Heavy Tails and the Central Limit Theorem

Thomas B. Fowler, Sc.D.

Heavy-tailed distributions have become very important in fields as diverse as telecommunications and economics. They often occur in situations where one would expect that the Central Limit Theorem should apply. This paper investigates how the Central Limit Theorem fails, and shows one mechanism by which heavy-tailed behavior can arise—by the addition of what the paper defines as "hypercorrelated" random variables. Such hypercorrelation is necessary since heavy-tailed behavior cannot arise from the addition of linearly-related random variables.

Introduction

The Central Limit Theorem is one of the most profound and important results in mathematics. It tells us that under fairly general conditions, the net result of many actions or events can reliably be described by a normal distribution. Some common examples are:

- Height of male or female freshman students at a college (result of the combined action of many independent genetic and environmental factors)
- Time to make a certain long trip by automobile (result of weather, road conditions, accidents, and other independent factors)
- Electricity consumption in a city (result of a large number of independent consumers)

Other cases which would seem to fit (but, in fact, do not) are:

- Height of ocean waves (result of the sum of many different forces)
- Stock market prices (result of many trades by independent investors)
- Length of Internet messages (result of many different activities by disparate users)

The Central Limit Theorem is used to estimate the distribution of some quantity when it is the result of many small contributions from various sources. The requirements for the Central Limit Theorem to be applicable are: [1]

- Condition 1: Variables summed must be independent
- Condition 2: All variables must have finite mean and variance
- Condition 3: No variable can make an excessively large contribution to the sum

Despite its success in many instances, the Central Limit Theorem is known to fail as a descriptive tool in cases where its applicability seems assured. In many of these cases, we find that we are dealing with heavy-tailed distributions, and decisions made on the basis of normal distribution characteristics often lead to catastrophic results. Perhaps the most famous recent example is the meltdown of a hedge fund, Long-Term Capital Management, in 1998. The fund assumed that certain currency movements would be governed by the normal distribution on account of the (presumed) applicability of the Central Limit Theorem. In reality, a set of movements occurred which, according to the fund managers, was a "10 sigma event," triggering huge leveraged losses which nearly destabilized the global financial system. [2] Of course, 10 sigma events do not occur, since their probability is 7.62 x 10^{-24}. Indeed, if the event in question is measured on a daily basis, which includes settlement of most market-based accounts, the event would occur once every 1.31 x 10^{23} days. This corresponds to about 3.6 x 10^{20} years—roughly 10 orders of magnitude longer than the age of universe. What happened was that a heavy-tailed distribution, not a normal distribution, governed the phenomenon.

Another case of great interest in the telecommunications arena is the distribution of the file size of Internet messages. Indeed, many parameters associated with Internet traffic are best described by heavy-tailed distributions, including page requests per site, reading time per page, and session duration. [3] Heavy-tailed distributions cause serious problems in the design of Internet Protocol (IP)-based systems, because the usual queuing theory methods employed to determine the required size of network components to meet performance specifications break down when variance and higher-order moments cannot be calculated. It has been necessary to develop special techniques, such as the Transform Approximation Method (TAM), in order to cope with this problem. [4, 5] The TAM utilizes the fact that one does not need actual infinite times and infinite variances, but can utilize finite approximations to get useful results.

What goes wrong, and why? Which of the three conditions is the problem? How can heavy tails emerge from sums of small-contribution random variables? We must examine conditions 1 through 3, given above, to determine which is likely to fail in practical cases, and how such a failure would manifest itself. Clearly, condition 2 is not of much interest because if the mean and variance of one or more components do not exist, we already have a heavy tail for all intents and purposes. Nor is condition 3 likely to be of great significance, because we are interested in the case where heavy tails emerge from the sum of many small contributions. Therefore, condition 1 would seem to be an ideal candidate for investigation, because correlated variables may be able to exhibit large tails. Presumably, positive feedback mechanisms of some type give rise to the larger-than-expected probabilities of events. However, it is not obvious that this is the cause, because adding random variables with the distributions that are just multiples of each other (100 percent correlated) yields a similar distribution with different mean and variance, but not a heavy tail. Thus more is needed than simple correlation. Specifically, what is needed is some type of nonlinear feedback mechanism which can yield large changes. Therefore the relationship among the variables being summed would include this kind of nonlinear feedback, and that would ultimately generate the heavy tails.

Heavy-Tailed Distribution and Infinite Variance

The size of files sent over the Internet has been determined to have what is known as a "heavy-tailed" probability distribution; [6] that is, the probability of a given file length occurring falls off very slowly with increasing file length, unlike ordinary distributions, where the rate of fall-off is much faster. Technically, given a probability distribution function *f(x)*, if for large x values, its cumulative distribution function $F(x)$ has the property that its complementary distribution

$$1 - F(x) \approx \kappa_1 x^{-\beta}$$

where $\kappa_1 > 0$ and $\beta \in [0,2]$, then the distribution function is said to be *heavy-tailed* because it falls off very slowly with increasing values of x. [7, 8] In turn, this has an important consequence. Setting $\beta = 2$ and differentiating the above equation,

$$f(x) = 2\kappa_1 x^{-3}.$$

Since $Var(x) = E(x^2) - [E(x)]^2$ and $[E(x)]^2$ is fixed, it follows that the variance is determined by

$$E(x^2) = \int_{-\infty}^{\infty} x^2 f(x) = \int_{-\infty}^{\infty} \frac{2\kappa_1}{x} dx = \infty \cdot$$

That is, the variance is infinite.

The physical significance of infinite variance can be readily understood. Recalling that variance measures central tendency, for any *finite* variance, values are known to be clustered around a central measure, the mean. Increasing the coarseness of the horizontal scale will result in the distribution appearing more "peaked," i.e., clustered around the mean. For example, if one plots a histogram of some normally distributed characteristics, and then smoothes the resulting curve, it will appear as shown in Figure 1 for various horizontal (abscissa) scale ranges. Note that the probability of large deviations from the center or mean is very small. However, in the case of *infinite* variance, there is no such clustering, and regardless of the scale on which measurements are made, there is no change in their central tendency—in effect, all scales look the same. [3] The histograms are illustrated in Figure 2. Physically, this means that it is difficult or impossible to put limits on the values of random variable that one may observe. Such values can become arbitrarily large in absolute value with a much higher frequency than is the case with better-behaved distributions such as the normal distribution.

The simplest heavy-tailed distribution (and the most famous in packet network analysis) is the Pareto distribution, [7] which has the general form

$$F(x) = 1 - \left(\frac{k}{x}\right)^{\alpha}, \quad \alpha, k > 0 \quad x \geq k$$

$$p(x) = \frac{dF(x)}{dx} = \alpha k^{\alpha} x^{-\alpha-1}$$

As α decreases, the "heavy-tail" effect increases. For $\alpha < 1$, the variance becomes infinite. This distribution is a type of negative power law distribution, similar to the one discussed above.

Certifying infinite variance in practical cases is very difficult. In many cases, infinite variance is a much stronger condition than necessary; often simply having a very large variance is enough to cause significant financial, queuing, or other problems. In such cases, the ratio of the standard deviation to the mean can become quite large. A common way to identify heavy tails is to look for a fall off of the density function that takes the form of a power law (straight line on a logarithmic plot). This can be easily recognized (see Figure 3). Note the approximately linear power law behavior of the Pareto distribution for larger values along the abscissa, and its much slower rate of decrease.

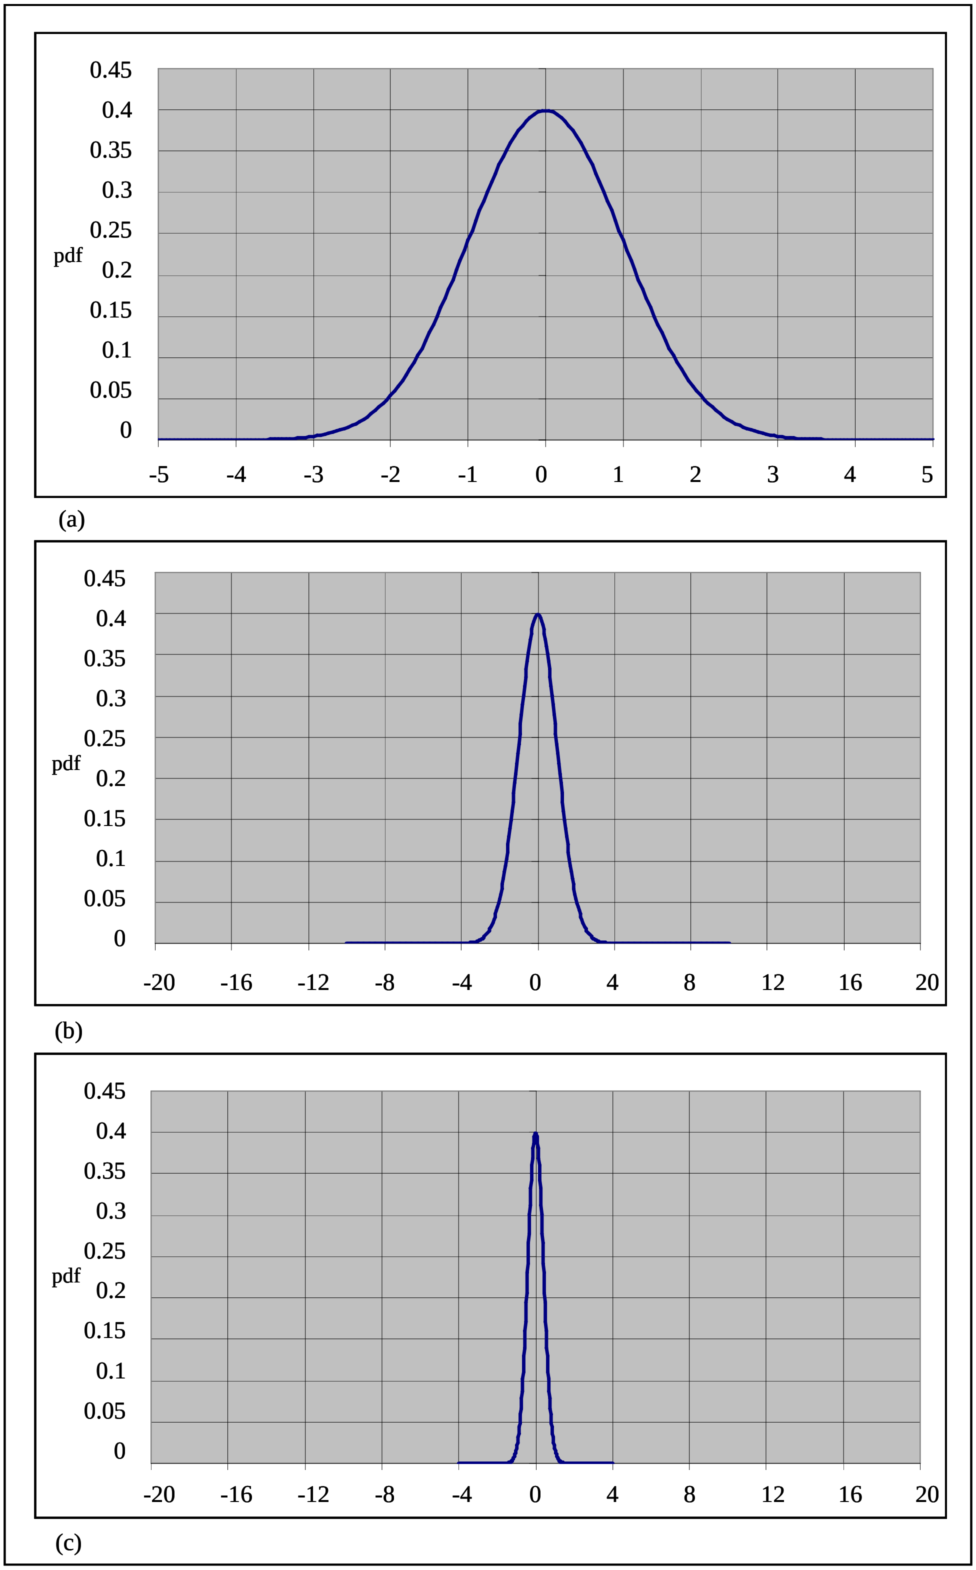

Figure 1. Appearance of Distribution Function of Random Variable With Well-Defined Variance, Illustrated on Different Abscissa Scale Ranges

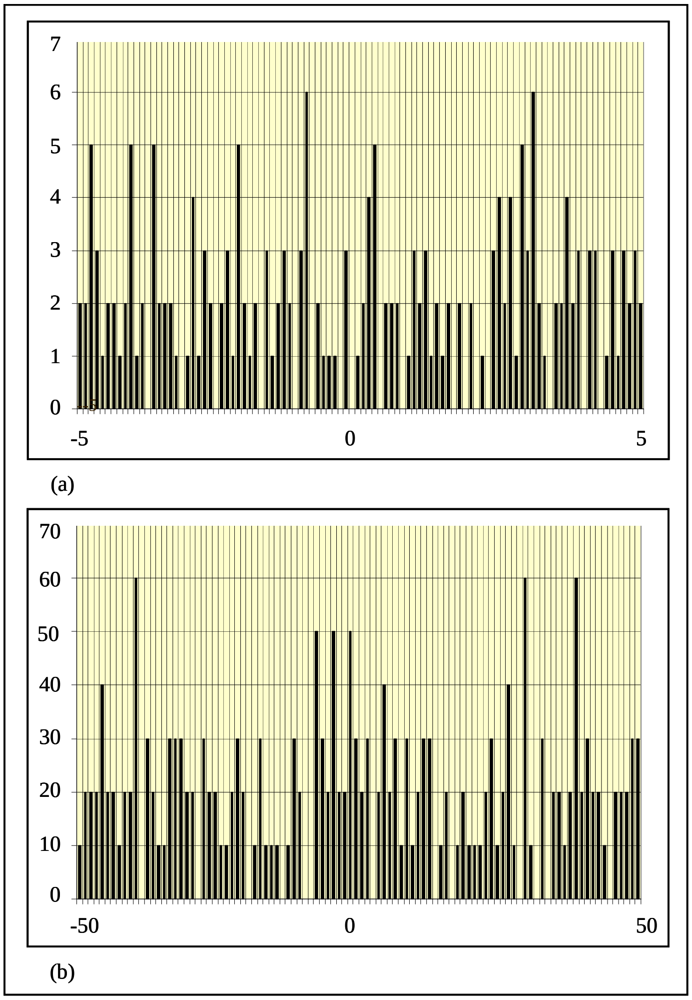

Figure 2. Histogram of Random Variable Without a Variance, Showing Same Pattern on Different Horizontal Scales

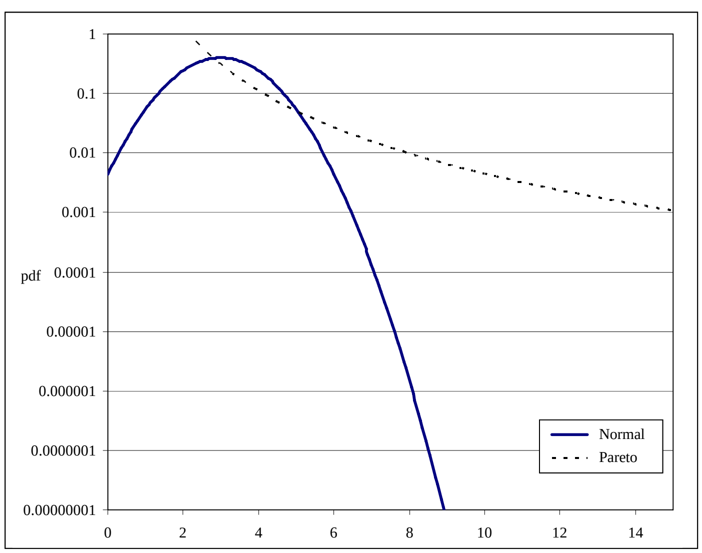

Figure 3. Comparison of Normal and Pareto Distributions With Respect to Fall Off for Large Values of X

Failure of the Central Limit Theorem

To understand how the failure of Central Limit Theorem manifests itself, and why more is required than simple correlation of random variables added, consider the following two cases.

1. Random variables *X*, *Y* are completely correlated, i.e., $\rho = 1$. In that case, $Y = aX + b$. A typical data set might be:

X	1	2	3	4
Y	2	4	6	8

Here $Y = 2X$.

2. Random variables *X*, *Y* which are independent, but which have similar distributions:

$X: (\mu, \sigma^2)$
$Y: (2\mu, 4\sigma^2)$

In this case, though mean values are related as $\mu_Y = 2\mu_X$, correlation $\rho = 0$.

Now observe what happens when these are added. Let $Z = X + Y$. Then for the two cases:

1. $\mu_Z = \mu_X + \mu_Y$ since expected value is always additive. Thus $\mu_Z = \mu_X + 2\mu_X = 3\mu_X$

For the variance we have

$$
\begin{aligned}
Z^2 &= (X+Y)^2 = X^2 + 2XY + Y^2 \\
E(Z^2) &= E(X^2) + E(2XY) + E(Y^2) \\
&= E(X^2) + E(2X \cdot 2X) + E([2X]^2) \\
&= E(X^2) + E(4X^2) + E(4X^2) \\
&= E(9X^2) = 9E(X^2)
\end{aligned}
$$

Then Var(Z) can be calculated as:

$$
\begin{aligned}
Var(Z) &= E(Z^2) - [E(Z)]^2 \\
&= 9E(X^2) - [3E(X)]^2 \\
&= 9E(X^2) - 9[E(X)]^2 \\
&= 9Var(X)
\end{aligned}
$$

2. $\mu_Z = \mu_X + \mu_Y$
$\sigma_Z^2 = \sigma_X^2 + \sigma_Y^2$

For the above case, this becomes

$$
\begin{aligned}
\mu_Z &= \mu_X + \mu_Y = \mu_X + 2\mu_X = 3\mu_X \\
\sigma_Z^2 &= \sigma_X^2 + \sigma_Y^2 = \sigma_X^2 + 4\sigma_X^2 = 5\sigma_X^2
\end{aligned}
$$

Note that in the calculation for variance of the correlated variables, the middle term $E(2XY)$ does not drop out as it would for independent variables, since in general for correlated variables $E(XY) \neq E(X)E(Y)$. In the case of independent variables we have

$$
\begin{aligned}
V(X+Y) &= E(X+Y)^2 - [E(X+Y)]^2 \\
&= E(X^2 + 2XY + Y^2) - [E(X) + E(Y)]^2 \\
&= E(X^2) + 2E(X)E(Y) + E(Y^2) - [E(X)]^2 - 2E(X)E(Y) - [E(Y)]^2 \\
&= E(X^2) - [E(X)]^2 + E(Y^2) - [E(Y)]^2 \\
&= V(X) + V(Y)
\end{aligned}
$$

In general case, for *n* random variables $X_1 \ldots X_n$, where $X_i = a_i X_1$, $a_i > 0$ for all a_i and, where, for all a_i and

$Z = \sum_{i=1}^{n} X_i$, we have

$$E(Z) = \sum_{i=1}^{n} E(X_i) = \sum_{i=1}^{n} E(a_i X_1) = \sum_{i=1}^{n} a_i E(X_1) = E(X_1) \sum_{i=1}^{n} a_i$$

$$V(Z) = \left(\sum_{i=1}^{n} a_i \right)^2 V(X_1) \text{ or } \sigma_Z = \sigma_{X_1} \sum_{i=1}^{n} a_i$$

In effect, this says that the *standard deviation scales at the same rate as the mean*. Thus the *shape* of the distribution is unchanged, and specifically, it does not exhibit the flattening required for a heavy tail.

But if the random variables are independent, *E(Z)* is the same, but

$$V^*(Z) = \sum_{i=1}^{n} a_i^2 V(X_1) < V(Z)$$

That is, standard deviation scales as $\sqrt{\sum_{i=1}^{n} a_i^2}$, which is *less* than the rate at which the mean scales. This implies that the shape of the distribution is being compressed—exactly the opposite behavior to that required for heavy tails. Naturally such a result is just what we would expect, since the Central Limit Theorem governs this case.

In general, comparing these, we have

$$\frac{V(Z)}{V^*(Z)} = \frac{\left(\sum_{i=1}^{n} a_i \right)^2}{\sum_{i=1}^{n} a_i^2} > 1$$

This is true for variables regardless of their distribution, so it shows that *heavy tails cannot arise by linear combinations of correlated random variables*. The new summed distribution has the same shape as the old, only scaled up. When converted to pdf, it will not have a heavy tail, defined as variance tending to ever larger values in the pdf. Obviously, adding independent random variables will not lead to heavy tails because the standard deviation scales at a smaller rate than the mean, so the distribution tends to contract. The two situations are illustrated in Figures 4 through 6. Note in Figure 4 that the correlated sum and X overlap.

Hypercorrelation

The fact that the variance of a sum of correlated random variables is greater than the corresponding sum for independent random variables suggests the next step. If normal correlation cannot create heavy tails, but only replicate the shape of the original distribution, perhaps some type of *amplified* correlation can do so. One way in which an "amplification" can occur in correlated distributions is if a power-law relationship exists for the variables. So we shall consider relationships of the form $Y = aX^n$, and we will define ρ' as the standard correlation coefficient ρ for the natural log of variables Y, X where the values of Y and X have been fitted using an exponential regression, so that a and n are known. That is, ρ' will measure the correlation for $\ln Y = \ln a + n \ln X$.

This is equivalent to finding ρ for W and Z where $W = a' + b'Z$. Then the *hypercorrelation* between X and Y is given by $\rho^* = n\rho'$. Thus if ρ' is 0, there is no hypercorrelation. Conversely, if ρ' is 1, there is hypercorrelation of n.

It is clear that n measures the strength of the amplification effect, i.e., how much faster Y increases than X. In real-world situations, we would want to assume a distribution for a and a distribution for n, and then determine under what circumstances the heavy-tailed behavior would arise. This could be done by fitting the tail of the summed hypercorrelated random variables to a Pareto distribution.

Specifically, let us assume that we have a large number of hypercorrelated variables Y_i, such that $Y_i = a_i X^{n_i}$, each of which has an associated hypercorrelation ρ_i^*. We should be able to show that if a large number of Y_i's are summed, $\sum_{i=1}^{k} \rho_i^* / k$ increases, and the ratio of the standard deviation to the mean continues to increase. Thus heavy tails emerge since the sum of the Y_i's is bounded on the lower end.

Simulation

Efforts to verify the creation of heavy-tailed distributions by means of hypercorrelation have thus far been

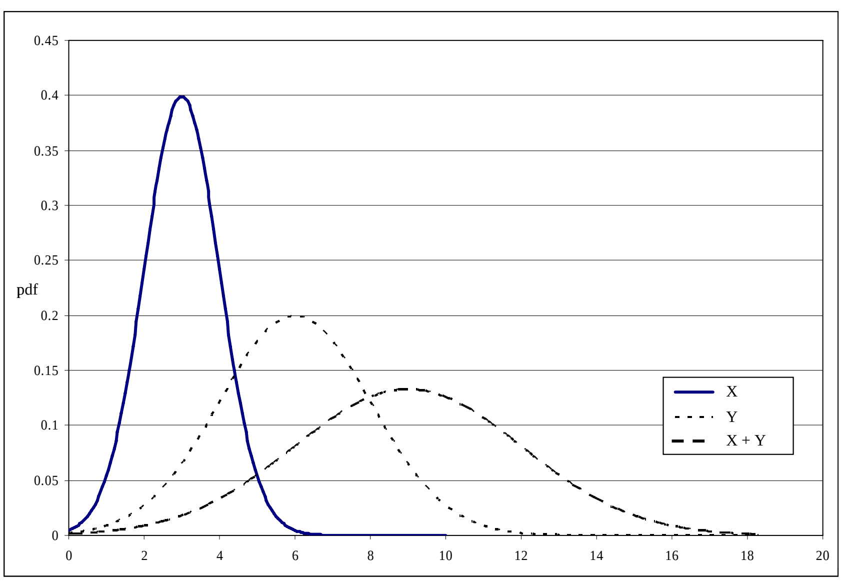

Figure 4. Two Correlated Random Variables and Their Sum

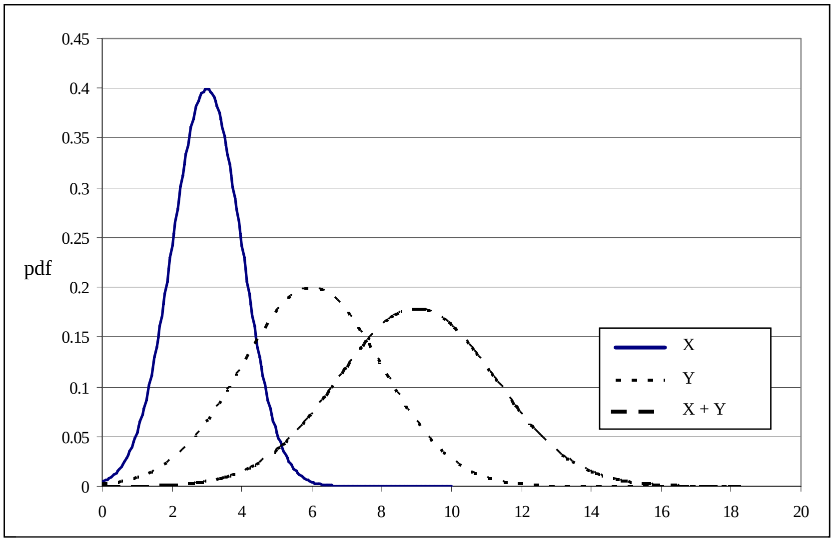

Figure 5. Two Independent Random Variables and Their Sum

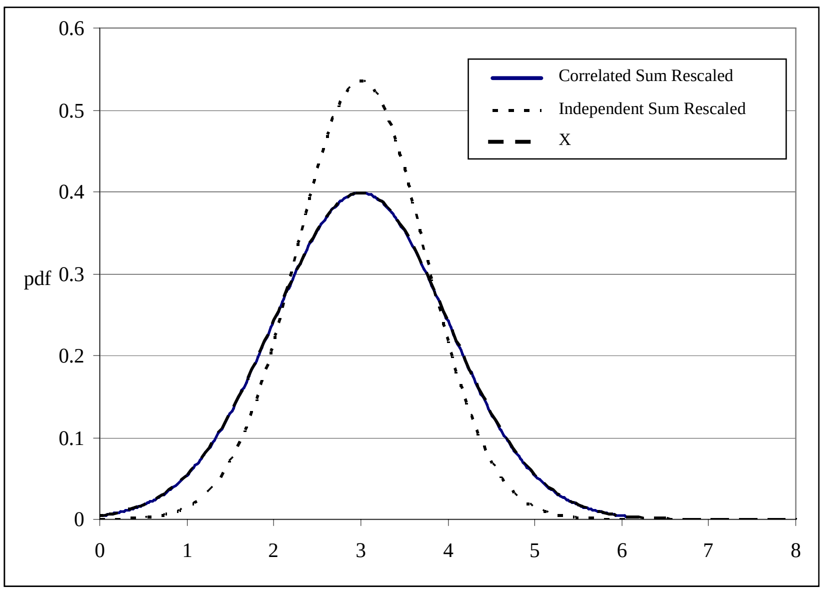

Figure 6. Rescaled Sums Compared to Original Random Variable X

done by means of numerical simulation. In these experiments, hypercorrelated random variables are generated, and then added. The resulting distributions are then analyzed for heavy-tailed behavior. Such behavior is quite obvious on logarithmic plots, especially when compared to the usual result of adding random variables, namely a normal distribution. The procedure is as follows:

1. Assume a functional form of $Y_{ij} = a_i X^{n_j}$
2. Generate 5 normally distributed random values for the a_i, and, for each of those, 10 normally distributed random values for the n_j
3. Generate a normally distributed set of values for X
4. Calculate the Y_{ij} terms and add, $Y_i = \sum_{j=1}^{10} Y_{ij}$
5. Calculate the distribution of the sum $Z = \sum_{i=1}^{5} Y_i$

6. Fit a normal and a Pareto distribution to the distribution for *Z*
7. Record Pareto parameters and average hypercorrelation
8. Repeat for new set of *a, n* values with different mean, standard deviation

Some typical graphs resulting from step 6 are shown in Figures 7 and 8 on both linear and log scales. These graphs illustrate clearly the heavy-tailed behavior of the summed hypercorrelated variables, as compared to a normal distribution. Note that in Figure 7(a), the normal and hypercorrelated distributions look almost identical until well past the mean; this could easily fool an observer into thinking that the process in question was really normally distributed when, in fact, it has a heavy tail. The difference, of course, is immediately obvious when the distribution is plotted on a logarithmic scale. A smaller value of *n*, which corresponds to less hypercorrelation and presumably less heavy-tailed behavior, gives correspondingly a higher value for the Pareto constant α. Note also in Figure 8 that the degree of heavy-tailed behavior is much lower than in Figure 7, as illustrated by the decreased divergence between the hypercorrelated (and Pareto) curves and the normal curve. This result is expected, since the exponent *n* in the runs of Figure 7 is 6, whereas it is 2 for the runs of Figure 8. As the exponent *n* approaches 1 from the upper side, the heavy-tailed behavior will disappear.

A summary plot of many runs, given in Figure 9, shows the relationship between average hypercorrelation value and Pareto constant α. As expected, the relationship is inverse but well-defined, as indicated by the regression line.

How Heavy Tails Might Arise Through Hypercorrelation

It is useful to consider how heavy tails due to hypercorrelation might arise. First, consider the case of investors, whose behavior is known to be subject to crowd influences ("if everyone is doing it, I should too!"). So Investor 1 says, "I'm buying 10 shares of company A." Investor 2 says, "I see what Investor 1 is doing, and I like it, so I'm going to beat him and buy 100 shares of company A." Investor 3 says, "Everybody else is doing it, so I'm going to beat them and buy 1,000 shares of company A!" Thus if the number of shares bought is increasing by a factor, so that we have N, N^2, N^3, etc., all added together, we have a hypergeometric distribution sum and a heavy tail condition. Heavy-tailed behavior can be particularly dangerous in today's financial environment, where the amount of derivatives in existence is estimated at $250 to $370 trillion—about 10 to 12 times the world's annual gross domestic product (GDP). [9, 10] If hypercorrelated linkages exist in even a small portion of these derivatives—and very little is known of how they are linked—then the results could be catastrophic for the world financial system, as they almost were with the long-term capital management hedge fund crash. Since hypercorrelated behavior can mimic ordinary normal behavior under many circumstances, it might lie there undiscovered. In May 2005, there was a "correlation crisis" in the market for a type of security derivative known as a synthetic collateralized debt obligation (CDO). [11] In this case, there was apparently a single hedge fund (highly leveraged) caught by an unexpected stock market move. Facing huge losses, the fund tried to liquidate its positions, revealing the hypercorrelation—

> The hedge fund, desperate to avoid a liquidity crunch, was reduced to hawking its positions around the broader marketplace, and met a distinct lack of demand. Hedge funds and investment bank traders "began to see the writing on the wall and started to try and offload similar positions themselves," in the words of one CDO banker. [12]

While this did not bring down the financial markets, no one knows how severe a correlation crisis might become in the future, possibly due to hypercorrelation. And, in fact, there is further evidence of just such hypercorrelation in the derivatives arena. Regarding the CDOs, an analyst has made the following rather astonishing observation—

> In effect, that means one asset—such as a mortgage loan—is being used and reused many times over to create new trading and hedging opportunities. It is akin, in a sense, to how a small amount of sugar can be spun up into a huge cone of candy floss. [13]

Given that derivatives now have a total value of 802 percent of world GDP, and 75 percent of its liquidity, [14] and that the foregoing observation suggests hypercorrelation of a significant amount of these derivatives, an unexpected turn of events could quickly lead to a potentially dangerous situation due to heavy tails (more "10 sigma events").

In the case of Internet traffic, the psychology of users and its impact on traffic has not been well-studied, and the heavy-tailed nature of the traffic may have multiple sources. But some ideas come to mind. If success breeds a desire for more of the same, for instance, if successful retrieval of one Web page leads to requests for larger numbers by the same user, or pages with more information, a type of hypercorrela-

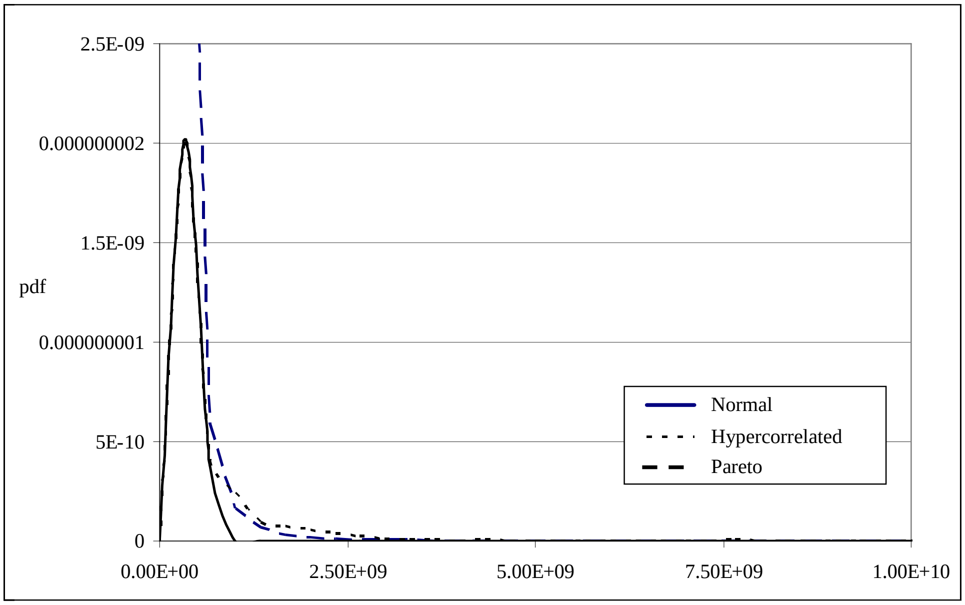

(a) Linear Scale

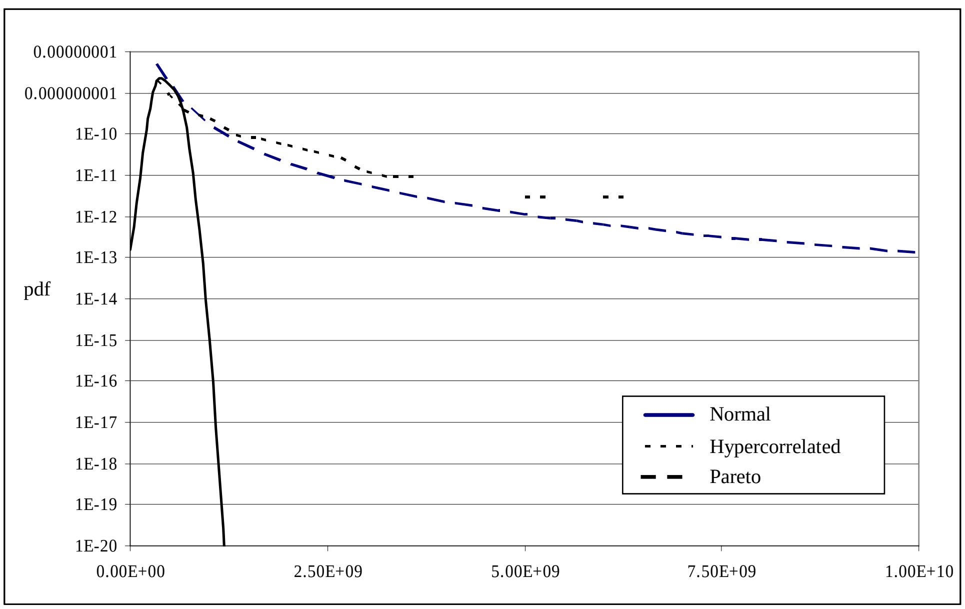

(b) Log Scale

Figure 7. Summed Hypercorrelated Variables

Note: Distribution of a: μ=10, σ=3; distribution of n: μ=6, σ=2; average hypercorrelation value: 5.923; Pareto constant α=2.1

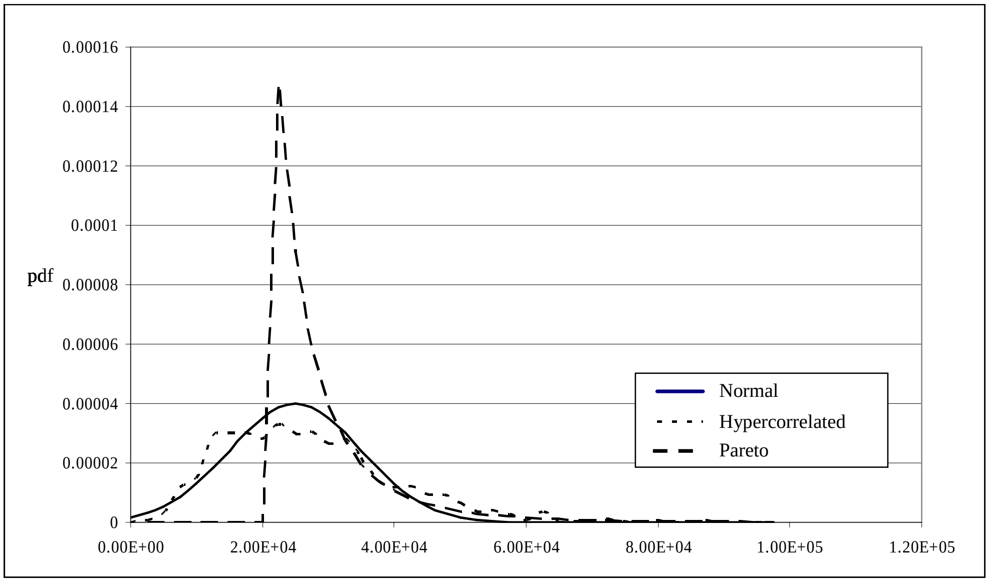

(a) Linear Scale

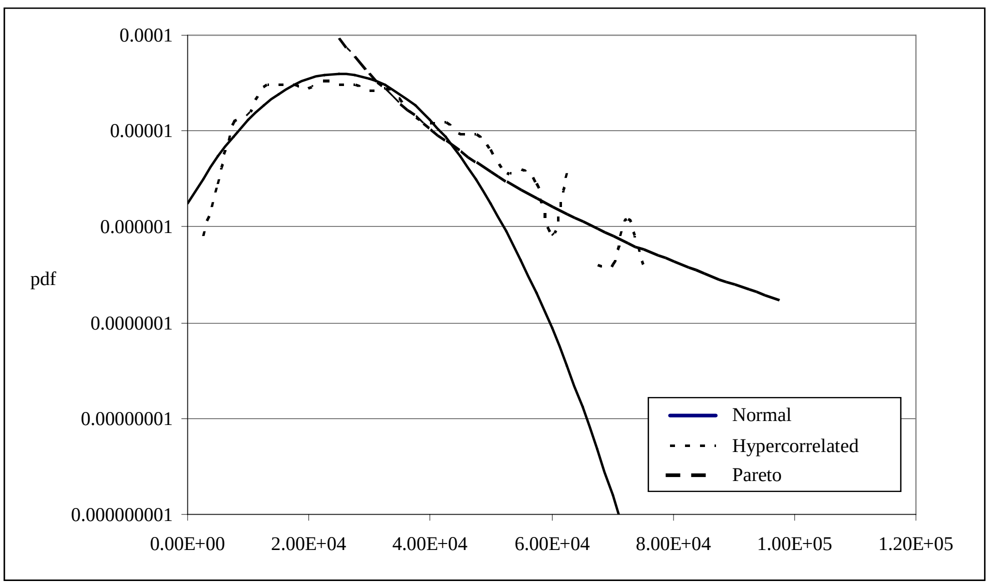

(b) Log Scale

Figure 8. Summed Hypercorrelated Variables

Note: Distribution of a: μ=10, σ=3; distribution of n: μ=3, σ=1; average hypercorrelation value: 1.868; Pareto constant α=3.6

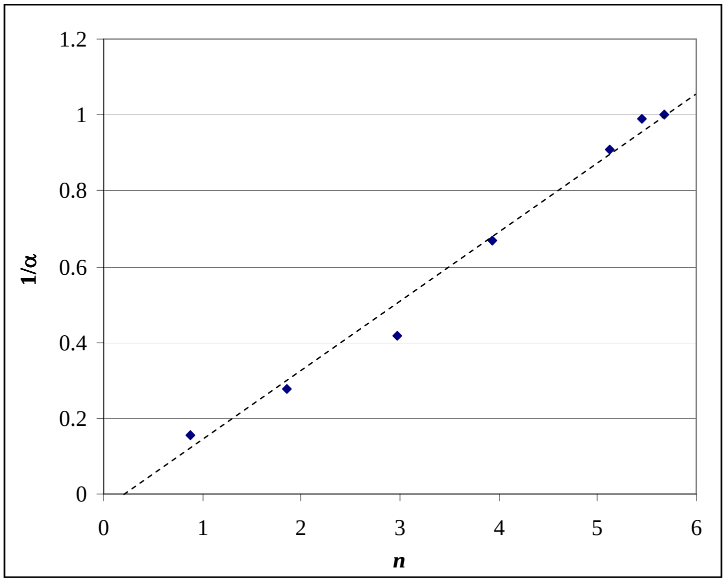

Figure 9. Relationship Between Hypercorrelation n and Pareto Constant α Based on Simulation Runs

lation will emerge. As an example, let us consider how most commercial sales Web sites are set up. If a user is seeking information about a product that may be of interest, the user would first arrive at a summary page which gives highlights of the product. Such a page is typically designed to load fast. Then the user might want to drill down a bit, and ask for a more detailed specifications page. Later, the user might ask for a product brochure download, or a page with photos or images of some type. The ratio of file size here could easily be 1:5:25, or something on that order. This would explain, at least in part, the heavy-tailed character of traffic, page requests per site, reading time per page, and session duration. Another example might be the "chain-letter" effect. An individual emails five friends, who then email five of their friends, and so forth. Since sending emails to multiple persons is as easy as pasting a list in the address field, this type of behavior can come about with very little effort on the part of the user (as opposed, say, to phone trees, where more physical effort is required to dial five separate numbers and leave five voice messages).

Very large ocean waves (~30 m or more) have become a concern recently, as they occur with much greater frequency than a normal distribution of wave heights would suggest. As recently as 1995, they were thought to occur only once every 10,000 years; now it is believed that one or more may be occurring at any time somewhere in the oceans. [15] Such waves, known as "rogue waves," "monster waves," or "king waves," can swamp a supertanker or large passenger ship. Between 1985 and 2004, more than 200 supertankers have sunk, many of them believed to be the victim of rogue waves. [16] Existing theories about the origin of these waves, focusing in space in time, current focusing, and nonlinear focusing, have not yielded convincing explanations. [17] One theory is that since the rogue waves are often found near a strong current, such as the Gulf Stream off of the east coast of North America, the current acts like a lens which focuses wave energy by combining small waves into a larger wave. Converging weather fronts, where waves from different systems often combine, could also be a cause. In both of these cases, since smaller waves are combining, if the mechanism can amplify the contributions of the smaller waves so that they become hypercorrelated, a heavy tail for wave height could result. In such a case, rogue waves could appear with much greater frequency than intuition would suggest. The actions of the combining mechanism, with respect to the possibility of hypercorrelation, deserve further study.

Conclusions and Future Work

Heavy-tailed probability density functions can arise through summations of hypercorrelated variables, which are variables of the general form $Y = aX^n$. The larger the value of *n*, the greater the degree of heavy-tailed behavior. Distributions of such sums, $Z = \sum_i Y_i$, can often mimic normal distributions for certain values *Z*. Therefore application of the Central Limit Theorem must be carefully monitored to ensure that hypercorrelation is not present. Predictions made or actions taken on the assumption of normal distribution behavior could yield catastrophic results when events thought to be "10 sigma" occur. Heavy-tailed behavior can be made more apparent by plotting the sums on a logarithmic scale. This research only deals with the case of heavy-tailed distributions arising from sums of random variables; heavy tails may also originate in other ways. Future work will concentrate on a more theoretical understanding of the emergence of heavy tails from hypercorrelated random variables. It will also look at empirical reasons for the emergence of heavy-tailed behavior in order to better understand the circumstances under which such behavior can be expected.

Notes and References

1. Meyer, Paul, *Introductory Probability and Statistical Applications*, Second Edition, Reading, MA: Addison-Wesley, p. 251, 1970.
2. *http://www.investopedia.com/terms/l/longtermcapital.asp.*
3. Fowler, T. B., "A Short Tutorial on Fractals and Internet Traffic," *The Telecommunications Review*, Volume 10, Noblis, 1999.
4. Fischer, M. J. and C. M. Harris, "A Method for Analyzing Congestion in Pareto and Related Queues," *The Telecommunications Review*, Volume 10, Noblis, 1999.
5. Fischer, M. J., D. Gross, D. M. B. Masi, and J. F. Shortle, "Analyzing the Waiting Time Process in Internet Queueing Systems With the Transform Approximation Method," *The Telecommunications Review*, Volume 12, Noblis, 2001.
6. Crovella, M. and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Transactions on Networking*, Volume 5(6), pp. 835–846, December 1997.
7. Willinger, W. and V. Paxson, "Where Mathematics Meets the Internet," *Notices of the American Mathematical Society*, Volume 45, Number 8, pp. 961–970, August 1998.
8. Crovella, M. and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Transactions on Networking*, Volume 5(6), pp. 835–846, December 1997.
9. *OTC Derivatives Market Activity in the Second Half of 2004*, Bank for International Settlements, p. 7, May, 2005, *www.bis.org/publ/otc_hy0505.pdf.*
10. Tett, Gillian "Derivatives Market Grows by 25pc," *The Australian-FT Business*, November 20, 2006, *http://www.theaustralian.news.com.au/story/0,20867,20785202-36375,00.html.*
11. A CDO is a group or pool of debt instruments, such as derivatives, loans, or bonds, which is used as the basis for the issuance of new securities.
12. "Not Very Sexy, Just Very Profitable," *Financial Times*, p. 19, December 29, 2006.
13. "Should Atlas Still Shrug? The Threat That Lurks Behind the Growth of Complex Debt Deals," *Financial Times*, p. 9, January 15, 2007.
14. Ibid.
15. "New Theory (And Old Equations) May Explain Causes of Ship-Sinking Freak Waves," September 13, 2006, *http://www.physorg.com/news77381892.html.*
16. Walker, Cameron, "Monster Waves Surprisingly Common, Satellites Show," *National Geographic News,* August 10, 2004, *http://news.nationalgeographic.com/news/2004/08/0810_040810_rogue_waves.html.*
17. Dysthe, K. B., H. E. Krogstad, H. Socquet-Juglard, and K. Trulsen, "Freak Waves, Rogue Waves, Extreme Waves and Ocean Wave Climate," 2001, *www.math.uio.no/~karstent/waves/index_en.html.*

About the Author

Thomas B. Fowler is a senior principal engineer at Noblis where his experience includes computer-aided education systems, air traffic control, information systems, telecommunications systems, and next generation government telecommunications and forecasting trends and usage patterns. He received his doctorate of science degree in systems and control theory from The George Washington University. Contact him at tfowler@noblis.org.