

Déjà Vu All Over Again

By Paul D. Kaplan

When risk models fall short, advisors need to look no further than the historical record to plan for the next 100-year flood.

“We seem to have a once-in-a-lifetime crisis every three or four years.”

—Leslie Rahl, founder of Capital Market Risk Advisors¹

The dramatic events on Wall Street and in financial centers around the world that started on “Black Sunday,” Sept. 14, have upset many common assumptions about the global financial system. What started as a mortgage crisis spread to nearly every corner of the financial system when Lehman Brothers collapsed, Merrill Lynch sold itself to Bank of America, and AIG became strapped for cash—all in a single weekend. These and the events that followed have shaken investor confidence to the core. As of Dec. 31, the Dow Jones Industrial Average was down 22.4% since Black Sunday. The yield spread on junk bonds over LIBOR reached an unprecedented 16%. The markets for many assets have become illiquid, and credit is dried up for nearly anyone who needs it. The U.S. Federal Reserve, the U.S. Treasury, and their counterparts around the world have taken dramatic steps to restore liquidity to asset markets, stimulate lenders to make loans again, shore up investor

confidence in equity markets, and avoid a deep global recession.

If you need to be reminded how bad things are, listen to our political and fiscal-policy leaders as they describe the crisis with phrases that begin with the ominous words “once in a” As they were pushing their \$700-billion bailout package last fall, members of the Bush administration said that the crisis was a “once-in-a-century event,” and this was echoed in November by Henry Paulson, the former secretary of the U.S. Treasury, who said the meltdown was a “once- or twice-in-a-100-year event.” Former Federal Reserve chairman Alan Greenspan characterized the crisis as a “once-in-a-century credit tsunami.”

There’s little doubt that aspects of this crisis are unique and that the economy is facing its hardest challenge since the Great Depression, but are severe economic crises the rare events Paulson, Greenspan, et al., have suggested? A study of capital market history suggests no. To see this, you need to look no further than the Ibbotson Stocks, Bonds, Bills, and Inflation poster from Morningstar hanging on your wall.

Take, for example, the poster’s depiction of the compound annual return of the S&P 500 Index, identified on the chart as Large Stocks.^{2,3} The growth of \$1 to \$2,049 over 83 years is impressive (a rate of 9.6% per year), but the record is peppered with several long and severe declines, some in the not-too-distant past.

To illustrate our point, we isolated the S&P 500 line of the poster and added blue areas that show the highest level that the cumulative value of the S&P 500 had achieved as of that date (Exhibit 1). Wherever a blue area is shown, the S&P 500 was amid a decline relative to its most recent peak. The deeper the gap, the more severe the decline; the wider the gap, the longer the time until the S&P 500 returned to its peak. Wherever a blue area is not shown, the S&P 500 was climbing to a new peak.

Not surprisingly, the granddaddy of all market declines started with the Crash of 1929 and did not recover until 1945. The S&P 500 lost more than 83% of its value in about three years and took 12½ years to recover. What may be

¹ As quoted by Christopher Wright, “Tail Tales,” *CFA Institute Magazine*, March/April 2007. ² We obtained the historical monthly total returns from Morningstar EnCorr, an institutional asset-allocation software and data package. ³ We use a logarithmic scale for all growth of \$1 charts.





Exhibit 1: Mind the Gaps U.S. large-cap stocks have made impressive gains over the years, but several significant declines have interrupted the S&P 500's trajectory.



Growth of \$1 includes reinvested dividends. Monthly data used to calculate returns.

more sobering, however, is that the second-greatest decline took place within the past decade. With the crash of the Internet bubble in 2000, the S&P 500 lost almost 45% of its value over a two-year period and took four years to return to its peak value.

In all, including the current crisis, the S&P 500 has suffered eight peak-to-trough declines of more than 20% since the mid-1920s. Two of the three greatest declines occurred in the past eight years. To suggest that the current crisis is a once-in-a-century event ignores the record.

Measuring Risk: The Standard Model

With 20% declines occurring, on average, every decade or so, you'd think that the standard risk models that investors use to make their asset-allocation decisions would assign a significant probability that these events will occur. Think again. To see why, we need to look at the history of how these models were formed.

To help make sense of the highly complex capital markets, financial economists in 1960s and 1970s developed a set of mathematical models of the markets that are used to this day

throughout the investment profession. The best known of these models are the capital asset pricing model of expected returns and the Black-Scholes option pricing model. These models' creators have won the Nobel Prize in economics for their path-breaking work. Each of these models starts by making an assumption about the statistical distribution of stock market returns. The CAPM assumes that returns follow a normal, or bell-shaped, distribution. The Black-Scholes model assumes that returns follow a lognormal distribution.⁴

⁴ For returns to follow a lognormal distribution means that logarithm one plus the return in decimal follows a normal distribution.





Exhibit 2: Cracks in the Bell Standard risk models assume S&P 500 returns follow a bell-shaped distribution, even though the index has experienced more than 10 declines of at least -13% .



Histogram shows the frequency of monthly returns for the S&P 500 from January 1926 to November 2008.

With these standard models, the primary measure of risk is standard deviation. If returns follow a normal distribution, the chance that a return would be more than three standard deviations below average would be a trivial 0.135%. Since January 1926, we have 996 months of stock market data; 0.135% of 996 is 1.34—that is, there should be only one or two occurrences of such event.

But the record of the stock market tells a different story. The monthly returns of the S&P 500 have been more than three standard deviations below average 10 times since 1926. In other words, the standard models assign meaninglessly small probabilities to extreme events that occur five to 10 times more than the models predict.

We can illustrate the problem further by overlaying a lognormal model of returns on a histogram of monthly total returns on the S&P 500 (Exhibit 2). The model says that declines of more than negative 13% have almost no chance of happening—yet they have occurred at least 10 times since 1926.

An Alternative Approach: Log-Stable Distributions

In the early 1960s, Benoit Mandelbrot, a mathematician teaching economics at the University of Chicago, was advising a doctoral student named Eugene Fama. Mandelbrot had developed a statistical model for percentage changes in the price of cotton that had “fat tails.” That is, the model assigned nontrivial probabilities to large percentage changes. In his doctoral dissertation, Fama applied

Mandelbrot’s model to stock prices and obtained promising results.⁵ Until recently, however, the work of Mandelbrot and Fama had been largely ignored.⁶

In his dissertation, Fama assumed that the logarithm of stock returns followed a fat-tailed distribution called a “stable Paretian distribution,” or stable distribution.⁷ Hence, we refer to the resulting distribution of returns as a “log-stable distribution.”

We can illustrate an example of Fama’s work by using the same S&P 500 histogram in our earlier exhibit but with a log-stable distribution curve overlaying it instead of a lognormal curve.⁸ The log-stable model (Exhibit 3) fits the empirical distribution much closer than the lognormal both at the

⁵ For an account of the work of Mandelbrot and Fama during this period, see Benoit Mandelbrot and Richard L. Hudson, *The (Mis)Behavior of Markets*, New York: Basic Books, 2004. ⁶ The idea of using fat-tailed distributions to model asset returns is starting to gain some traction. FinAnalytica was founded to provide investment analysis and portfolio construction software based on Mandelbrot and Fama’s work. Morningstar added distribution charts and forecasting models based on it to Morningstar EnCorr. ⁷ Strictly speaking, the assumption is that the logarithm of one plus the return in decimal form follows a stable Paretian distribution. ⁸ This chart can be produced in Morningstar EnCorr Analyzer using the log-stable feature.





Exhibit 3: It's a Fat-Tailed World, After All A log-stable distribution does a good job of modeling the empirical returns of the S&P 500, especially at the center and the tails.



Histogram shows the frequency of monthly returns for the S&P 500 from January 1926 to November 2008.

center and the tails. In particular, note the close match between the density curve and the histogram between negative 13% and negative 29%.


The tails of a stable distribution are so fat that its variance is infinite. In other words, the concepts of standard deviation and variance are not defined for stable distributions. You might find the idea of an infinite variance counterintuitive, because it is possible to calculate a standard deviation for any finite set of data. However, the underlying mathematical distributions that we use to model asset returns assign probabilities over the range from negative infinity to positive infinity.⁹ Some distributions that cover this infinite range assign so little probability out in the tails that variance can be defined. These are “thin-tailed”

distributions, the normal or bell-shaped distribution being the best-known example. Other distributions assign so much probability to the tails that variance is infinite. Such is the case with stable distributions.


The manner in which a stable distribution assigns probability to its tails is very close to what is known as “power law.” When a distribution of a loss follows a power law, a plot of logarithm of the magnitude of loss (x) versus the logarithm of the probability of the loss turning out to be x or worse is a downward-sloping straight line. Therefore, while the probability of loss decreases with the magnitude of loss, it does so gradually.

In Exhibit 4, we plot the magnitude of loss versus the logarithm of the probability of

Exhibit 4
Power Law Tails: Unlike a normal distribution, a stable distribution approaches the straight line of a power law, indicating that it has “fat tails.”



⁹ That is the probability distribution of one plus the return on an asset return in decimal form. The lowest possible return on an unlevered position in an asset is negative 100%, which is negative 1 in decimal form. Adding one we get 0. The logarithm of 0 is $-\infty$.




loss for a normal distribution, a stable distribution, and a power law distribution. The line for the normal distribution curves down, indicating that it has thin tails. In contrast, the line for stable distribution approaches the straight line of the power law because it is very similar to a power law for large losses.

These results show that the log-stable distribution does a good job of modeling the empirical returns distribution of the S&P 500. The better fit of the log-stable distribution demonstrates that the S&P 500 has fatter tails than predicted by the lognormal model. It also calls into question commonly used portfolio construction techniques such as the mean-variance optimization, which relies on the assumption of a finite variance.

If the log-stable model does such a better job in describing the distribution of asset returns, why has it not received more acceptance? There are several possible reasons. First, the mathematics is challenging. Second, the variances and all higher moments of stable

Exhibit 5
Role of Time: The log-stable model indicates that there's a 4% to 5% probability that the S&P 500 will lose 50% or more over extended time periods. The lognormal model puts the odds much lower.



Hard Eight

The S&P 500 has suffered eight peak-to-trough declines of more than 20%.

Peak	Trough	Decline %	Recovery
August 1929	June 1932	83.41	January 1945
August 2000	September 2002	44.73	October 2006
December 1972	September 1974	42.64	June 1976
October 2007	November 2008	40.89	To Be Determined
August 1987	November 1987	29.58	May 1989
November 1968	June 1970	29.16	March 1971
December 1961	June 1962	22.28	April 1963
May 1946	November 1946	21.76	October 1949

Table shows the worst cumulative peak-to-trough declines in percentage terms since December 1925. Based on monthly returns.

random variables are infinite. The lack of a finite variance means that most portfolio theories and most portfolio construction techniques are invalid, including those based on alternative risk measures such as “downside risk.” Finally, there is no single obvious way to estimate the parameters of stable distributions as there is with normal distributions.

Risk Measures versus Risk Models

For advisors, the lesson here is not that they should throw away the standard ways of summarizing risk using measures such as standard deviation and downside deviation.¹⁰ Nor should advisors run to embrace Fama’s log-stable models.

Instead, we think advisors should understand the limitations of standard risk measures and have a basic understanding of what Mandelbrot’s and Fama’s work says about describing risk. Rather than solely relying on a few summary statistics to characterize the risks of


an investment, advisors would benefit by beginning to think about a more complete risk model. A complete risk model allows investors to consider three questions about a potential decline in value simultaneously:

- ▶ How likely might a decline occur?
- ▶ How long might it last?
- ▶ How bad might it get?

It is already common practice in some segments of the financial-services industry to use a risk model to measure “value at risk”—that is, how bad a loss might be over a given length of time and with a given probability.

As you can appreciate through our study of historical stock market declines, time horizon is a key dimension of risk not explicitly addressed by standard risk measures. A complete risk model can be used to explicitly take time horizons into account.

¹⁰ In recognition that return distributions may not be symmetric, measures such as skewness and kurtosis are sometimes presented alongside standard deviation. However, like variance, these measures are not defined for stable Pareitian distributions.



For example, in Exhibit 5, we plot the probability of a cumulative loss of 50% or more over various time horizons using the lognormal distribution for the S&P 500 that we show in Exhibit 2 and the log-stable distribution in Exhibit 3. The lognormal model shows that the risk of such a severe decline over an extended period is negligible. The log-stable model, on the other hand, indicates that such a loss over an extended period has a probability of 4% to 5%—numbers significant enough to gain the attention of risk-averse advisors and investors who might want to be prepared for such a scenario.

Conclusion

In every financial crisis, investors relearn the same message—there isn't a magic risk measure or model that can account for or predict every significant drop in the market. Economists and quantitative analysts have made incredible strides over the decades engineering new ways to explain the distribution of returns. These developments provide investors with valuable information to help them decide how to allocate their portfolios for any number of investing scenarios and mitigate risk. But they are not perfect.

As we've shown, the record contains a much bumpier ride than many risk models would suggest. In addition to preparing clients' portfolios for these occasional severe declines and taking other precautions, advisors would do well to keep reminding their clients of the risks they face as investors. Clients should be fully prepared to take on the 100-year floods they will surely face in the future. ■■■

Paul D. Kaplan, Ph.D., CFA, is Morningstar's vice president of quantitative research and a frequent contributor to *Morningstar Advisor*.

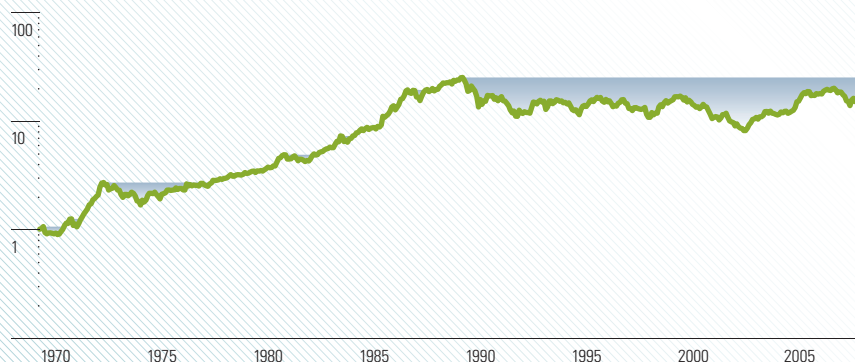
We Are Not Alone

The uneven performance of the stock market is hardly unique to the United States. Severe declines—mostly within the past decade—have occurred in developed markets since January 1970. Here are the worst declines for seven countries.


Country	Peak	Trough	Decline %	Recovery
Germany	February 2000	March 2003	67.89	April 2007
Japan	December 1989	April 2003	67.62	To Be Determined
U.K.	August 1972	November 1974	64.73	January 1977
Italy	June 1973	December 1977	59.39	September 1980
Spain	April 1974	November 1979	58.81	March 1984
France	August 2000	March 2003	58.28	March 2007
Canada	August 2000	September 2002	47.11	September 2005

Source: Morgan Stanley Capital International and Morningstar EnCorr. Chart shows monthly return data in local currency for major stock-market index in each country.

The Japanese market has yet to recover from its peak in December 1989.



The markets in four of the seven countries have performed worse since October 2007 than the U.S. market, which has fallen 40%.



Data through December 2008. Based on monthly returns.