of historical data and fault tree analysis has been used to estimate their frequency. The three representative outcomes and their associated frequencies are summarized in Table 3. Another parameter that affects the consequences of the incidents is prevailing wind conditions. We will assume eight possible wind directions that are given an equal probability of occurring.

Detailed descriptions of the possible incidents may be found in ref. [1], but we briefly outline them here. The main elements of the facility are a storage tank, a rail tank car, and associated transfer equipment. A small leak of liquid chlorine (∼2 kg/s for 10 min) might arise from a defective hose or valve, or an impact to a transfer pipe. A small vapor leak (∼0.2 kg/s for 20 min) might arise from the same sources. A large vapor leak (∼2 kg/s for 60 min) might occur due to a lifting of the relief valve under the stress caused by an external fire. In all three cases, the primary concern is the toxic effects of the released chlorine; the loading facility is located 100 m west of a residential area 400 m square, containing a uniformly distributed population of 400 persons.

3.2.2. Point scale for damage events

We use the accident hazard index (AHI) due to Khan and Abbasi [4,5]. While their approach provides a means to rank
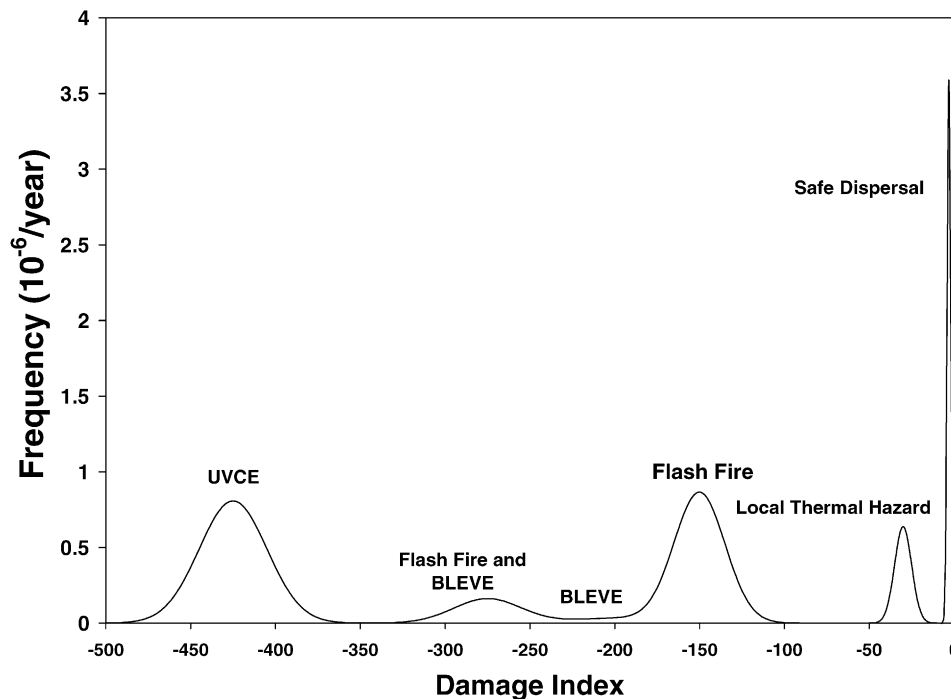

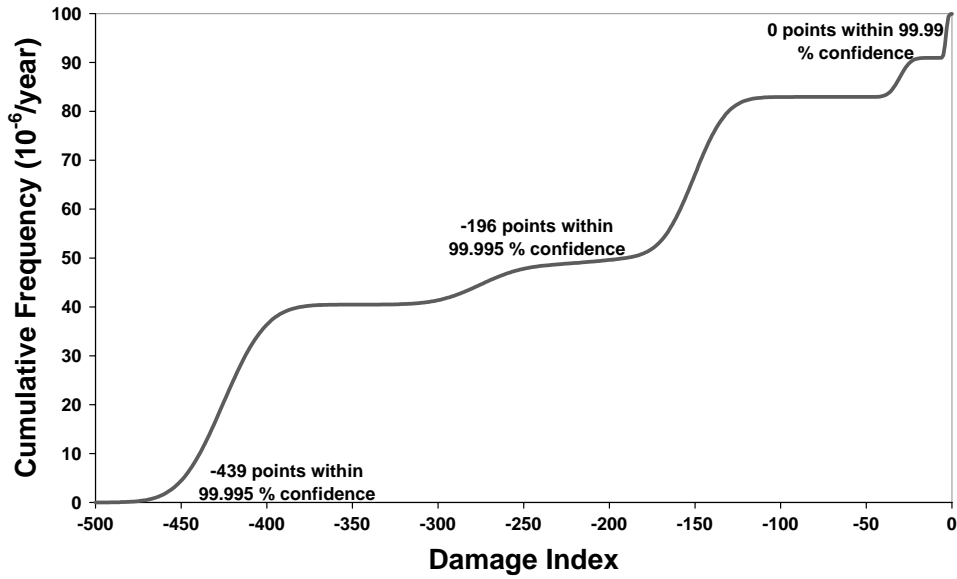Fig. 7. Probability density function for the LPG leak (Gaussian uncertainty).

Fig. 8. Cumulative distribution function for the LPG leak (Gaussian uncertainty).

three types of damage, namely thermal, mechanical (blast), and toxic, we will focus on toxic damage for this example problem of chlorine release. Khan and Abbasi's procedure for determining the contribution to the AHI of a toxic load involves the following steps. First, a parameter R is estimated from

$$R = \left(\frac{q}{\mathrm{LC}_{50}}\right)^{1/3} \tag{1}$$

where LC_{50} is the concentration $(\mathrm{kg/m^3})$ of chlorine vapor that is expected to be lethal to 50% of the exposed population and q is the total quantity (kg) released. The value of R is

used as input to a function that produces a dimensionless severity factor X. If the event is completely contained in the process area, this severity factor X is then the AHI. If an external effect (such as harm to population areas) is a concern, a population impact factor must be integrated with the severity factor X to produce the final AHI.

In this example, the direction of the prevailing wind during a release event is an extra stochastic factor. If the wind carries the chlorine vapor into the nearby residential area, an impact factor must be included. We assume that this will happen when the wind blows towards the northeast, east, and southeast (a total of 37.5% of the time). There
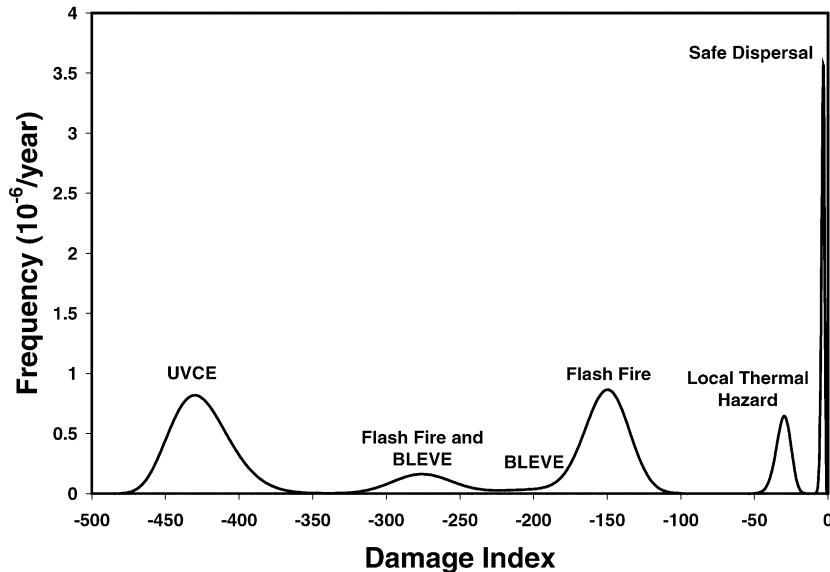

Fig. 9. Probability density function for the LPG leak (beta uncertainty).

Table 3
Data for the chlorine rail car problem

Chlorine potential accidents	Estimated frequency (per year)	Gas release rate (kg/s)	Total gas release (kg)	LC_{50} (kg/m^3)
Liquid leak	5.80E−04	2.7	1620	6.81E−04
Vapor leak	6.60E−04	0.26	156	6.81E−04
Relief valve discharge	3.00E−06	2.4	8640	1.67E−03
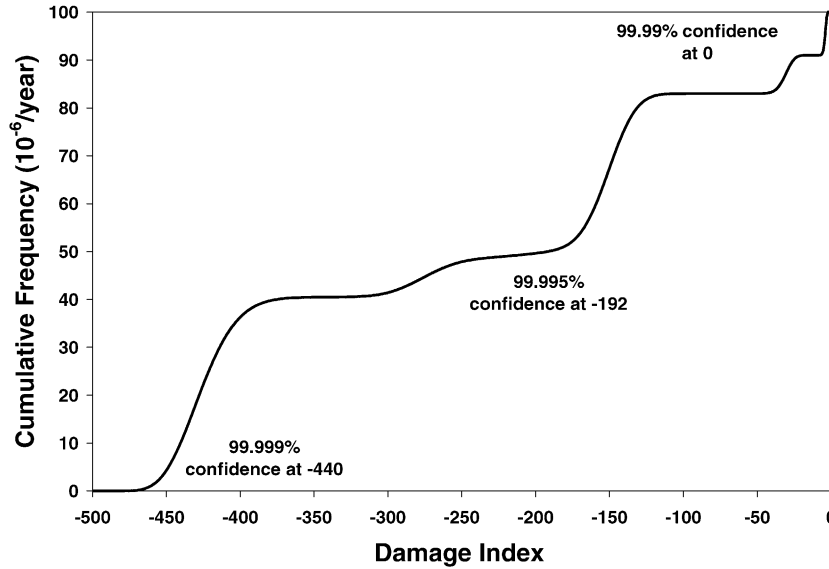
Fig. 10. Cumulative distribution function for the LPG leak (beta uncertainty).

are now two possibilities for the AHI associated with each event, one with the population impact factor and one without. The population impact factor is derived from a special formula derived from Khan and Abbasi; the input parameters are population density, which is the number of people (thousands) per square kilometer [4,5].

3.2.3. Analysis of scenario

Since no uncertainty in the hazard index was available, we carried out a simple probability mass function analysis for the VaR plot, as in Section 3.1.3.

The probability mass distribution function featuring the three unwanted events (with and without the population
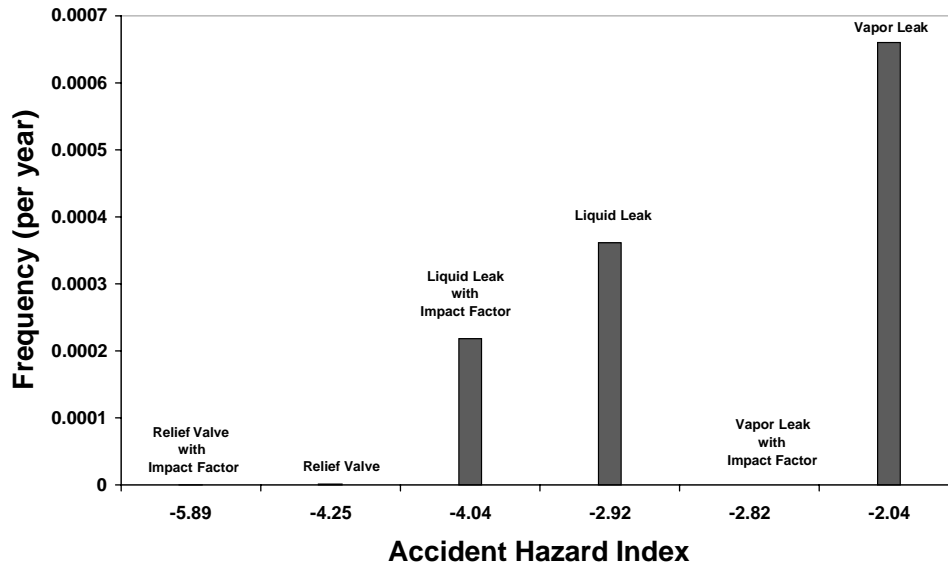

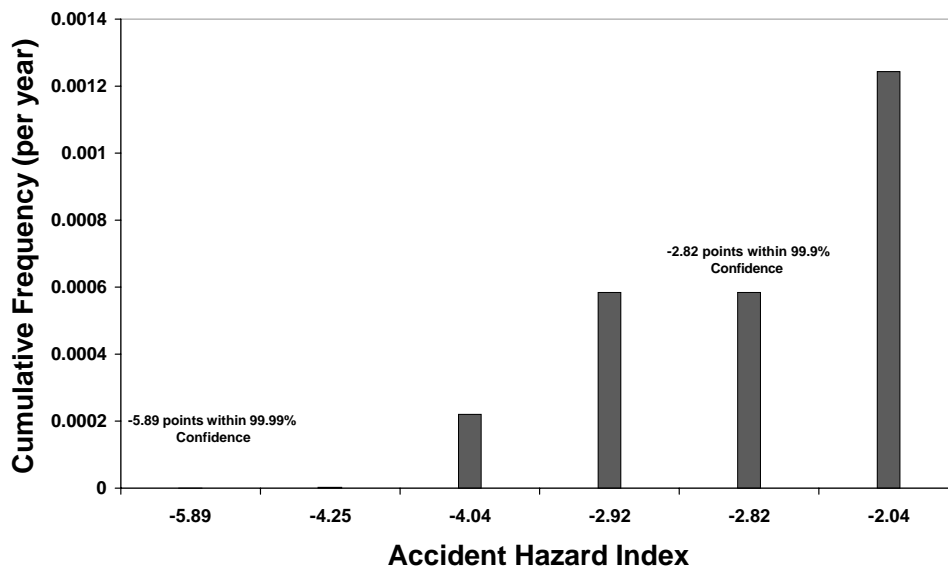Fig. 11. Probability mass function for the chlorine rail car problem.

Fig. 12. Cumulative distribution function for the chlorine rail car problem.

damage input) is shown in Fig. 11. The relief valve discharge had the highest hazard index, followed by the vapor and liquid leaks. The vapor indices had the greatest frequency. However, the wind did not affect the vapor leak's AHI, because the rate of gas release ($\sim$0.2 kg/s) was too small to be a hazard to a residential population 100 m away. The resulting probability mass function plot is shown below in Fig. 11. The cumulative mass density plot is shown in Fig. 12.

The following VaR statements may be made from the data.

- Over a 1-year time horizon, to a 99.9% confidence level, our value at risk from toxic leaks at the tank car facility is 2.82 on the AHI.
- Over a 1-year time horizon, to a 99.99% confidence level, our value at risk from toxic leaks at the tank car facility is 5.89 on the AHI.

4. Conclusions and future directions

We discussed how VaR concepts from finance might be used to make a better business case for process safety in the CPI. We demonstrated the procedure on two example problems from the CPQRA literature, creating VaR curves based on valuation with different damage/hazard indices (literature-based and customized). The effects of uncertainty in damage associated with possible events were included.

Future work will involve application to larger-scale, real-world problems, perhaps using an automated tool.

We will develop more complex and realistic approaches to valuation, including the use of different hazard indices (e.g. Mond, Dow) and a monetary scale. Other technical improvements will include the modeling of uncertainty in event frequency, as well as consequence.

References

[1] Center for Chemical Process Safety of the American Institute of Chemical Engineers, Guidelines for Chemical Process Quantitative Risk Analysis, American Institute of Chemical Engineers, New York, 1989.
[2] P. Jorion, Value at Risk: The New Benchmark for Managing Financial Risk, 2nd ed., McGraw-Hill, 2000.
[3] A.F. Barbaro, M. Bagajewicz, Managing Financial Risk in Planning under Uncertainty, AIChE J. 50 (5) (2004) 963–989.
[4] F.I. Khan, S.A. Abbasi, Hazard identification and ranking: a multi-attribute technique for hazard identification, Report CPCE/RA 22/97, Pondicherry University, Pondicherry, 1997.
[5] F.I. Khan, S.A. Abbasi, Accident hazard index: a multi-attribute method for process industry hazard rating, Process Safety Environ. Protect.: Trans. Inst. Chem. Eng. Part B 75 (4) (1997) 217–224.