The Use of the Pareto Shape Parameter as a Leading Indicator of Process Safety Performance

Fred Henselwood
NOVA Chemicals
Calgary, Alberta, Canada
henselfw@novachem.com

ABSTRACT

Metrics addressing process safety incident performance typically focus on frequency and severity statistics. Often these lagging metrics are not overly sensitive to actual performance, making trending and forecasting difficult. This paper presents the results from a statistical study of a large incident dataset where changes in the Pareto Shape parameter were observed as a function of time. This approach has been found to give far better insight into process safety performance than traditional incident metrics and readily relates back to concepts such as the "Incident Triangle" and "Layers of Protection". Through the application of this approach, trends within Process Safety Incident performance have been observed earlier, and more accurate forecasting has allowed for the identification of anomalies. In turn, these critical observations have allowed for the better structuring and targeting of process safety programs. Although incident data is generally considered as a lagging indicator this approach has clearly reduced the lag time associated with this type of data and has given valuable insight into the current status of process safety performance.

1. INTRODUCTION

Incidents are typically characterized by two key parameters with these parameters being severity and frequency, (i.e. the incident was this "bad" and occurs this "often") with aggregate data being used to assess process safety performance. For example, there were X incidents last year associated with a total cost of Y dollars. Unfortunately these traditional metrics carry with them significant statistical noise that can make both trending and forecasting of incident performance very difficult to accurately conduct. If one assumes that the number of incidents that occurs in a given time period is a function of random chance (with the expected frequency reflecting actual safety performance) a common modeling approach is to use the Poisson distribution function for estimating the observed number of occurrences.

The Poisson probability distribution function can be written as follows:

$$p(y) = \frac{e^{-\mu}\mu^{y}}{y!}$$

Where *p(y)* is the probability of *y* occurrences in a given time period and μ is the expected number of occurrences for that time period. When looking at aggregate incident data *y* is the measured number of incidents while μ is the expected number of incidents. Although the expected value for *y* is μ the standard deviation associated with *y* is $\mu^{1/2}$ and as such there can be sufficient statistical noise associated with *y* so as to blur the estimate of μ. For example, if the safety performance associated with a facility is expected to generate 12 incidents each year, observing 9 incidents one year and 15 incidents the next year would appear to represent a significant shift in performance, yet these values are well within the noise associated with random chance. As such, attempts to gain resolution beyond the square root of the expected sample size are not readily possible. Increasing the sample size by a factor of 4 will reduce the associated noise by a factor of 2, therefore ***increasing sample size*** is one strategy that can be used to reduce noise, but there are limits to the improvements in resolution that can be made through this approach.

When viewing incident severity totals such as total incident cost, not only are there limitations associated with the underlying number of event occurrences there are also issues associated with the magnitude of each individual incident. For example, some incidents might only cost a few dollars while other incidents can be on the order of millions of dollars. As such, when this data is aggregated one or two significant events may dominate over the rest of the data pool. Again this situation can be readily described using the Poisson equation in that when one looks at the most extreme incidents where only a few incidents of that magnitude are expected in a given time period the noise associated with random chance is quite large. As an example, if 4 incidents over a million dollars were expected, observations between 2 and 6 occurrences would not be unusual. If the million dollar plus events dominate the observed total incident cost fluctuations between $2,000,000+ one year and $6,000,000+ the next year would not be unexpected and as such this statistical noise would blur any true changes in underlying incident performance.

Despite these limitations incident data can provide significant insight into process safety performance. Improvements in resolution can be achieved by moving away from traditional measures such as frequency and severity. The focus of this paper is on an approach that looks at the observed distribution of incidents and how this distribution function relates to process safety performance.

2. THEORY

If one views an incident as a collection of mistakes and/or accidents, incident distributions such as those described by the incident triangle can readily be explained through an assumption that the severity of the incident is a function of the number of mistakes and/or accidents that occurred to produce the incident. Similarly, if one considers the concepts behind the layer of protection analysis (LOPA) model the severity of an incident relates to the number of layers of protection that failed.

Based on these models, the severity of an incident can be mathematically expressed by the following function:

$$Severity = L \times J^n$$

Where *L* is a constant representing a base level of severity, *J* represents the jump in incident severity associated with each additional failure, and *n* represents the number of failures that have occurred. For example, if *J* equals 10 (an order of magnitude) each additional mistake or accident associated with an individual incident would increase the severity of the resulting incident by an order of magnitude. Graphically you can view this as the bottom layer of an incident triangle representing $n = 1$, the next layer of the incident triangle being $n = 2$, and so on and so on, with each layer of the incident triangle representing incidents which are an order of magnitude greater in severity than the previous layer. Within the LOPA model this equation assumes that the magnitude of the incident severity is a function of the number of layers of protection that have failed.

Building from the LOPA model and working with the following assumptions: that each layer of protection is independent; that each layer of protection is responsible for the same jump (*J*) in severity; and that each layer of protection has the same likelihood of failure, the following series of equations can be developed.

$$P(1) = 1 - P(J)$$

Where *P(1)* represents the probability that only 1 layer of protection will fail and *P(J)* represents the probability that a typical layer of protection will fail. If the probability of a layer of protection failing was 1 in 10 then there is a 90% chance that only the original layer of protection will fail. Alternatively, there is a 10% chance that one or more additional layers of protection will fail. If we look at the probability of the second layer of protection failing and not the third layer of protection failing the following equation can be derived:

$$P(2) = P(J) - P(J)^2$$

Again if there is a 10% chance of failure for each layer of protection failing the probability that only two layers of protection will fail is 9%. When looking at the probability that the n^{th} layer will fail (and not the n^{th} +1 layer) the following equation can be derived:

$$P(n) = P(J)^{n-1} - P(J)^n$$

Which can be expressed as the following cumulative distribution function:

$$F(N) = 1 - P(J)^n$$

Rearranging the severity and the cumulative distribution functions in terms of *n* yields the following two equations:

$$n = \frac{\ln(Severity/L)}{\ln(J)}$$

$$n = \frac{\ln(1 - F(N))}{\ln(P(J))}$$

Equating the two equations to each other gives:

$$(Severity/L)^{1/\ln(J)} = (1 - F(N))^{1/\ln(P(J))}$$

Solving for *F(N)* gives:

$$F(N) = 1 - (L/Severity)^{-\ln(P(J))/\ln(J)}$$

Which has the same form as the Pareto distribution function with the following parameters:

$$Location = L$$
$$Shape = -\ln(P(J))/\ln(J)$$

As such if one assumes that the likelihood of each layer of protection is approximately the same (*P(J)*) and that the increase in the magnitude of the severity (*J*) is the same for each layer of protection then a curve fit of incident data, using a Pareto distribution function, should provide insight into the effectiveness of the related process safety performance.

3. DATA TREATMENT

Eight years of incident data (1999-2006) were utilized in the following study, representing just over 1,000 events each with a severity of $30,000 (2007 USD) or greater. (Incident costs were corrected for inflation and currency issues.) The dataset primarily represents two olefin manufacturing regions (Joffre Alberta and Sarnia Ontario) the size and operation of which was relatively stable through this time period with some expansion occurring in 2000 within the Joffre area. The selection of $30,000 as a minimum incident threshold value represented a balance between maximizing the amount of data available for the study while ensuring high data quality and full incident reporting. Lower value incidents are less likely to be fully reported and fully investigated relative to larger incidents and as such data quality typically decreases for lower value incidents. Further, the $30,000 threshold placed an emphasis on process safety related incidents relative to occupational health related incidents and generally minimized the portion of the reported incident costs associated with the related incident investigation.

The incident data was sorted to form a cumulative distribution plot. This function was then transformed to a linear function allowing for standard curve fitting, which was conducted using Excel.

$$\ln\left(Cost/Location\right) = m \times \ln\left(1 - F(X)\right)$$

Where *Cost* is the observed incident severity measured in 2007 USD, *Location* is the minimum incident value for the study ($30,000), *F(X)* is the measured percentile divided by 100, and *m* is the slope of the resulting plot. Further, *m* can be related to the Pareto Shape parameter through the following equation:

$$m = -1/Shape$$

As such curve fitting of the observed incident data should readily give information with respect to the quality of the related process safety programs.

Although the Shape parameter and related layer of protection failure probability have been linked to process safety performance the usefulness of these parameters is limited in that the data they express is not readily tangible as their meaning is somewhat abstract. As an example, people cannot readily relate to a change in Shape parameter from 1.0 to 1.2. As such, to improve the insight gained through this analysis an expected cost function has been utilized and represents the summation of a series of uniformly distributed incidents (on a percentile basis) as described by the Shape parameter and an estimated event frequency. The expected cost function is given below:

$$ECF = \sum_{i=1}^{I} Location\left(\frac{1}{1 - \left(\left(i - \frac{1}{2}\right)/I\right)}\right)^{1/Shape}$$

Where *I* is the expected number of incidents.

The benefit to the use of this metric is that it approximates the median observed total incident cost resulting in a value that is not overly influenced by the noise associated with extreme incidents within the dataset. This metric also has a benefit relative to the median incident cost for a Pareto distribution function (*Location x 2 ^ (1/Shape)*) in that as the number of expected incidents increases the extreme percentile events become more important to the estimated value and act to skew the expected cost function towards higher values (i.e. two incidents at the 50th percentile do not add to give the same value as one incident at the 25th percentile and one incident at the 75th percentile). As such, as the number of expected incidents rises the average incident cost also rises (to a maximum value as defined by the above Pareto distribution function median value). Although this metric gives a good approximation of the expected cost of a basket of incidents this parameter does suffer from those limitations previously discussed with respect to the noise associated with the expected number of incidents used in the calculations.

4. OBSERVATIONS

Figure 1 illustrates the linear Pareto distribution fit obtained for the entire dataset.

Figure 1.

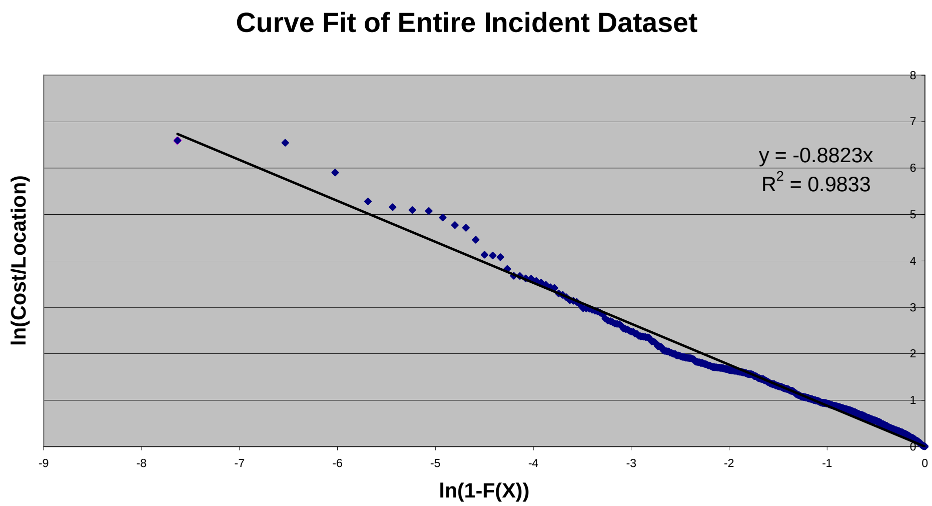

Figure 1 shows the entire incident dataset and the generally good agreement obtained with respect to the alignment between the Pareto distribution function and the observed incident data. Working from the measured slope the Shape parameter can be determined (*Shape* = 1.1334) which then allows for the determination of the probability of a layer of protection failing. (*P(J)* = 1 in 13.6 +/- 0.1, assumes *J* = 10 for each layer of protection).

In addition to being able to fit the entire dataset, each month was then analyzed based on the data recorded for the preceding year. In this way it was possible to observe changes in the Shape parameter as a function of time, independent of seasonal factors, so as to assess changes in process safety performance.

Based on observations that the linear fit was at times significantly influenced by the one or two largest incidents within a given dataset, particularly when looking at one-year blocks of data, the linear fits were developed using the 0-98th percentiles of the reviewed datasets. In this way the largest incidents influenced the assignment of the percentiles but did not directly influence the linear fit. The difficulty in fitting these extreme data points is consistent with the limitations already discussed in the introduction with respect to the noise associated with low probability events. As an example, some one-year blocks of data may include a 1 in 10 year event while other may not even include a 1 in 6 month event.

Figure 2 shows the measured Shape parameter as a function of time as measured for the one-year blocks of data by month.

Figure 2.

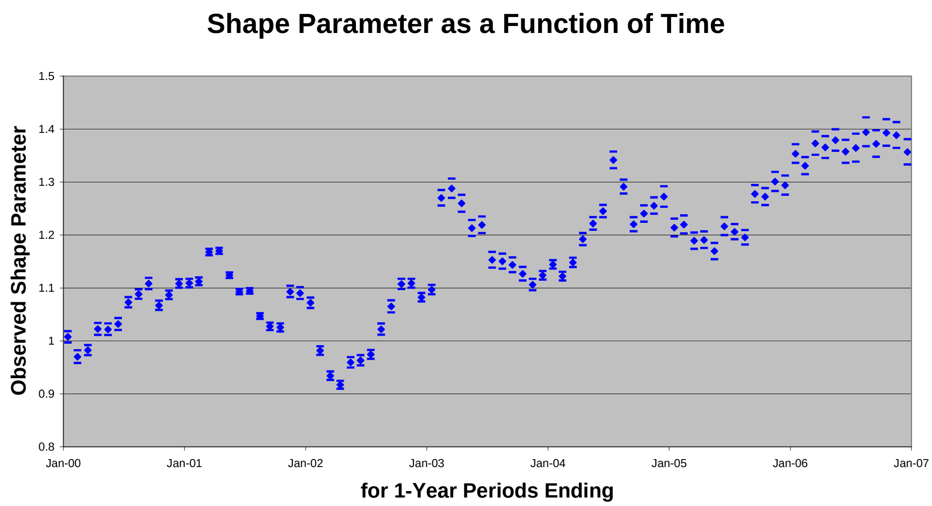

Figure 2 shows the observed trend in the Shape parameter as a function of time. Although associated with some variability the Shape parameter is clearly trending towards higher values. During the study period the likelihood of a layer of protection failing ($1/P(J)$) has gone from approximately 1 in 10 to 1 in 25 (assumes J = 10).

As a comparison Figure 3 shows the traditional metrics associated with process safety performance (Frequency and Total Cost).

Figure 3.

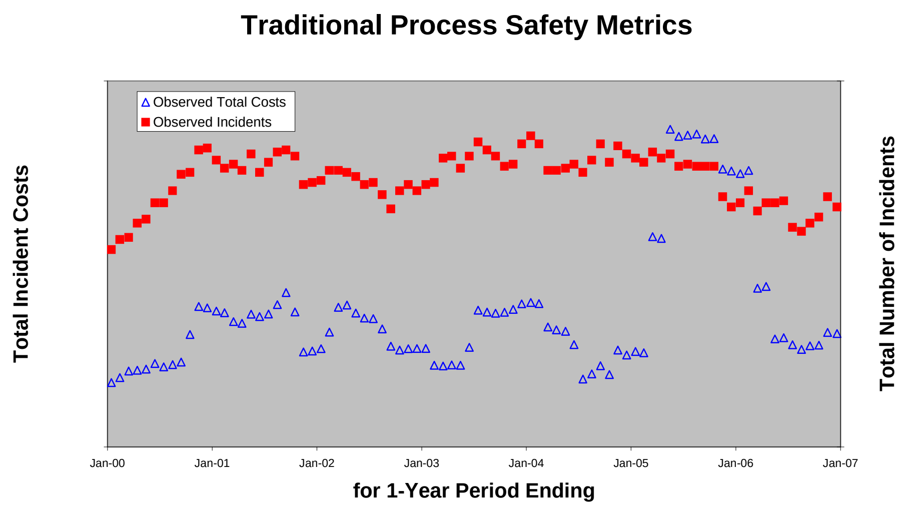

Figure 3 shows the number of incidents observed for the previous year by month (closed boxes) and the total costs associated with those incidents (open triangles).

Clearly the ability to spot trends with these traditional metrics (incident count and total incident cost) is significantly limited relative to the resolution given by the plot of the Shape parameter. Figure 4 illustrates a comparison between the observed total incident costs for one-year periods by month and the expected cost function for these same periods.

Figure 4.

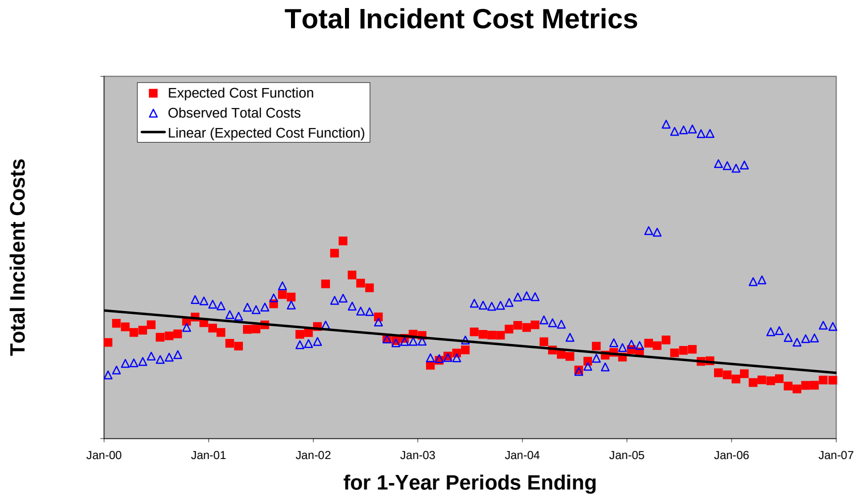

Figure 4 shows the expected cost function calculated for one-year periods by month (closed boxes) and the observed total costs for these same periods (open triangles). In addition a linear trend line has been plotted for the expected cost function.

Clearly much of the noise associated with the observed data has been removed using the expected cost function. The linear fit of the expected cost function indicates that aggregate incident costs have been reduced by over 50% during the course of the study period, representing a significant improvement in process safety performance. More detailed comparison of the incident datasets for 2000 and 2006 shows that approximately 15% of the estimated cost savings is related to the occurrence of fewer incidents (frequency reductions) and the remaining 35% reduction is related to a 40% reduction in the average cost per incident (severity reductions).

A further benefit of the ability to model incident data is that incident forecasting is now possible and the likelihood of incidents of a given magnitude can now be calculated. Figure 5 shows a comparison between the number of expected significant incidents and the observed number of significant incidents as calculated based on the obtained Shape parameters and measured incident frequencies (significant incidents have been defined based on the cost per incident exceeding a set threshold).

Figure 5.

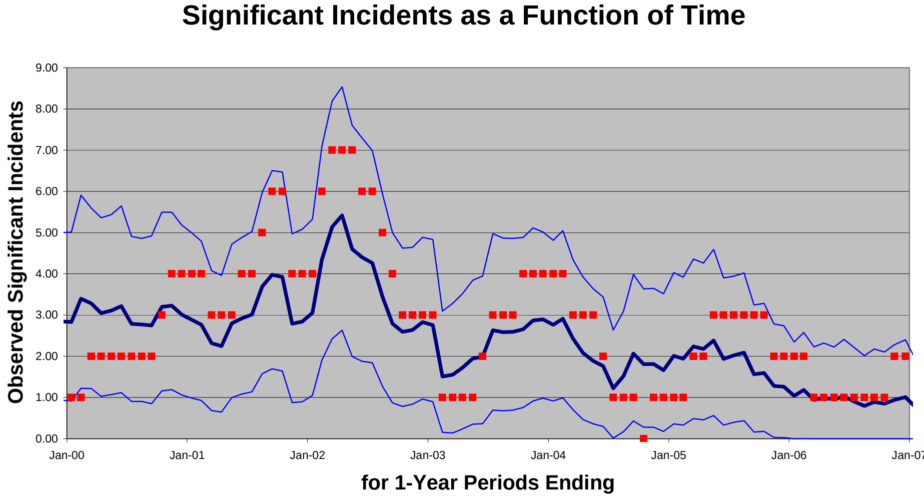

Figure 5 shows the modeled number of significant events for one-year periods by month (heavy line) along with confidence limits representing approximately two standard deviations (light lines). The closed boxes represent the observed number of significant incidents for each period.

As with the expected cost function plot the modeled number of significant events clearly indicates a downward trend. Very notably the number of significant incidents has been reduced from approximately 3 per year to 1 per year over the study period. As significant incidents require multiple layers of protection to fail the frequency reductions observed at this level are greater than those observed for smaller incidents. This is consistent with the originally proposed severity equation where there was an exponent relationship with the number of layers of protection that must fail to produce an incident of a given magnitude.

5. UNEXPECTED BENEFIT

The model approach and assessment of metrics such as the Shape parameter, typical layer of protection failure probabilities, and the expected cost function have given additional insight into the underlying process safety performance. Trends that are detectible with these functions clearly exist over the entire study period and are very clear with respect to the last five years. By comparison, traditional metrics failed to give this type of insight as the associated noise with these metrics overwhelmed the underlying signals they were intended to convey.

When modeling the expected number of significant events (including severities greater than that illustrated in Figure 5) there was a bias towards observing more incidents than what the modeling was predicting. This bias became even greater when larger severities were considered. A review of the incidents potentially making up these anomalies revealed that a few of these incidents were not consistent with the original assumptions. These "freak" incidents typically represented catastrophic events that were initiated outside of the control of the corporation. Examples include extreme weather events and major power outages. In these situations the assumption that a sufficient number of layers of protection existed to prevent the escalation of these incidents does not hold to be true. Figure 6 illustrates a comparison between the expected cost function and the observed total costs after the removal of these freak events from the dataset.

Figure 6.

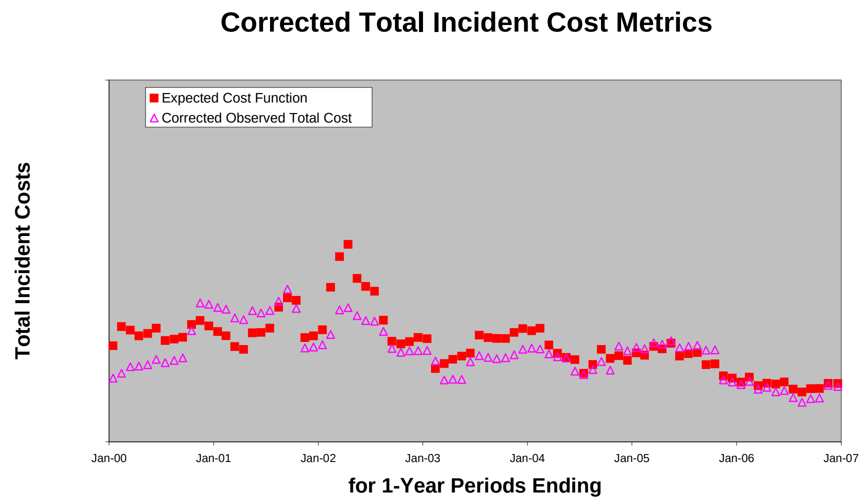

Figure 6 shows the expected cost function calculated for one-year periods by month (closed boxes) and the corrected observed total costs for those same periods (open triangles).

The alignment between the corrected observed total costs and the expected cost function readily shows the strength of this modeling approach. Clearly these freak incidents have always represented a portion of the overall incident costs. However, as the incident costs associated with "normal" incidents involving multiple layers of protection has been significantly reduced the percentage of the total incident costs represented by these freak incidents has increased. Without the use of metrics affording this level of resolution these types of observations are not readily possible and as such the allocation of process safety resources cannot be fully optimized.

6. CONCLUSIONS

The derived relationships and the agreement between theory and the observed data indicates that incidents as a group can be readily modeled (just over a thousand incidents were described by a two parameter curve fit with a R^2 value 0.9833). Further, the derived relationships readily generated far better resolution with respect to the ability to trend process safety performance than that which could be achieved using traditional metrics such as incident counts and related total incident costs. As an example trends showing improvements in performance were not observed based on total incident cost data whereas the derived expected cost functions showed very clear trends. In addition to identifying trends within the data it was possible to identify incidents (less than 1% of all incidents) that were inconsistent with the underlying assumptions through gaining significant insight into the observed distribution of incidents this then allows for the appropriate targeting of process safety related programs.

Secondary to the modeling successes the study also readily demonstrated and quantified the significant improvements that have been made with respect to process safety performance during the study period with the greatest improvements being made with respect to the reduction in the frequency and severity of major incidents.

7. REFERENCES