

Elementary Statistics of Tail Events

William Nordhaus

April 8, 2011

[This note is excerpted from an article on Weitzman for general economists. It describes some of the technical features of fat tailed distributions. The longer piece is forthcoming in the *Review of Environmental and Economic Policy*.]

The Economic Importance of Tail Events

From time to time, something occurs which is outside the range of what is normally expected. For reasons that will soon become clear, I call this a tail event. Some tail events are unremarkable, such as a bogus email about a large inheritance that awaits you in Nigeria. Others may change the course of history. Momentous tail events include the detonation of the first atomic weapon over Hiroshima in 1945, the sharp rise in oil prices in 1973, the 23 percent fall in stock prices in October 19, 1987, the destruction of the World Trade Towers in 2001, and the meltdown of the world financial system in 2007-08. A tail event is an outcome which, from the perspective of the frequency of historical events or perhaps only from intuition, should happen only once in a thousand or million or centillion years.

Tail events are more than statistical curiosities. In some cases, they may be so important that they dominate the way we think about our options and our strategies. Obviously, tail events dominate thinking about nuclear weapons. Less obvious is how to deal with tail events in economics. One example of how tail risk has changed economic policy is in the area of finance. In response to the meltdown of the banking system in 2007-2008, the theoretical approach to bank regulation has moved toward containing “systemic risk” rather than individual bank risk.

The Problem of Fat Tails

Low-probability, high-consequence events can dominate the impacts and societal concerns for many issues, of which climate change is a signal example. This is the problem known as “fat tails.”

To illustrate the problem of fat tails, it is helpful to first picture a probability distribution such as the common “bell curve” or normal distribution. The normal distribution has most observations clustering around the center, with few showing highly divergent results. Take the height of American women as an example. This variable has close to a normal distribution with a mean of 64 inches and a standard deviation of 3 inches. Based on the properties of a normal distribution, 95 percent of

American women will be between 58 and 70 inches tall. How likely is it that you will observe a 11-foot-tall woman? This is about 23 standard deviations from the mean, which is an exceedingly small number for a normal distribution (about 10^{-230}). Indeed, the world's tallest woman is reported to be about 8 feet tall.

Other probability distributions have the property that from time to time a very unlikely looking event occurs. When this is the case, we say that we have witnessed a "tail event" and that the tail of the distribution is "fat" rather than medium as in the case of the bell curve.

"Multi-Sigma Events"

People sometimes refer to "four-sigma" or "six-sigma" events. These are shorthand terms for how many standard deviations from the average something is. Returning to women's height, if you see a woman who is 6 foot tall, that is a three-sigma event. For example, in a normal distribution, a three-sigma positive shock (or a shock three standard deviations above the mean) will occur about once every 200 observations. So, this suggests that only 1 in 200 women will be taller than 6 feet.

We can examine daily changes in stock prices to illustrate this point. Prices on U.S. stock markets fell approximately 23 percent on October 19, 1987. An estimate of the daily standard deviation of price change over the 1950-1986 period shows a standard deviation of 1 percent. If this were a normal distribution, then we would see a 5 percent change in prices once every 14,000 years and a 7.2-sigma change about once in the life of the universe. However, 23-sigma events, like 11-foot people, simply do not occur for a normal distribution.

Yet, these large deviations occur much more frequently than would be predicted by the normal distribution. An interesting example would be the returns on the U.S. stock market. I have calculated the monthly returns for stock prices for the 140-year period from 1871 to 2010. As shown in Table 1, the actual maximum and minimum increases are much larger than would be found with a normal distribution. The maximum is a "10 sigma event," which would almost never happen with a normal distribution (probability less than one day in the life of the universe). So people who think that financial markets follow the bell curve will from time to time find themselves very surprised.

Largest increase in 140 years	
Actual	40.7
Normal distribution	14.3
Largest decrease in 140 years	
Actual	-30.8
Normal distribution	-13.7

Table 1. Comparison of normal and actual distribution of stock price changes

This table displays the phenomenon of fat tails for the stock market. I looked at 140 years of stock price changes measured in percent per month (from 1871 to 2010). I then took a series with the same mean and standard deviation. The largest actual increase in the sample was 40.7 percent, whereas the largest with a normal distribution would be 14.3 percent. Similar numbers are seen for declines as well. (Source: Robert Shiller and DRI data base. Increases are logarithmic.)

It has been known for many years that there are large deviations from the normal or bell-curve distribution for the stock market as well as for many other phenomena. Statisticians have developed both probability distributions that have fat tails and techniques to estimate these other distributions. A particularly interesting probability distribution that may have fat tails is one that is known as the “power law.” This refers to a distribution in which the probability is proportional to a value to a power or exponent. Power law distributions were introduced into economics by Benoit Mandelbrot (1963), and are widely used in the natural and social sciences. One example of this is the power law for earthquakes, the sizes of cities and firms, solar flares, moon craters, wars, incomes, wealth, and commodity prices.

The Element of Surprise

One way to think about tail events is to note that you can be extremely surprised by an outcome when a process has fat tails. This “surprise” can be measured by asking how deviant of an observation might turn up when you have many observations. For example, suppose that you have looked at housing prices for 50 years and observe that they have never declined. So you place your bets on housing prices continuing to rise. But then you get a draw from a fat-tailed distribution and housing prices fall – not just a little, but 20 or 40 percent.

Now suppose you were in the oil market in the early 1970s. Suppose further you were a trader who eschewed any economic theories and just looked at historical data (which is not an absurd approach given the unpredictability of oil prices). Oil prices had been pretty stable, and the “sigma” was about 5 percent on a monthly basis. But then, in 1973, we had, by historical standards, a 37-sigma event (see Figure 1, which also shows the high-sigma stock-market and oil-price surprises events on a monthly basis). No wonder many people thought the economic world as we knew it was coming to an end in 1973. If we could have a 37-sigma event, then almost nothing consistent with the laws of nature could be ruled out if we think the distribution is normal.

Enter fat tailed distributions.

There has been much historical and statistical research on “surprises” or fat tails over many years. The conclusion of this research – on oil prices, stock prices, earthquake size, war fatalities, and many other phenomena – is that we were surprised for the wrong reasons. We might have thought that the spikes in Figure 1 were near-zero-probability events. They were not. Rather, we were surprised because the distributions were not normal. That is, they had fat tails, which means that the probability of the “way out” events was much greater than would be predicted by the normal distribution.

How Can We Know Which Distributions Are Fat-Tailed?

This research solved one problem but raised another. The new problem is, which are the fat-tailed distributions and which are the thin-tailed ones? When should we be on the lookout for high-sigma events? And how can we answer these questions before the high-sigma event occurs?

I have already mentioned the fat-tailed probability distribution associated with the power law. This is known in statistics as the Pareto distribution, after the Italian economist Vilfredo Pareto (who also introduced the important concept of a Pareto optimum). The important point about this distribution is that the probability of high-sigma events declines slowly relative to distributions like the normal.

Technical Definitions of Fat Tails

There is no generally accepted definition of the term “fat tails,” which is also sometimes called “heavy tails.” One set of definitions (Schuster 1984) divides distributions into three classes. A thin-tailed distribution has a finite upper limit (such as the uniform distribution), a medium-tailed distribution has exponentially declining tails (such as the normal distribution), and a fat-tailed distribution has power-law tails (such as the Pareto distribution). This definition is proposed in Eugene F. Schuster (1984).

To further illustrate these distinctions, it is helpful to go a step further and present the mathematics. Begin with the Pareto distribution, where the probability of an event is $P = k_1 X^{-(1+a)}$, where a is the Pareto “shape parameter,” which reflects the importance of tail events; X is the variable of interest; and k_1 is a constant that ensures that the sum of probabilities is 1. If a is very small, then the tail is very fat and the variable has a highly dispersed distribution. As a gets larger, the tail looks more like a normal distribution.

Figure 2 shows an example for a normal distribution and a Pareto distribution with a slope parameter of $a = 1.5$. The curves show the probability that an event will be at least N sigma, where sigma is a measure of dispersion. Note that the probability of a surprise is very small for normal distributions once a 4 or 5 sigma threshold is reached. For this version of the Pareto distribution, which is found for many economic and physical variables, the 4-sigma probability is still substantial.

Applications

There are many applications of the statistics of tail events. An interesting recent one is earthquakes and tsunamis. It turns out that earthquake energy is a Pareto distribution with a very fat tail (a Cauchy distribution in which $a = 1$). If earthquakes really follow this distribution, then the expected value of earthquake power is unbounded! You can see why the Japanese were surprised in March 2011.

Another interesting application is finance. Stock price changes are Paretian or power law. Luckily, unlike earthquakes, the expected value of the change is finite, but the variance is infinite! That makes a mockery of “mean, variance” analysis in portfolio theory and other theories which require a finite variance of price changes. There is a beautiful paper by Paul Samuelson on this issue but this is largely ignored.

A final example is the distribution of income and wealth, which are also power law. The issue here is why? What process leads to this? This is an unsolved problem in economic theory (or at least one I haven’t seen the answer to).

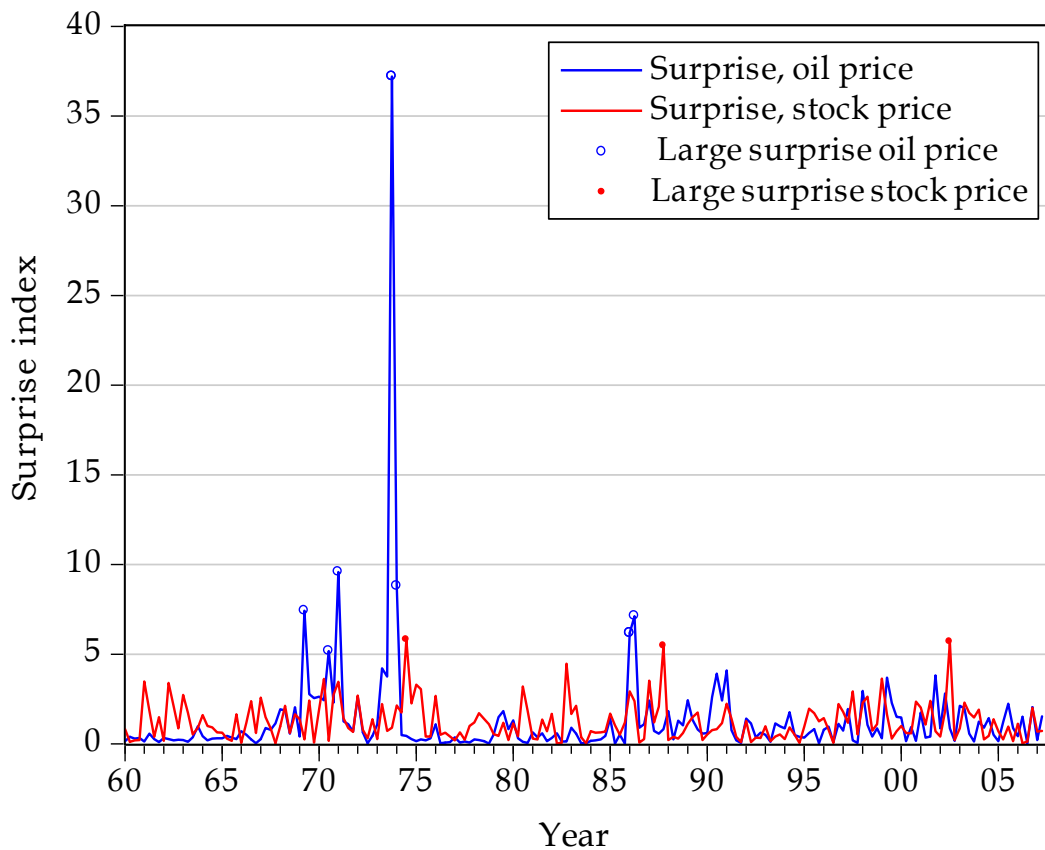


Figure 1. Surprise index for oil prices and stock prices, 1960-2007

Notes: The “surprise index” is measured as the 3-month change in the logarithm of the price divided by the 20-year moving average volatility, where each is measured monthly. The circles show the periods when the surprise was more than five moving standard deviations, with open circles for stock prices and solid circles for oil prices.

Sources: Nordhaus (2007a)

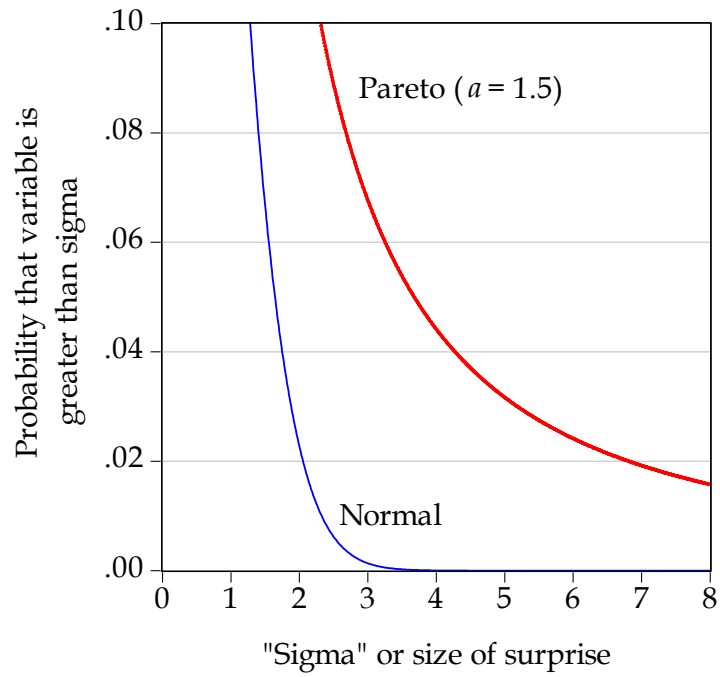


Figure 2. Illustration of fat tails for a normal distribution and a Pareto distributions with scale parameter $a = 1.5$.

References

Benoit Mandelbrot. 1963. "The Variation of Certain Speculative Prices," *The Journal of Business*, 36, No. 4, 394-419.

Nordhaus, William, 2007a. "Who's Afraid of a Big Bad Oil Shock?" *Brookings Papers on Economic Activity*, 2007:2.

Schuster, Eugene F. 1984. "Classification of Probability Laws by Tail Behavior," *Journal of the American Statistical Association*, Vol. 79, No. 388, Dec., 936-939.