Are U.S.A. School Shootings Black Swans?

By Ted G. Lewis

Abstract

This article quantitatively analyzes risk of death by school shootings in the USA using data from 1900 to 2023 for the purpose of determining if school shootings are black swans. Additionally, the risk assessment method demonstrated here is general enough to be used in other extreme incidents characterized by a power law distribution. It shows that some USA shootings qualify as potentially black swans. Two of the largest incidents are shown to be black swans, and two are not. The "black swan-ness" of an incident is shown to be relative to a history of incidents as captured by a power law with fractal dimension q. Incidents with $1 \leq q \leq 2$ are capable of being black swans. I use a test proposed by Giordana et al. to decide which of the largest shootings (by number of deaths) qualify as black swans. The result of statistical risk analysis is that two major shooting incidents qualify as black swans: the Virginia Tech shooting in 2007, and the Sandy Hook shooting in 2012, see Table I. Why is it important to distinguish events as black swans? The short answer is that risk analysis is useful for resource allocation. We want to allocate resources where they will do the greatest good. In a risk-informed strategy special consideration is given to incidents prone to black swans – they are resourced more than others. The purpose of this article, then, is to illustrate the implementation of an optimal resource strategy when applied to school shootings.

Suggested Citation

Lewis, Ted G. "Are U.S.A. School Shootings Black Swans?" *Homeland Security Affairs* 19, Article 5 (December 2023) www.hsaj.org/articles22539.

Introduction

The risk of dying in a school shooting is rather small compared with dying in a traffic accident on the way to school. None the less, we need to examine the risk and ask, "are school shootings black swans?" Black swans were qualitatively defined by Nassim N. Taleb as,

> …an event with the following three attributes. First, it is an *outlier*, as it lies outside the realm of regular expectations, because nothing in the past can convincingly point to its possibility. Second, it carries an *extreme impact*…. Third, in spite of its outlier status, human nature makes us concoct *explanations for its occurrence after the fact*, making it explainable and predictable. [1]

Taleb identifies two key attributes of black swans – size and surprise. Black swans are big and unexpected.

Subsequently, more rigorous definitions based on statistical analysis have quantified the meaning of grey and black swans. A grey swan is *an event that is possible and known, potentially extremely significant but* considered not very likely to happen.[2] The frequency of incidents versus consequence (size) must obey a long-tailed distribution such as a power law.[3] The data plotted as a power law must contain an "outlier" point (the black swan) that indicates an unexpectedly large deviation from the power law itself. In addition, the severity of an incident, relative to the "norm" established by the power law, places the black swan point above the power law curve, i.e., the likelihood of an event is greater than predicted by the power law. This leads to "shades of grey" and "black" following Taleb's initial coloring of swans. If school shootings are not black swans, might they be grey swans?

More formally, Giordana et al. define grey and black swans as incidents falling "far above" the straight line obtained by fitting to a log-log plot of frequency of incidents.[4] They define "blackness β", indicating the occurrence of a black swan, or not, due to a low-probability, high consequence data point in the data analyzed by fitting consequences to a power law. Black swans appear in the long-tailed power law distribution as points above the power law curve at the extreme right end, or tail. White swans are low consequence events lying on the curve, gray swans are higher consequence lying slightly above the curve, and black swans are extreme consequence events lying far to the right above the curve. Giordana et al. propose a metric for determining if an incident qualifies as a black swan.

Giordana et al. claim that only four incidents – the Turkish Airlines disaster, Lionel Messi's performance in soccer, the 9/11 terrorist attacks, and World War I deaths – qualify as black swan events based on their size and fit to a power law. The Turkish Airlines disaster where 583 passengers died in 1977, the number of soccer field goals scored by Lionel Messi (806), casualties in WWI (~40 million), and the 9/11 attack on America in 2001 (2,996) all qualify as black swans according to Giordana et al.

But, not all catastrophic incidents qualify as black swans. WW2 and the 1987 Black Monday stock market crash do not qualify because they fall on the power law distribution where expected, i.e. they fall on the power law curve. Both size and expectation of size must be considered when categorizing incidents as black swans. Size alone is insufficient to define an incident as a black swan. Clearly, the size of deaths in WW2 and the stock market crash of 1987 merit attention as a disaster, but not necessarily an unexpected disaster.

Given the wide disparity in sizes of events in terms of consequence, and the fact that the incidents analyzed by Giordana et al obey a power law distribution, how are black swans different? According to Giordana et al. the difference lies in the closeness of fit (or non-fit) of a black swan represented as a point on the power law graph, see Figure 2. Black swan points depart from the power law by being unexpectedly large and extremely low probability.

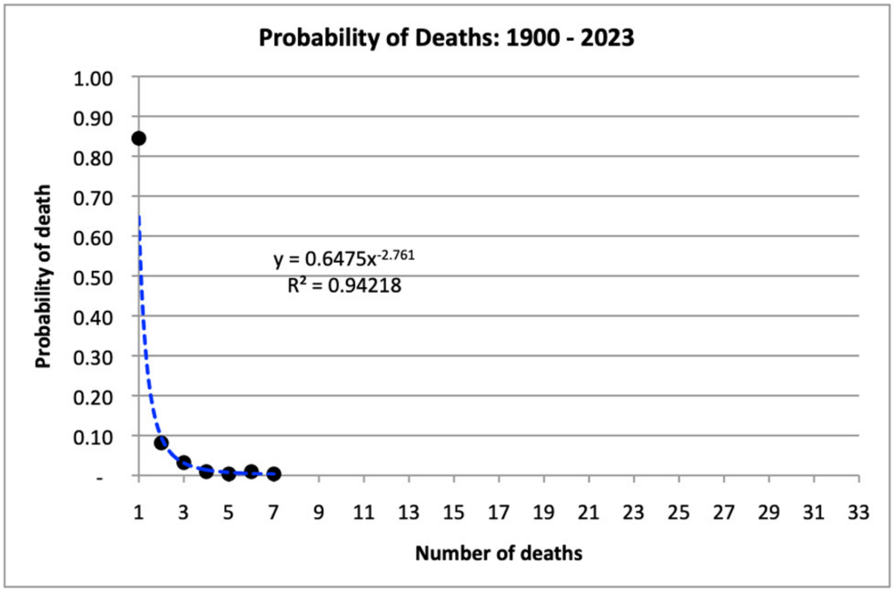

Figure 1: Frequency of deaths.

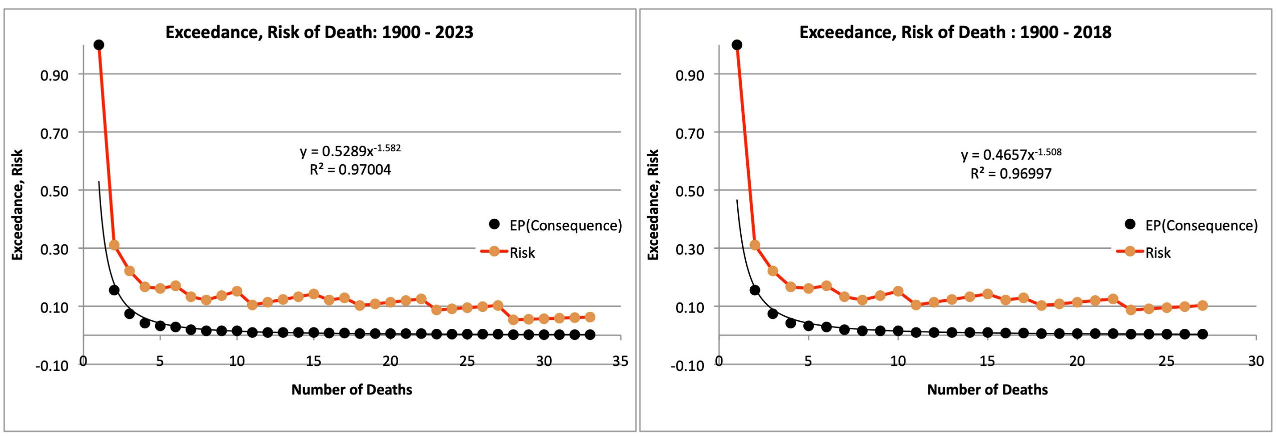

Figure 2: Exceedance, Risk of Death

(b). Left graph: exceedance and risk based on data including the Virginial Tech Shooting. Right graph: exceedance and risk based on the data excluding Virginial Tech Shooting but including the next-worst incident resulting in 27 deaths.

Analysis of the time series data shows a power exceedance in 2023 of q = 1.582, and in 2018 of q = 1.508.

Might the killing of 33 people in the 2007 Virginia Tech incident qualify as a black swan?

> A 23-year-old student, Seung-Hui Cho, armed with two pistols, a Glock 19 and a Walther P22, killed thirty-two students and faculty members at Virginia Tech and wounded another seventeen students and faculty members in two separate attacks before committing suicide. The first attack was on the second floor of West Ambler Johnston Hall, where he shot a young girl and another student who came to help. He

then changed his clothes, recorded multiple videos stating that he 'had to do it for his brothers and sisters', and made several pictures of himself with hammers, large knives, and guns. He then chained the doors of Norris Hall shut and shot and killed thirty more students and staff. He also shot seventeen more students and staff, wounding them. Six more students were also wounded from non-gunshot injuries. The incident is the deadliest school shooting, in modern U.S. history.[5]

The aim of this article is to apply a rigorous analysis of the data to determine the level of "black swan-ness" posed by school shootings. This analysis shows that school shootings in the U.S. are classical long-tailed incidents that are grey swans at most, but do not qualify as black swans according to the following analysis. Except for the Virginia Tech incident, analyzed below, the data simply fall too close to a perfect power law versus falling above the power law curve, meaning that most school shootings do not qualify as black swans. But, how can we tell?

Much has been written about school shootings, but as far as the author knows, a rigorous analysis of the data has never been performed and results quantitatively analyzed. Additionally, the following analysis may provide practitioners with a method applicable to other incidents of a stochastic nature, such as terrorism, industrial accidents, and criminality. The author does not want to claim too much generality, so the reader must decide how applicable this analytical technique is to their domain.

We first define exceedance and risk in terms of deaths, and then model the exceedance probability by fitting the data to a power law. This is easily done in Excel and reveals an important parameter – the slope of the data expressed by a graph: exceedance versus consequence. As a byproduct, risk can also be determined from the graph.

This relatively simple process reveals the possibility of black swans because the power law has a long tail, and the slope of the power law may indicate the potential of a black swan occurring. We apply the method of determining the "blackness β" of Giordana et al. to the U.S. school shootings data from 1900 to 2023 resulting in a rating of "black swan" or less. The worst school shooting incident in America occurred in 2007 resulting in 33 deaths. However, does this indicate that a black swan has occurred? What do the data say about the potential for black swans?

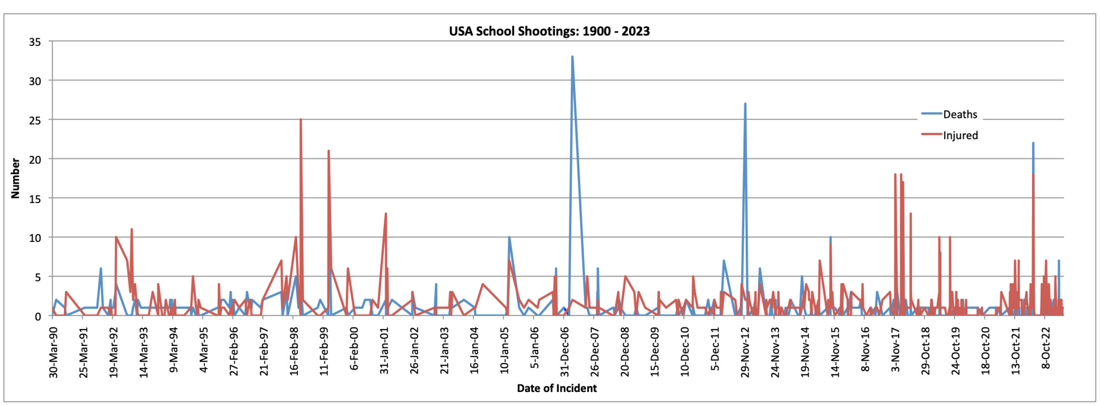

Figure 3: Number of deaths and injuries $x_i, i = 0, 1, \dots n.$ in USA school shootings since 1900.

Data and Methods

The data used here are from Wikipedia and shown graphically in Figure 1.[6] They consist of n = 724 incidents spanning 12,111 days with incidents separated by an average of 60 days and a maximum of 1611 days apart. The question we ask here is, "what is the risk of being murdered in a school shooting in the USA?" Additionally, "was the 2007 incident that resulted in 33 deaths, a black swan?" We apply the test suggested by Giordana et al, to the worst incident recorded between 1900 and 2024.

In the following, x_i designates the consequence (deaths) due to the i[th] incident. The maximum number of deaths occurred in 2007 resulting in 33 deaths. Hence, $1 \leq x_i \leq 33$. The *frequency* of incidents $f(x_i)$ is simply the normalized number of times incident x_i occurred in the time series. That is,

$$f(x_i) = \frac{number\ of\ deaths\ of\ size\ x_i}{total\ number\ of\ deaths}$$

$$\sum_{i=1}^{i=n} f(x_i) = 1$$

Frequency may also be used as a measure of *vulnerability* assuming vulnerability is equated with likelihood of dying during an incident. In this case, $\Pr(x_i) = f(x_i)$. Figure 2(a) shows probability of dying is also a power law with rapidly declining slope versus number of deaths per incident.

Normalized frequencies are converted into *exceedance* EP – the tail sum over frequency,

$$EP(x_i) = \sum_{j=i}^{j=n} f(x_j)\,; n \leq i < 0$$

Note that,

$$EP(x_0) = \sum_{j=1}^{j=n} f(x_j) = 1$$

Alternatively,

$$EP(x_n) = f(x_n)$$

$$EP(x_{n-i}) = EP(x_{n-i+1}) + f(x_{n-i}); i = 1, 2, \dots n$$

Lewis defines *risk* as the expected exceedance value obtained by multiplying x-axis and y-axis to obtain R(x).[7], [8]

$$R(x_i) = x_i EP(x_i);$$

This is further validation of "black swan-ness": if a black swan exists in the data, $R(x_i)$ will show an upturn in value versus x_i.

We find that in most cases, exceedance EP fits a long-tailed power law.[9],[10] Mathematically, power law distributions are a function of inverse consequence x, raised to an exponent of the form,

$$EP \sim A/x_i^q$$

Where slope q is obtained by fitting an optimal least squares (OLS) straight line to the logarithmic form of the power law;

$$\log(EP) = \log(A) - q log(x_i)$$

Here, A and q are obtained by numerical estimation of the OLS line.[11] That is, we simplify the analysis by capturing exceedance and risk in the form of an analytic formula with a single parameter, q – the slope of the straight line – called the *fractal dimension.* Fractal dimension determines the nature of risk as follows:

If $2 < q \leq 3$ EP has no mean or variance, hence predictions are unknown.
If $1 \leq q \leq 2$ EP has a variance, but no mean, hence incidents are potentially black swans.[12]
If $q < 1$ EP has a mean value and variance, hence mean (and variance) can be used as predictors.

In this case, q = 1.58, hence school shootings fall into the category of unknown mean value. This suggests that prediction of future deaths due to school shootings is uncertain. Moreover, it indicates the potential for the process to produce a black swan. We show that the Virginia Tech shooting was a black swan incident according to the metric proposed by Giordana, thus confirming Newman's claim that $1 \leq q \leq 2$ indicates the possibility of a black swan.

Analysis

Risk analysis of the data in Figure 1 is carried out using the following methodology, see Figure 2b:

- Calculate normalized frequencies, $f(x_i);\ x_i \in [1, 33]$.
- Calculate EP and R; plotting them on linear and log-log charts to obtain slope q.
- Categorize the results in terms of risk:
    - What does q tell us about risk?
    - Is school shooting a white, grey, or black swan event?
- Apply Giordana's test to confirm that a black swan occurred in 2007.

Figure 2 shows a near perfect fit of the data to a power law with fractal dimension of 1.58 and R-squared of 0.97. This value of q indicates the potential for a black swan. Next, we analyze the most extreme incident on the Virginia Tech campus in 2007. Before the Virginia Tech shooting, the largest consequence in modern times was the Columbine High School shooting with consequence of 15 deaths on April 20, 1999. Comparing these two data sets (Virginia Tech vs. Columbine) one can determine if the Virginia Tech attack was a black swan. Subsequently, additional "large and unexpected" attacks occurred on the dates with the associated deaths as listed in Table I.

Are the major incidents listed in Table I all black swans? Only if each subsequent incident exceeded the previous incidents by a "an unexpectedly large size". Giordana applies a "before and after" test to determine "unexpectedly large size", as follows. Compare the power law OLS curve of the data with the black swan candidate versus the power law curve and OLS fit including the potential black swan. *If the ratio exceeds 1.0, the incident is a black swan*. That is,

$$\beta = \frac{\frac{A}{x_b^{q}} - 1}{\frac{A_0}{x_b^{q_0}} - 1} \ where\ x_b = potential\ black\ swan\ consequence$$

$\beta > 1$ implies black swan.

A, A_0, and q, q_0 are the before and after power law parameters, respectively, determined by the data recorded up to the time of the incident under consideration, see Table I. Plugging in the numbers from OLS curve fitting for 33 and using data up to 2023 versus data up to 2007;

$$\beta = \frac{\frac{.529}{33^{1.578}} - 1}{\frac{.4657}{33^{1.508}} - 1} = 1.00027 > 1$$

Thus, the Virginia Tech incident was a black swan. What about the others? Table I summarizes the calculations of β for each of the largest incidents. Each incident is compared with the previous incident data to determine β. Parameters A_0 and q_0 are obtained using data up to the time of the incident considered as a black swan. Black swans are relative to previous incidents, as determined by their power law parameters.

Table 1: Most consequential US shootings and their risk parameters.

When/Where the attack occurred	Number of deaths	A	q	β
4/16/07 (Virginia Tech)	33	.529	1.578	1.00027
12/14/12 (Sandy Hook)	27	.4657	1.508	1.00038
2/14/18 (Parkland, FL.)	17	.5267	1.592	0.99954
5/24/22 (Uvalde, Texas)	22	.4778	1.524	0.99948

According to this analysis, U.S. shootings are prone to black swans ($1 \leq q \leq 2$), and some qualify as black swans while others do not. We can determine which incidents qualify as black swans using the analysis presented here. A black swan is an incident in which the consequence time exceedance is significantly greater than expected per its power law.

Conclusion

Quantitative risk analysis shows some school shootings in America qualify as black swans and others do not. It depends on the fractal dimension of their distribution. This does not reduce the seriousness or demean the severity of non-black swan school shootings. It simply means that they do not technically qualify as black swans.

Policy makers should consider balancing prevention and response using quantitative risk assessments like the one used here. Of course, *response* is essential, but how many resources should be applied to *prevent* catastrophic incidents like school shootings? This implies that incidents that are prone to black swans like school shootings should receive additional resources aimed at preventing worse-case consequences.

This technique is general and can be applied to most any historical data to assess risk. Given a time series of incidents and their consequences, construct the exceedance probability as illustrated above, then using Excel, find the power law slope q. Figure 3 illustrates this for major floods over the past 1.8 million years. While the incidents, consequences, and exceedance are different, the technique is the same.

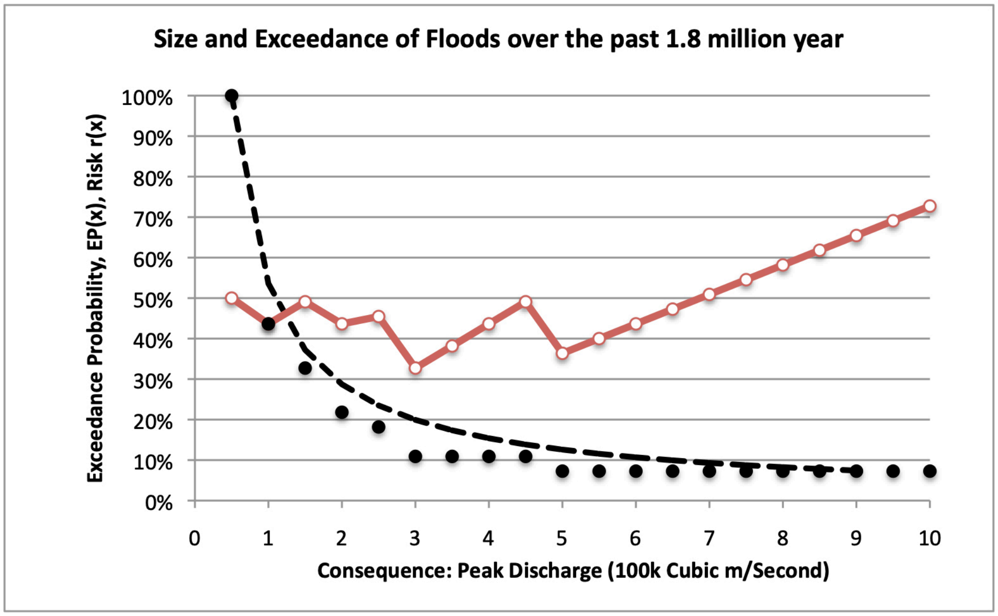

Figure 3. Illustration of risk assessment applied to large historical floods over the past 1.8 million years on earth. Note the upswing in risk as consequence increases. This suggests impending black swans in the past and likely to occur in the future.[13]

About the Author

Ted G. Lewis is a retired professor of computer science and former executive director of the Center for Homeland Defense and Security at the Naval Postgraduate School. He spent forty years in academic, industrial, and advisory capacities, ranging from academic appointments at the University of Missouri-Rolla, University of Louisiana, and Oregon State University, to senior vice president of Eastman Kodak Company, to CEO and president of DaimlerChrysler Research and Technology, North America. Dr. Lewis has published over thirty books and 100 research papers. He is the author of *Critical Infrastructure Protection in Homeland Security: Defending a Networked Nation* (2006, second edition 2014), *Network Science: Theory and Applications* (2009), *Bak's Sand Pile* (2011), and *Book of Extremes* (2014). He received his PhD in computer science from Washington State University. Dr. Lewis may be contacted at tedglewis@icloud.com.

Bibliography

Cioffi-Revilla, Claudio, and Manus I. Midlarsky. Power Laws, Scaling, and Fractals in the Most Lethal International and Civil Wars. 2013. https://www.press.umich.edu/pdf/047211395X-ch1.pdf.

Giordano De Marzo et al. "Quantifying the unexpected: A scientific approach to Black Swans, PHYSICAL REVIEW RESEARCH 4, 033079 (2022).

Lewis, T.G. The Mathematics of Catastrophe. Applied Math 2022, 2, 480–500. https://doi.org/10.3390/appliedmath2030028.

Lewis, Ted G. The Many Faces of Resilience, Communications of the ACM, January 2023, Vol. 66 No. 1, Pages 56-61. 10.1145/3519262. https://cacm.acm.org/magazines/2023/1/267957-the-many-faces-of-resilience/fulltext.

Lewis, Ted G. Critical Infrastructure Resilience and Sustainability Reader. (John Wiley & Sons, 2023).

Newman, M. E. J. "Power laws, Pareto distributions and Zipf's law". Contemporary Physics. 46 (5) (2005): 323–351. arXiv:cond-mat/0412004. Bibcode:2005. ConPh..46..323N. doi:10.1080/00107510500052444. S2CID 202719165.

Taleb, Nassim N. The Black Swan: The Impact of the Highly Improbable. (2010) Random House. ISBN 978-0-8129-7381-5.

Notes

1. Nassim N. Taleb, *The Black Swan: The Impact of the Highly Improbable,* (2010), Random House, ISBN 978-0-8129-7381-5.

2. https://www.investopedia.com/terms/g/grey-swan.asp .

3. Power laws are distributions that obey a power, q, as in $Pr(x) \sim A/x^q$.

4. Giordano De Marzo et al, "Quantifying the unexpected: A scientific approach to Black Swans**,** PHYSICAL REVIEW RESEARCH 4, 033079 (2022).

5. Description: https://en.wikipedia.org/wiki/List_of_school_shootings_in_the_United_States_(1900–2022) .

6. Ibid.

7. T.G. Lewis, "The Mathematics of Catastrophe," *AppliedMath*, 2022, 2, 480–500, https://doi.org/10.3390/appliedmath2030028 .

8. T.G. Lewis, "The Many Faces of Resilience," *Communications of the ACM*, January 2023, Vol. 66 No. 1, Pages 56-61. 10.1145/3519262. https://cacm.acm.org/magazines/2023/1/267957-the-many-faces-of-resilience/fulltext.

9. Claudio Cioffi-Revilla and Manus I. Midlarsky, Power Laws, Scaling, and Fractals in the Most Lethal International and Civil Wars,. https://www.press.umich.edu/pdf/047211395X-ch1.pdf.

10. Observed patterns in terrorist attacks, wealth and poverty in developing societies, political instability, foreign aid distributions, and aspects of domestic and international conflicts obey power laws.

11. OLS is optimal least-squares.

12. M.E.J Newman, *"Power laws, Pareto distributions and Zipf's law,"* Contemporary Physics. 46 (5) (2005): 323–351. arXiv:cond-mat/0412004. Bibcode:2005. ConPh..46..323N. doi:10.1080/00107510500052444. S2CID 202719165.

13. Ted G. Lewis, *Critical Infrastructure Resilience and Sustainability Reader.* (John Wiley & Sons, 2023).

Copyright