

POWER LAW ANALYSIS IMPLICATIONS OF THE SAN BRUNO PIPELINE FAILURE

Dr. Ken Oliphant, P.Eng.
JANA Corporation
Aurora, ON Canada

Wayne Bryce, P.Eng.
JANA Corporation
Aurora, ON Canada

William Luff, M.A.Sc.
JANA Corporation
Aurora, ON Canada

ABSTRACT

When major pipeline incidents occur there is always a question as to how applicable the learnings from that incident are across the industry. To address this question for the San Bruno pipeline failure in 2010, an analysis of historical transmission pipeline industry events was conducted to determine if San Bruno was consistent with past industry performance or whether it was an outlier event. This paper draws on Power Law analysis to generate a characteristic curve of past transmission pipeline accidents in the US. Power Law, or Pareto, behavior has been observed for a wide variety of phenomenon, such as fire damage, earthquake damage and terrorist attacks. The size of these events is seen to follow not the typical normal distribution but the Power Law distribution, where low probability – high consequence (LPHC) events play a more significant role in the overall risk picture. Analysis shows that the consequences of pipeline incidents in a variety of pipeline industries (gas distribution, gas transmission, gas gathering and hazardous liquid pipelines) are seen to exhibit Power Law behavior. The Power Law model is seen to capture the distribution of the size of consequences from pipeline incidents and defines the relationship between the size of an incident and its frequency. Through characterization of these distributions, it is possible to project the likelihood or expected frequency of events of a given magnitude and to assess if a given incident fits within historical industry patterns; i.e. whether the incident is consistent with past observations or is an outlier.

The Power Law analysis shows that the San Bruno incident, which caused eight fatalities and an estimated \$380 million in property damage in 2010, is not an outlier. Rather, this incident lies on the Power Law curve for historical transmission pipeline incidents, with an estimated frequency of once every 40 years. The event is consistent with the history of gas transmission pipeline consequences in the US. This paper argues that the San Bruno incident, therefore, provides lessons relevant to the industry as whole.

INTRODUCTION

When major incidents occur, the question arises as to how reflective the incident is of the state of the industry and, hence, as to how applicable the learnings from the incident are to the entire industry. In this analysis, this question is considered for the San Bruno pipeline failure. Was this event is an outlier event or is it more reflective of the general state of gas transmission pipelines? Power law analysis shows that this event is consistent with the history of gas pipeline failures in the US and in not an outlier incident. This suggests that the learnings from the San Bruno incident are broadly applicable.

NOMENCLATURE

Pipeline risk management, pipeline incident consequences, Power Law

There is a broad range of potential consequences for a pipeline incident. From the PHMSA database for gas transmission pipeline incidents, the property damage for reported incidents ranges from a few thousand dollars to over \$350 million. Similarly, for gas distribution pipelines, reported property damage due to pipeline incidents ranges from a few thousand dollars to greater than \$42 million. While there are deterministic factors at play, such as pipe size, operating pressure and pipe location (e.g. HCAs¹), there are also more random factors at play. For a gas distribution pipeline of the same size and operating pressure, for example, we can see leaks that result in very little consequence (e.g. those that are found by leak survey and repaired prior to a significant event), leaks of moderate consequence (e.g. those where gas accumulation and ignition occurs with limited property damage), right through to significant incidents (e.g. major property damage with injuries and/or fatalities). Each of these consequences will have an associated probability. Some will be more likely than

¹ High Consequence Areas

others – it is much more likely that a leak will be found and repaired than result in a significant incident, for example.

For a given future incident, therefore, there is a probability distribution of potential consequences that will be specific to the local environment surrounding that incident. In order to understand the risk associated with that incident, we need to understand this probability distribution. Likewise, for a pipeline system with multiple possible future leaks, there will be an overall probability distribution of potential consequences. It is the overall distribution that gives us insight into the true system risk. The question is, then, what do these potential consequence distributions look like and how do we estimate them?

Pareto Consequence Distributions in Pipeline Incidents – Power Law Behavior

In our work examining and modeling pipeline consequences, we observed that pipeline consequences appear to follow a very specific distribution. Pipeline consequences, along with many phenomena² such as fire damage, earthquakes, floods and power blackouts, follow Power Law or Pareto-type distributions where a small number of incidents account for the majority of the overall damage and, hence, risk. This type of behavior is often referred to as the 80/20 rule (or Pareto’s Law), where, for example, 80% of the damage comes from 20% of the incidents. While the specific ratios vary for different phenomena (95% of damage from 5% of incidents, 90% of damage from 10% of incidents, etc.), the concept is the same – a small number of events accounts for the majority of risk. This type of behavior gives rise to the low probability-high consequence events that can often dominate the risk picture.

Figure 1 shows the Power Law relationship for the frequency versus property damage for PHMSA³ reported gas distribution incidents in the US from 1992 to 2011, based on publicly available data from the PHMSA website⁴. The number of incidents resulting in different levels of property damage is shown for reported incidents with greater than \$100,000 damage. This lower bound was used to provide the best fit to the Power Law. This lower bound is believed to arise due to the requirements for size of incidents reported. The log⁵ of the frequency or number of events is plotted versus the log of the property damage that occurred for a total of 1095 reported incidents (the incident data for all causes) in a log-log plot. A

² A. Clauset et al, *Power Law Distributions in Empirical Data*

³ Pipeline and Hazardous Materials Safety Administration

⁴ <http://primis.phmsa.dot.gov/comm/reports/safety/PSI.html>

⁵ The log or natural logarithm of a number is the exponent to which the base 10 must be raised to produce that number. For example, the log of 1000 is 3, because 1000 is 10 to the power 3: $1000 = 10 \times 10 \times 10 = 10^3$. When data is plotted on a log scale, each increment is an order of magnitude higher than the previous – 1, 2, 3 on a log scale corresponds to 10, 100, 1000 on a linear scale. A relationship that is exponential in nature will plot as a straight line on a log-log plot.

strong Power Law relationship is observed with a 0.96 R² (96% of the data is described by the model). The same type of relationship is observed when the data is analyzed for individual utilities, by failure mode (e.g. third party damage, corrosion incidents, etc.). What this figure shows is that the majority of incident damage arises from a small number of incidents, as is typical for Power Law behavior.




Figure 1: Power Law Relationship for PHMSA Reported Gas Distribution Incidents

Note: Data plotted for reported pipeline incidents >\$100 k damage for 1992-2011 PHMSA incident statistics⁸

Figure 2 provides the same plot for the PHMSA reported data for Gas Transmission incidents based on the data from 2002 to 2011. Again, strong Power Law behavior is observed, with an R² of 0.97 (97% of the data is described by the model).




Figure 2: Power Law Relationship for PHMSA Reported Gas Transmission Incidents

Note: Data plotted for reported pipeline incidents >\$100 k damage for 2002-2011 PHMSA incident statistics⁸

Figure 3 provides the same plot for the PHMSA reported data for Hazardous Liquid Pipeline incidents based on data from 2002 to 2011. Yet again, strong Power Law behavior is observed, with an R² of 0.97 (97% of the data is described by the model). Figure 4 provides the data for gas gathering pipelines, with an R² of 0.95.

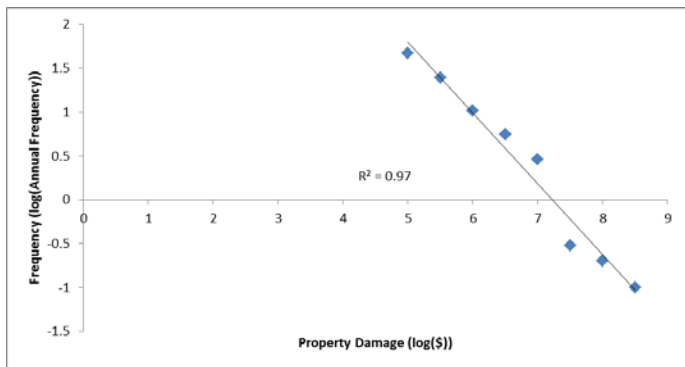


Figure 3: Power Law Relationship for PHMSA Reported Hazardous Liquid Pipeline Incidents

Note: Data plotted for reported pipeline incidents >\$100 k damage for 2002-2011 PHMSA incident statistics⁸




Figure 4: Power Law Relationship for PHMSA Reported Gas Gathering Pipeline Incidents

Note: Data plotted for reported pipeline incidents >\$100 k damage for 2002-2011 PHMSA incident statistics⁸

For four different pipeline industries: gas distribution, gas transmission, hazardous liquids and gas gathering, the same Power Law nature is observed for the distribution of incident size (measured in terms of PHMSA reported property damage) versus incident frequency. Similar Power Law behavior is observed for the distributions of number of injuries or fatalities versus frequency.

The Nature of Power Law Distributions

Power Law distributions have a unique form that differs significantly from the normal (or Gaussian) distributions that we are more accustomed to dealing with in statistical analysis. A classic example of a normal distribution is the variation in the height of women or men. As shown in **Figure 5**, the distribution of heights of North American men is normally distributed with a mean (or average) of just under 70". The majority fall between 65" and 74", and a few hit the extreme tails around 62" and 78", but there is a very low probability of someone falling outside this range. The ratio of these extremes, the tall end of the range divided by the short end of the range, is 1.3 (the ratio for the tallest and shortest men on record is around

5). The variation is uniformly distributed relatively closely around the mean.




Figure 5: Normal Distribution – Men’s Heights in North America

In contrast, the power law distribution for gas distribution incidents is shown in **Figure 6**. Those for gas transmission, hazardous liquids and gas gathering follow the same general form. Instead of being symmetrically distributed around the mean value, as observed in the normal distribution, there is a long tail to the distribution. It is this long tail that represents the low probability-high consequence events. The ratio of the largest to the smallest (the low end of the distribution is cut off at \$100,000 due to the artificial cut-off for reporting of incidents)⁶ is 430, indicative of the very broad range in potential consequences. The low probability-high consequence events dominate the risk picture – the top 20% of incidents is responsible for 60% of the property damage. The top 1% of incidents is responsible for 20% of the property damage.




Figure 6: Power Law Distribution for PHMSA Reported Gas Distribution Pipeline Incidents

Another key concept of distributions in general is that they describe a common population with common underlying drivers. Events that fall outside the distribution are outliers, impacted by factors other than those giving rise to the

⁶ The PHMSA cut-off is \$50,000 in property damage. The incidents in the \$50k to \$99k range were excluded from the analysis as they do not fit the power law model for the rest of the distribution. This is likely a consequence of the reporting process.

distribution. Once we define a distribution, therefore, we can use it to assess if a given incident is consistent with that distribution and hence part of the general population and driven by the same underlying factors or if the event is an outlier driven by other factors.

How Does Power Law Behavior Arise in Pipeline Incidents?

The mechanism underlying Power Law behavior in pipeline consequences is tied to the probability string that leads to a serious incident. Essentially, for a serious incident, there is a series of connected events that must occur, each with an associated probability: a leak, gas accumulation prior to location and repair, ignition, the presence of receptors (i.e. property, people, etc.), etc. Although the probabilities of each step vary with the specific environment, these are all essentially random events (for example, someone coming home after a leak to find gas accumulation in the house is a random event). It is the string of essentially random events, their associated consequences and their associated probabilities that results in Power Law behavior. All consequences are the result of a series of events occurring, each with an associated probability. Generally, the more severe the consequence, the longer the series of events that must occur, and mathematically the smaller the probability of that series occurring.

A simple analogy can be found in lotteries, which, whether we play them or not, we are all generally familiar with. While a lottery is what we could call a contrived⁷, or human-made system, it does provide a means of visualizing the Power Law nature of a string of probabilistic events. If we look, for example, at a lottery with six (6) numbers being drawn from a possible 49 numbers, we have a probability string leading to different outcomes as shown in **Figure 7**. To win, you need to have a ticket that matches a given number of the numbers drawn – the more numbers matched the greater the prize or consequence. There is a given probability for each step in the series of events⁸: p(1) for getting one number, p(2) for getting two numbers, p(3) for getting three numbers, etc. If we get only one number right, we get consequence 1 (c(1)) – in the lottery example, nothing. If we get two numbers right, we get c(2)... three numbers right, c(3)... etc... all the way up to the big consequence, six numbers right. The probability of getting one

number right is 0.12 (6/49)⁹, or roughly 1 in 8. As we go through the probability string, the probabilities decrease and the consequences increase. The probability of getting two numbers right is the product of the probability of getting one number right (6/49) times the probability of getting a second number right (5/48)¹⁰, or 0.013 (roughly 1 in 78). The probability of making it all the way through the probability chain – to getting all six numbers right and getting the big prize – is the product of the probabilities of getting one number times the probability of getting a second number times the probability of getting a third number, etc., and is very low. This is a low probability-high consequence event (the probability is 1 in 13,938,816).




Figure 7: Probability String for Lottery with Six Numbers Drawn

For a single event, we will have a single outcome – we will win the lottery (very unlikely), some money (a little more likely) or nothing (most likely). When we take the collection of all lottery players, we will have a distribution of outcomes that includes some winners (the minority) and some losers (the majority). It is in this distribution of the collection of outcomes that we see Power Law behavior. If we take the results from an actual lottery drawing¹¹, the number of winners in each step of the probability string and their consequences (or winnings), we see that this distribution does indeed follow a Power Law relationship. The data for an actual draw from this lottery are provided in **Figure 8**. The data fit the power law with an R² of 0.95.

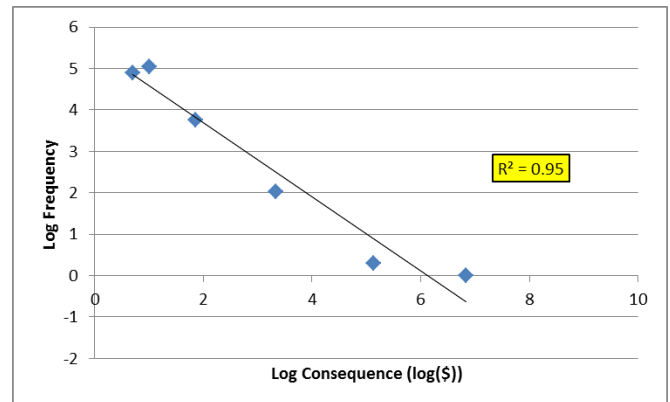


Figure 8: Power Law Behavior of Lottery with Six Numbers Drawn

⁷ Contrived in the sense that the prize money or consequences are set as a percentage of the overall pool of winning for each potential winning combination.

⁸ For this example, the lottery has 49 possible numbers, each number can only be drawn once and six numbers are drawn. The probabilities are, therefore: 6/49 for the first number, 5/48 for the second number, 4/47 for the third number, 3/46 for the fourth number, 2/45 for the fifth number and 1/44 for the sixth number. The overall odds for getting all six numbers is the product of these probabilities: (6/49)*(5/48)*(4/47)*(3/46)*(2/45)*(1/44) = 1 in 13,938,816

⁹ We have six chances (since we pick six numbers) out of 49 possible numbers for getting one number right, or a probability of 6/49 = 0.12

¹⁰ For the second number, we have five chances (five of the original six choices are left) out of a possible 48 numbers (since one is already gone), or a probability of 5/48 = 0.104

¹¹ Lotto 649 – August 10th, 2013 results

While this is a simple example, it provides some insight into how Power Law behavior emerges when we have a collection of events occurring where the consequences of each event follows a probability string.

If the probabilities along the probability string are similar for a group of events or incidents then that group will be defined by the same distribution. If the probabilities for a given incident are much lower or much higher than the general population, then the incident will be an outlier and fall outside the distribution for the general population. Applying this to San Bruno, if San Bruno falls within the general population of pipeline incidents, the probabilities along the probability string are consistent with the probabilities for the general population of pipeline incidents. San Bruno would then be reflective of the general population, suggesting the lessons from San Bruno are generally applicable to the overall industry. If San Bruno is an outlier then the probabilities along the probability string are different than the general population, suggesting San Bruno was driven by factors inconsistent with the general population of pipeline incidents.

San Bruno

Power Law analysis was used to assess the likelihood that an event the size of San Bruno would have been expected to occur somewhere in the US Gas Transmission industry based on historical industry performance. The Power Law distribution of property damage was developed by taking the historical data for gas transmission incidents in the PHMSA database reported prior to San Bruno (Figure x). There is an excellent fit to the data, showing a clear historical industry-wide relationship between the frequency of events and their size.

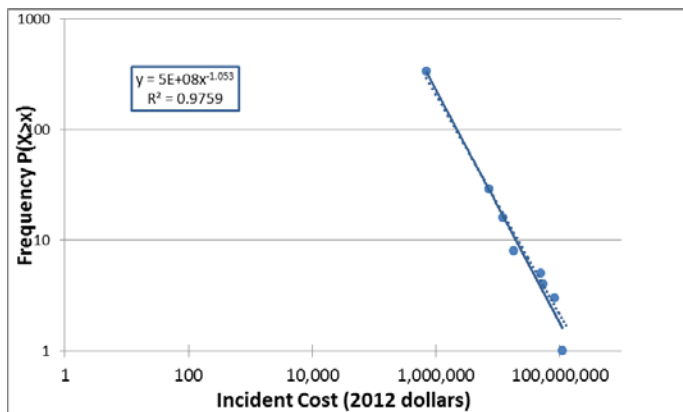


Figure 9: Power Law Model of PHMSA Reported Gas Transmission Incidents

This relationship can be used to predict the probability that we will see incidents of a certain size within a given time period. If there are no significant changes to the infrastructure or its management, using this past behavior should provide a reasonable projection.

The probability that a San Bruno magnitude incident would occur at any given utility is extremely low. When you look at all Gas Transmission utilities collectively, however, it is projected that an event the magnitude of San Bruno would be expected to occur roughly once every 40 years. There is a 96% probability that an incident the magnitude of San Bruno or greater will occur in the next 20 years (provided there are no significant changes to the infrastructure or its management). Figure x shows San Bruno on Power Law plot of incident frequency vs magnitude. It is clearly not an outlier; it is consistent with the historical data. An event the magnitude of a San Bruno, therefore, is not something unexpected or inconsistent with the historical performance of the overall gas transmission industry. Statistically speaking, it was just a matter of where and when an event of this magnitude would occur.

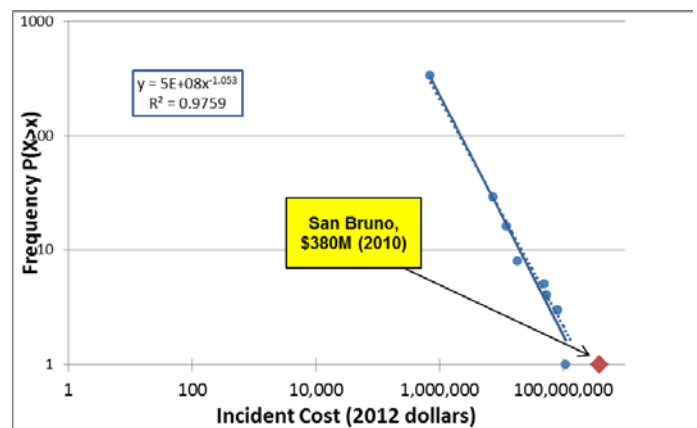


Figure 10: Power Law Model of PHMSA Reported Gas Transmission Incidents – San Bruno

In simple statistical terms, events that fall within the same statistical distribution have the same underlying drivers – events that fall outside a common statistical distribution do not. If we have a reliable statistical distribution for the height of people we know that if we measure someone’s height, it will fall within that distribution. If we measure the height of a rabbit, it will not – it is not part of the same population and does not have the same underlying drivers. The fact that San Bruno falls within the historical distribution of US pipelines incidents suggests that it is part of that same distribution and has the same underlying drivers. This suggests that the San Bruno incident provides lessons relevant to the industry as whole.

Conclusions

The consequences of pipeline incidents are seen to follow Power Law or Pareto distributions. That is, there is a direct relationship between the frequency of incidents and their size, and the distribution of incidents has a form that leads to the low probability-high consequence events dominating the risk picture. The San Bruno incident is seen to fall within the

distribution of historical gas transmission pipeline incidents, suggesting that it was not an outlier event and that the lessons from San Bruno are generally applicable to the overall industry.