A Comprehensive
Study of Risk Management

Risk Analysis for the Oil Industry

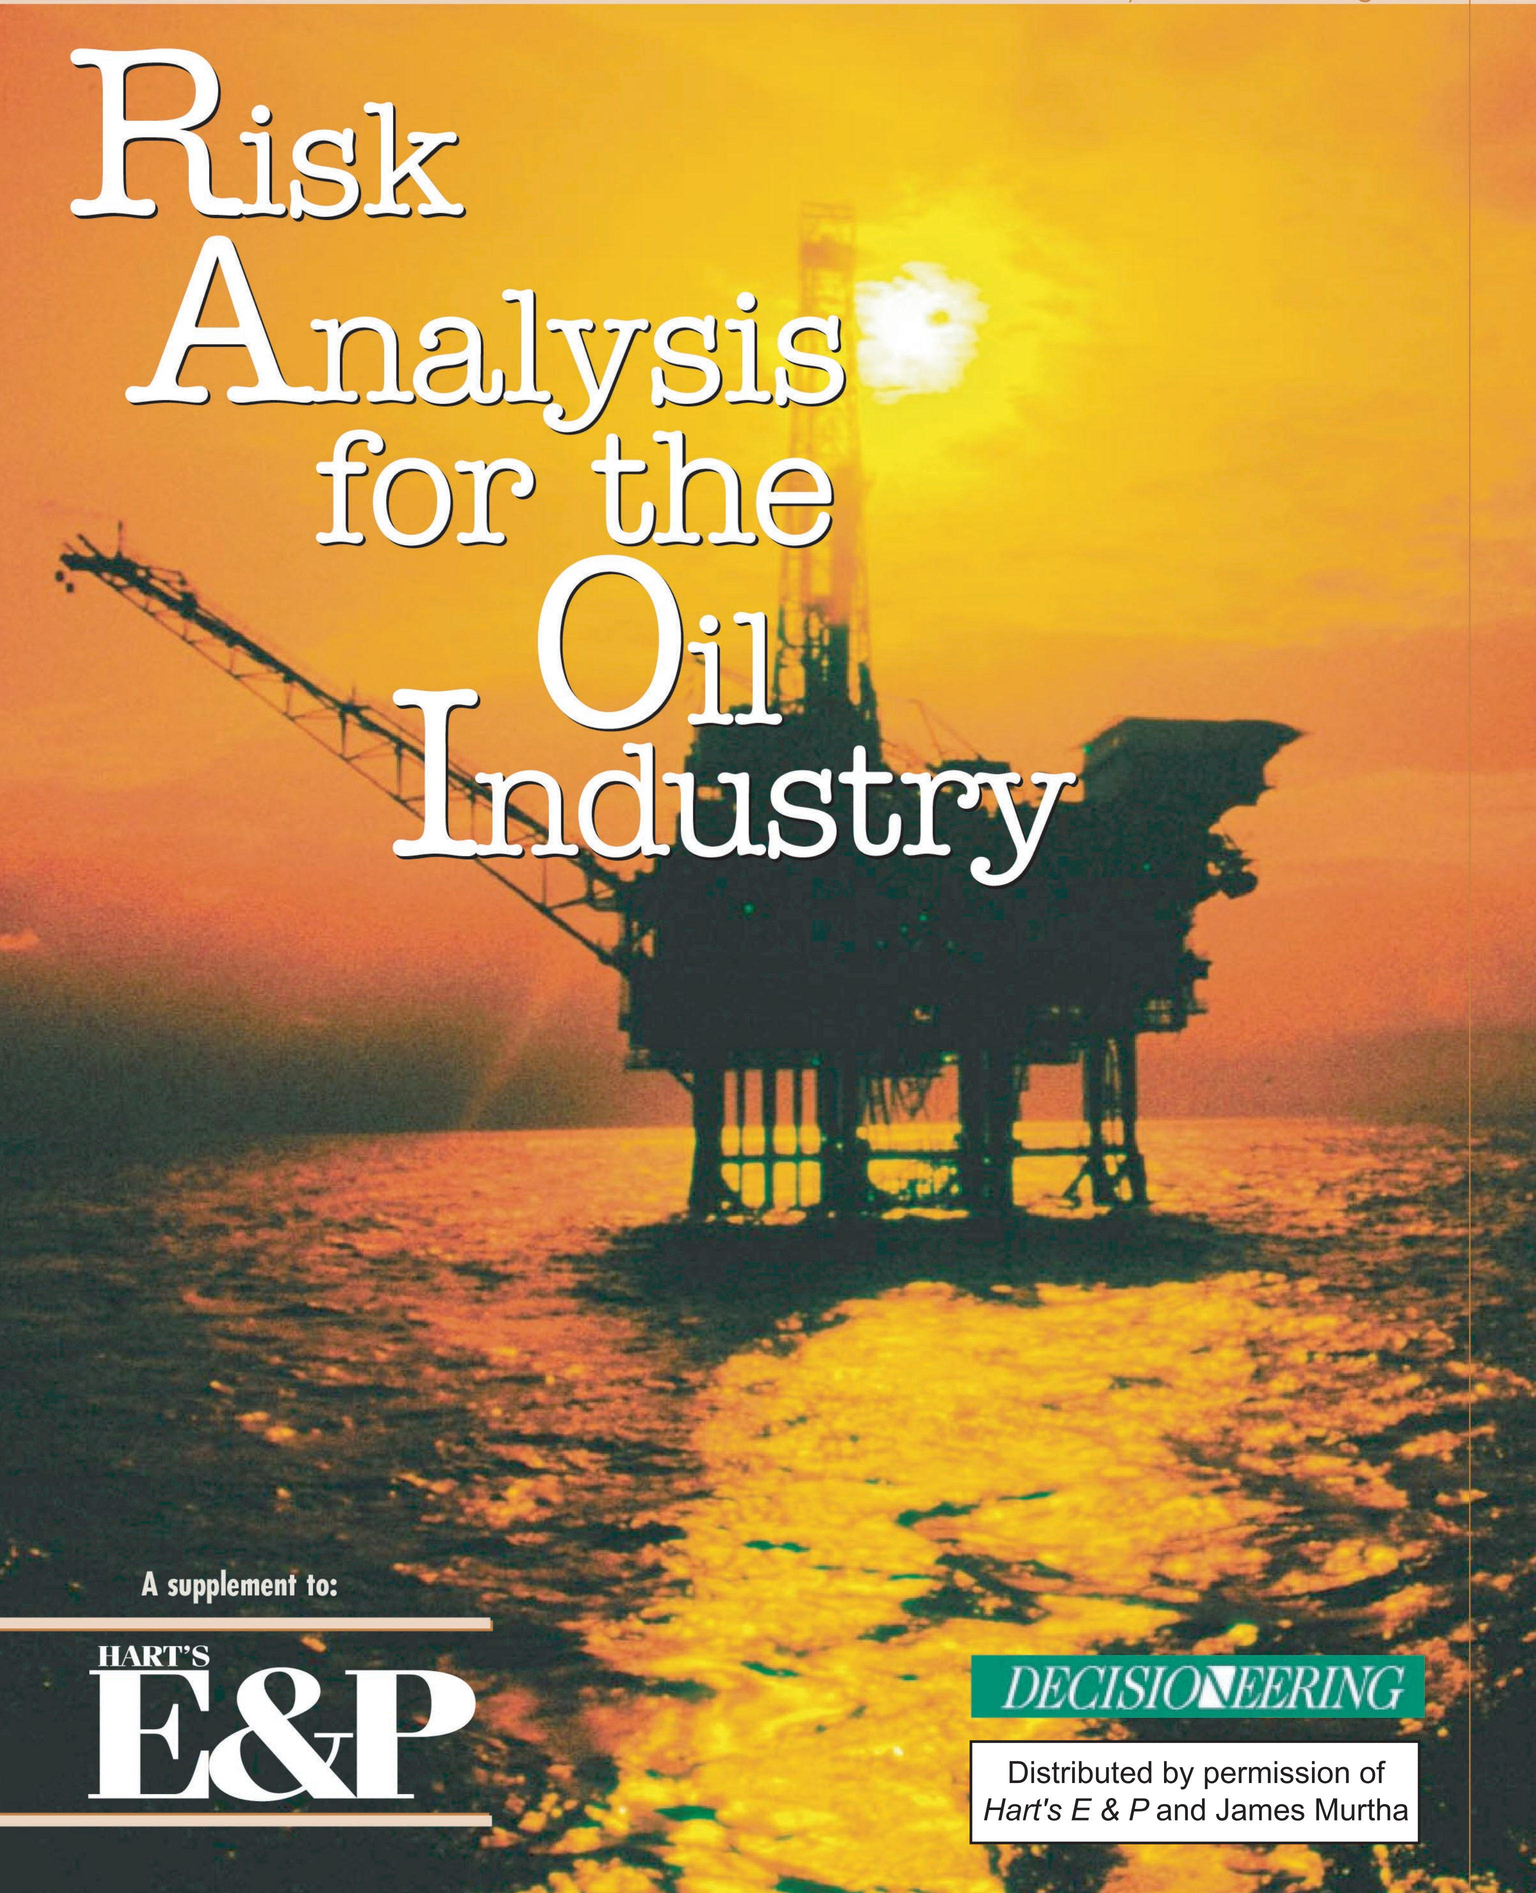

A supplement to:

HART'S
E&P

DECISIONEERING

Distributed by permission of
Hart's E & P and James Murtha

Biography

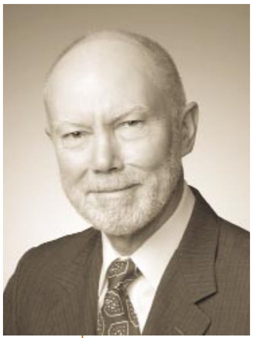

Jim Murtha, a registered petroleum engineer, presents seminars and training courses and advises clients in building probabilistic models in risk analysis and decision making. He was elected to Distinguished Membership in SPE in 1999, received the 1998 SPE Award in Economics and Evaluation, and was 1996-97 SPE Distinguished Lecturer in Risk and Decision Analysis. Since 1992, more than 2,500 professionals have taken his classes. He has published Decisions Involving Uncertainty – An @RISK Tutorial for the Petroleum Industry. In 25 years of academic experience, he chaired a math department, taught petroleum engineering, served as academic dean, and co-authored two texts in mathematics and statistics. Jim has a Ph.D. in mathematics from the University of Wisconsin, an MS in petroleum and natural gas engineering from Penn State and a BS in mathematics from Marietta College. ◆

Acknowledgements

When I was a struggling assistant professor of mathematics, I yearned for more ideas, for we were expected to write technical papers and suggest wonderful projects to graduate students. Now I have no students and no one is counting my publications. But, the ideas have been coming. Indeed, I find myself, like anyone who teaches classes to professionals, constantly stumbling on notions worth exploring.

The articles herein were generated during a few years and written mostly in about 6 months. A couple of related papers found their way into SPE meetings this year.

I thank the hundreds of people who listened and challenged and suggested during classes.

I owe a lot to Susan Peterson, John Trahan and Red White, friends with whom I argue and bounce ideas around from time to time.

Most of all, these articles benefited by the careful reading of one person, Wilton Adams, who has often assisted Susan and me in risk analysis classes. During the past year, he has been especially helpful in reviewing every word of the papers I wrote for SPE and for this publication. Among his talents are a well tuned ear and high standards for clarity. I wish to thank him for his generosity.

He also plays a mean keyboard, sings a good song and is a collaborator in a certain periodic culinary activity.

You should be so lucky. ◆

Table of Contents

Estimating Pay Thickness From Seismic Data

How do we estimate net pay thickness, and how much error do we introduce in the process? Usually, we specify two depths, to the top and bottom of the target interval and take their difference. The precision and accuracy of the thickness measurement, therefore, depends on the precision and accuracy of two individual measurements. The fact that we are subtracting two measurements allows us to invoke the Central Limit Theorem to address the questions of error. This theorem was stated in a previous article. A suitable version of it is given in the Appendix.

First, let us remind ourselves of some of the issues about measurements in general. We say a measurement is accurate if it is close to the true value, reliable or precise if repeated measurements yield similar results and unbiased if the estimate is as likely to exceed the true value as it is to fall short. Sometimes we consider a two-dimensional analogy, bullet holes in a target. If the holes are in a tight cluster, then they are reliable and precise. If they are close to the bullseye, then they are accurate. If there are as many to the right of the bullseye as the left, and as many above the bullseye as below, then they are unbiased.

With pay thickness, we are interested in the precision of measurement in a linear scale, which we take to mean the range of error. Our estimate for thickness will be precise if the interval of error is small. Consider the following situation.

Application: adding errors

A team of geoscientists estimated the gross thickness of a reef-lagoon complex that resides between two identifiable seismic events: the lower anhydrite at the top and a platform at the bottom. Logs from one well provide estimates of the distances from the top marker to the target facies and from the bottom of the facies to the platform – namely 100m and 200m respectively, and the facies thickness is 600m – the overall distance between the markers is 900m. In one particular offsetting anomaly, the depth measurement to the lower anhydrite is 5,000m, plus or minus 54m. The depth estimate to the platform is 5,900m, plus or minus 54m. What is the range of measurement for the thickness of the target facies?

First, we should ask what causes the possible error in measurement. In the words of the geoscientists, "If the records are good, the range for picking the reflection peak should not be much more than 30 milliseconds (in two-way time) and at 3,600 m/second that would be about 54m. This, of course, would be some fairly broad bracketing range, two or three standard deviations, so the likely error at any given point is much less. We would also hold that there is no correlation between an error in picking the lower anhydrite and the error in picking the platform."

A further question to the geoscientists revealed that the true depth would just as likely be greater than as it would be less than the estimate. That is, the estimates should be unbiased.

Solution

People who use worst case scenario arguments would claim that the top of the reefal facies could be as low as 5,154m and the bottom as high as 5,646m, giving a minimum difference of 492m. Similarly, they say the maximum thickness would be (5,754m-5,046m) = 708m. In other words, they add and subtract 2*54 = 108 from the base

But, is that what we really care about: the theoretical minimum and maximum?

Uncorrelated	*Normal*	*Triangular*	*Uniform*
99% confidence	[535m,665m]	[520m ,680m]	[503m, 697m]
95% confidence	[551m, 649m]	[539m, 661m]	[515m, 685m]
90% confidence	[558m, 642m]	[549m, 651m]	[526m, 674m]
Correlated (r=0.7)	*Normal*	*Triangular*	*Uniform*
99% confidence	[565m; 635m]	[555m, 645m]	[533m, 665m]
95% confidence	[573m, 627m]	[566m, 634m]	[548m, 652m]
90% confidence	[578m, 623m]	[572m, 629m]	[557m,643m]

Table 1. Effect of Distribution Shape on Ranges (maximum, minimum) for Net Pay Thickness

case of 600m to get estimates of minimum and maximum thicknesses.

But, is that what we really care about: the theoretical minimum and maximum? We may be more interested in a practical range that will cover a large percentage, say 90%, 95% or 99% of the cases.

A probabilistic approach to the problem says that the two shale depths are distributions, with means of 5,000m and 5,900m respectively. The assertion that there is no bias in the measurement suggests symmetric distributions for both depths. Among the candidates for the shapes of the distributions are uniform, triangular and normal. One way to think about the problem is to ask whether the chance of a particular depth becomes smaller as one moves toward the extreme values in the range. If the answer is yes, then the normal or triangular distribution would be appropriate, since the remaining shape, the uniform distribution, represents a variable that is equally likely to fall into any portion of the full range. Traditionally, errors in measurement have been modeled with normal distributions. In fact, K.F. Gauss, who is often credited with describing the normal distribution - but who was preceded by A. de Moivre – touted the normal distribution as the correct way to describe errors in astronomical measurements.

If the uncertainty by as much as 54m is truly a measurement error, then the normal distribution would be a good candidate. Accordingly, we represent the top and bottom depths as:

Bottom = **Normal(5,900m; 18m), where its mean is 5,900m and its standard deviation is 18m**
Top = **Normal(5,000m; 18m)**
Thick = **Bottom – Top –300m (the facies thickness is 300m less than the difference between the markers)**

We use 18 for the standard deviation because it is customary to regard the practical range of a normal distribution to be the mean plus or minus three standard deviations, which is 99.7% of the true (infinite) range. Thus the actual depths would be as much as 54m (=3 * 18) off from the estimated depths. Again, we must be careful to distinguish between the theoretical model limits and something of practical importance. Every normal distribution extends from –∞ to +∞. The standard practical range extends three standard deviations each side of the mean. Especially if we are talking about measurements, we know the values must be positive numbers.

Next, we run a very simple Monte Carlo simulation, in which we select 2,000 pairs of values from these two distributions and calculate Thickness. The results can be plotted up as a distribution, which the Central Limit Theorem (see below) predicts is itself another normal distribution, namely:

Thickness = Normal(600m; 25.5m)

In other words, the 3-sigma range of values for Thickness is [523m; 677m], not [492m; 708m] as the worst- and best-case scenarios would have us believe. In fact, 90% of the time Thickness will be between 558m and 642m, and 99% of the time the value of Thickness would fall between 535m and 665m.

Had we taken the distributions to be triangular (minimum, mode, maximum), the simulation would show that Thickness has a 99% probability range of [1,706.12ft; 2,231.08ft] or [520m; 680m]. In this case, we would use:

Top = **Triangular (5,046m; 5,100m; 5,154m)** and
Bottom = **Triangular (5,646m; 5,700m; 5,754m)** and
Thickness = Bottom – Top

The result is approximated by Normal(600m; 31.5m). The P5, P95 interval is [549m; 65m]. These and other approximations are summarized in Table 1.

Finally, we run this model using the uniform (minimum, maximum) distribution, specifically uniform (4946m,5054m) for the top and uniform (5,846m, 5,954m) for the bottom. Thanks to the Central Limit Theorem, the difference can still be reasonably approximated by a normal (not uniform) distribution – although as one can show, the triangular distribution type actually fits better – namely Normal (600m; 38.5m). The interval ranges shown in the last column of the table were obtained from a simulation with 6,000 iterations.

Summary and refinements

When we subtract one random variable from another the resulting distribution tends to be

We may be more interested in a practical range that will cover a large percentage, say 90%, 95% or 99% of the cases.

normal; its coefficient of variation (ratio of standard deviation to mean) reduces to about 30% from that of the original variables. When pay thickness is obtained by subtracting two depth measurements, the precision of thickness is better than the precision of the inputs. Thus rather than compounding errors, the effect is to reduce errors.

One aspect of Monte Carlo simulation that always needs to be considered is possible correlation between input variables. In the present context, the question is whether a large value for Top would typically be paired with a large or small value for Bottom. That is, if we overestimate the depth of Top, would we be more likely to overestimate or underestimate the value of Bottom or would the values be uncorrelated (independent)? When they were asked, the geologists said that if anything, the two values would be positively correlated: a large value for Top would tend to be paired with a large value for Bottom. Essentially, they argued that they would consistently select the top or the midpoint or the bottom of a depth curve. How would that impact the results?

We reran the simulation using a rank correlation coefficient of 0.7 between the two inputs in each of the three cases (normal, triangular, uniform shapes). The results, shown in the lower portion of Table 1, show a dramatic effect: the error intervals decrease by nearly 50% from the uncorrelated cases.

Conclusions

The Central Limit Theorem (see below) can be applied to measurement errors from seismic interpretations.

When specifying a range of error for an estimate, we should be interested in a practical range, one that would guarantee the true value would lie in the given range 90%, 95% or 99% of the time.

When we estimate thickness by subtracting one depth from another, the error range of the result is about 30% smaller than the error range of the depths.

If the depth measurements are positively correlated, as is sometimes thought to be the case, this range of the thickness decreases by another 50%.

The Central Limit Theorem

Let $Y = X_1 + X_2 + \ldots + X_n$ and $Z = Y/n$ where $X_1, X_2, \ldots, X_n$ are independent, identical random variables each with mean μ and standard deviation σ. Then:

- both Y and Z are approximately normally distributed;
- the respective means of Y and Z are $n*\mu$ and μ; and
- the respective standard deviations are approximately $\sqrt{n}\,\sigma$ and $\sigma/\sqrt{n}$

This approximation improves as n increases. Note that this says the coefficient of variation, the ratio of standard deviation to mean, shrinks by a factor of $\sqrt{n}$ for Y.

Even if the Xs are not identical or independent, the result is still approximately true: adding distributions results in a distribution that is approximately normal even if the summands are not symmetric. Moreover, coefficient of variation diminishes.

When is the approximation not good? Two conditions will retard this process: a few dominant distributions, strong correlation among the inputs. Some illustrations may help.

For instance, we take 10 identical lognormal distributions each having mean 100 and standard deviation 40 (thus with coefficient of variation, CV, of 0.40).

The sum of these distributions has mean 1,000 and standard deviation 131.4 so CV=.131, which is very close to 0.40/sqrt(10) or 0.127.

On the other hand, if we replace three of the summands with more dominant distributions, say each having a mean of 1,000 and varying standard deviations of 250, 300 and 400, say, then the sum has a mean of 3,700 and standard deviation, 575, yielding a CV of .16. The sum of standard deviations divided by the square root of 10 would be 389, not very close to the actual standard deviation. It would make more sense to divide the sum by the square of root of 3 acknowledging the dominance of three of the summands. As one can find by simulation, however, even in this case, the sum is still reasonably symmetric. ◆