A Better Model to Estimate Downside Risk Truncated Lévy Flight

James X. Xiong, Ph.D, CFA

Senior Research Consultant
Ibbotson Associates

The most common model of asset returns is assumed to be normally or Gaussian distributed (Bachelier, 1900). This model is natural if one assumes the return over a time interval to be the result of many small independent shocks, which leads to a Gaussian distribution by the central limit theorem. The model provides a first approximation of the behavior observed in empirical data. However, empirical studies have observed that the return distributions are more leptokurtic or fat-tailed than Gaussian distributions.

A normal distribution model assumes that asset return that is three standard deviations below their mean (three-sigma event) has a probability of only ~0.13%, i.e. once every 1000 times. For example, from January 1926 to April 2009 the S&P 500 total return index has had a monthly mean return of 0.91% and a monthly standard deviation of 5.55%. A negative three-sigma event would be a return of -15.74%. During this time period, there have been 10 monthly returns worse than -15.74% as shown in Table 1 (the three-sigma event). This implies the probability of a three-sigma event is 1% rather than 0.13%, or eight times greater than we would expect under a normal distribution. Hence, a normal distribution fails to describe the "fat" or "heavy" tails of the stock market.

Table 1. The worst 10 monthly returns for S&P 500 (from Jan 1926 to June 2009)

	S&P 500 (%)
Sep 1931	-29.73
Mar 938	-24.87
May1940	-22.89
May 1932	-21.96
Oct 1987	-21.52
Apr 1932	-19.97
Oct 1929	-19.73
Feb 1933	-17.72
Oct 2008	-16.79
Jun 1930	-16.25

Source: Morningstar Encorr

Many statistical models have been put forth to account for the heavy tails. Well-known examples are Mandelbrot's Lévy stable hypothesis (Mandelbrot, 1963), the Student's t-distribution (Blattberg and Gonedes, 1974), and the mixture-of-Gaussian distributions hypothesis (Clark, 1973). The latter two models possess finite variance and fat tails, but they are not stable which implies that their shapes are changing at different time horizons and that distributions at different time horizons do not obey scaling relations.

An alternative is the Lévy stable distribution model (Lévy, 1925). In 1963, Mandelbrot modeled the cotton prices with a Lévy stable process, and his finding was later supported by Fama in 1965. A Lévy stable distribution model has fat tails and obeys scaling properties, but it has an infinite variance which is in violation with empirical observations that the return variance is finite. The infinite variance would also complicate the task of risk estimations.

In the context of logarithm of asset returns (plus 1), the corresponding model is the lognormal and log-stable models. Kaplan (2009) illustrates that the lognormal distribution fails to fit the left tails of the S&P 500 return distribution, while the log-stable distribution does a much better job. However, the stable distribution suffers from an infinite variance as mentioned above and thus its tails are perhaps too fat as shown in Chart 1. The vertical axis of Chart 1 is in log scale with a base of 10, and this helps to view the tails of the distribution more clearly.

Chart 1: Historical Distributions of S&P 500 Monthly Returns Fitted by Log-Stable and Lognormal Models (from Jan 1926 to June 2009)

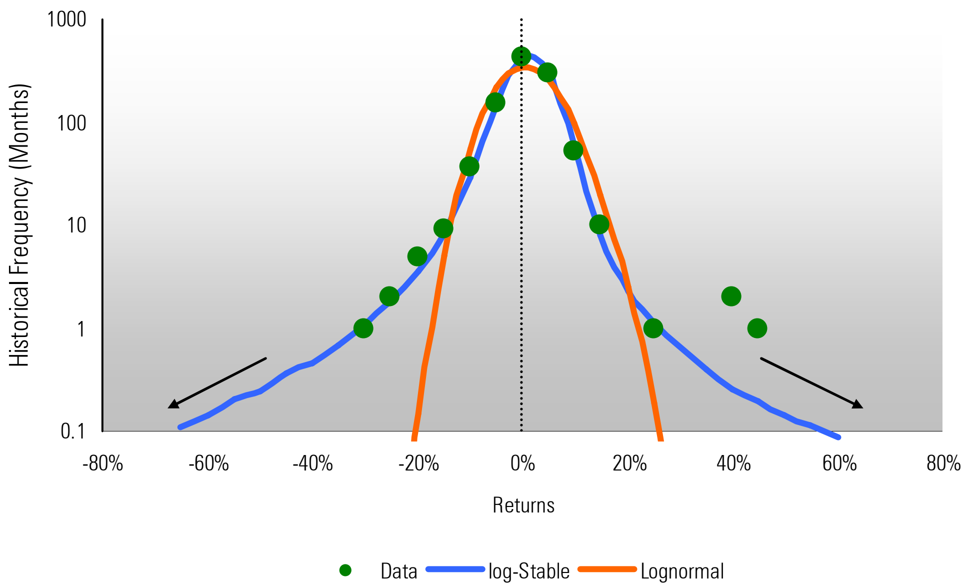

Do we have a better distribution model so that its tail is appropriately fat? Yes. A simple solution is to truncate the tails of the Lévy stable distribution, which results in what known as the Truncated Lévy Flight (TLF). The TLF distribution has finite variance, fat tails, and more importantly scaling property. The scaling property means that the shapes of distributions are the same at different time intervals.

Previous studies (Mantegna and Stanley, 1999) have demonstrated that the TLF model describes return distributions measured at small time horizons quite well. Chart 2 compares the log-TLF model with the lognormal model in fitting the historical *monthly* return distributions for S&P 500 (See Xiong (2009) for more details). It can be seen that the log-TLF model provides an excellent fit for S&P 500 in all aspects: center, tails, minimum and maximum.

Chart 2: Historical Distributions of S&P 500 Monthly Returns Fitted by Log-TLF and Lognormal Models (from Jan 1926 to June 2009)

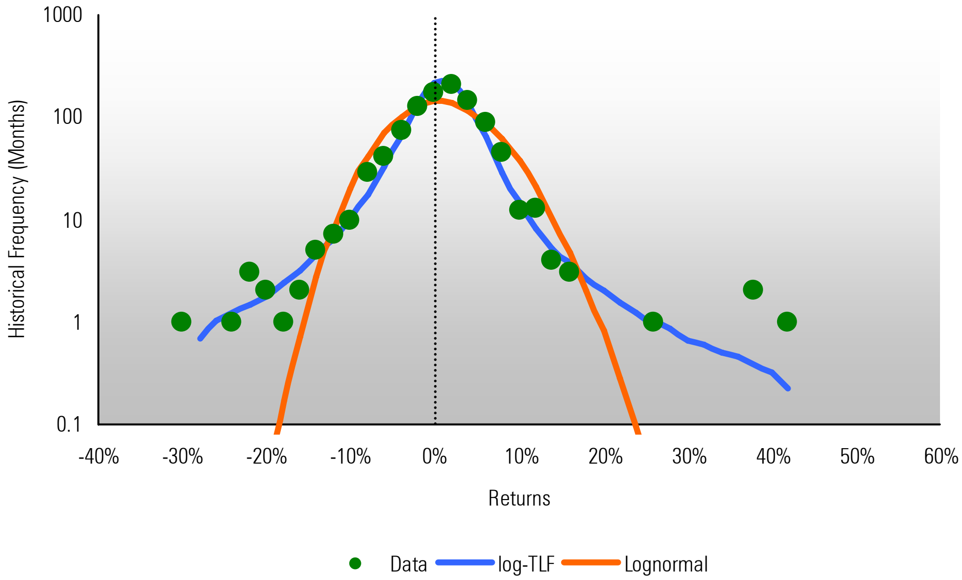

Based on our fit shown in Chart 2, the lognormal model underestimates the monthly CVaR by 2.27% for S&P 500. Assuming self-similarity of the monthly and annual return distributions, the annualized CVaR are underestimated by 7.86% ($= \sqrt{12} \cdot 2.27\%$) for S&P 500. This is a significant underestimate for the lognormal model!

We apply the same TLF model to bond indices, and Chart 3 shows an example of U.S. long-term government bond. Again, the log-TLF does an excellent job in fitting the entire return distributions of the bond. The monthly and annual CVaRs are underestimated by 0.48% and 1.66% by the lognormal model, respectively.

Chart 3.: Historical Distrib of US Long-Term Government Bond Monthly Returns Fitted Log-TLF and Lognormal Models (fJan 1926 to June 2009)

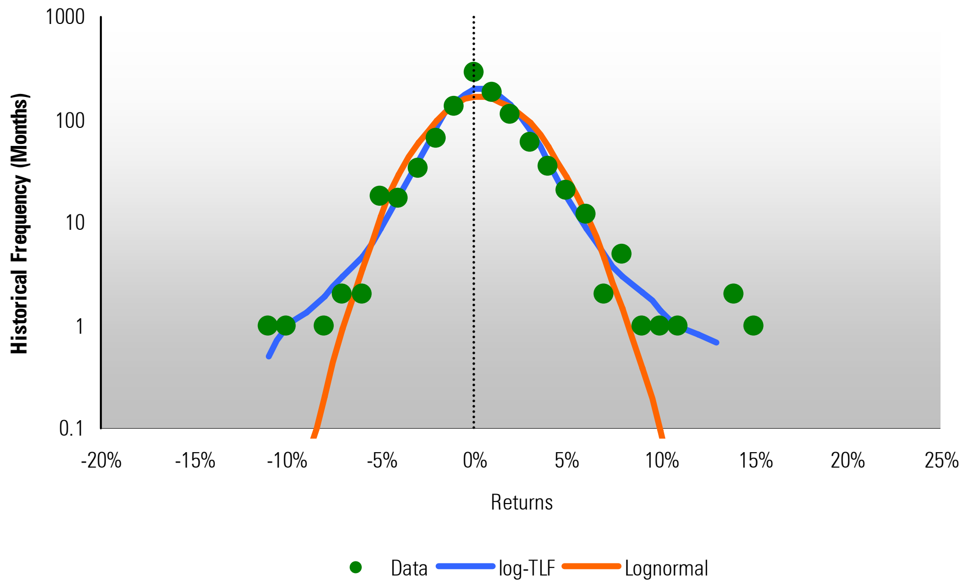

In summary, the log-TLF model is clearly superior to the lognormal model in describing the tail distributions. This is critical as risk estimations critically rely on the accuracy of the tail distributions.

Impact of Fat Tails on Portfolio Downside Risk

The return distribution models are critical on estimates of downside risk. The lognormal distribution model has thin left tail, and thus tends to underestimate the downside risk. A popular downside risk measure is value-at-risk (VaR), which is an estimate of the loss that we expect to be exceeded with a given level of probability (e.g. 5%) over a specified time period. Conditional value-at-risk (CVaR) is closely related to VaR and is derived by taking a weighted average between the value at risk (VaR) and losses exceeding the VaR. CVaR is also called the expected tail loss. Previous studies (e.g. Rockafellar and Uryasev, 2000) have shown that CVaR has more attractive properties than VaR, for example, CVaR is a coherent measure of risk. Therefore, we will choose CVaR as a measure of downside risk.

Armed with a better distribution model, we are ready for two common applications. The first one is to estimate downside risk for a portfolio, and the second one is to study the impact of fat tails on wealth accumulation.

We extend the univariate TLF model to a multivariate TLF model. We build three hypothetical portfolios: conservative (40% stocks / 60% bonds), moderate (60% stocks / 40% bonds), and aggressive (80% stocks / 20% bonds). Capital market assumptions are forecasted by Ibbotson Associates.

We generate a large sample of multivariate distributed returns (1,000,000) for the six asset classes, and then calculate the statistics for the three portfolios. The statistics are summarized in Table 3. We compare the portfolio return statistics under both log-TLF and lognormal distribution models.

Chart 3.: Historical Distributions of US Long-Term Government Bond Monthly Returns Fitted Log-TLF and Lognormal Models

Portfolios	**Mean**	**Std. Dev.**	**CVaR (Log-TLF)**	**CVaR(Lognormal)**	**CVaR diff.**
40/60 (Con.)	6.8%	10.4%	15.3%	11.8%	3.5%
60/40 (Mod.)	8.4%	14.9%	21.7%	17.2%	4.5%
80/20 (Aggr.)	10.0%	19.5%	28.7%	23.1%	5.6%

It is clear that CVaRs under the log-TLF distribution model are 3.5% ~ 5.6% higher than the corresponding CVaRs under the lognormal distribution model. The reason is that the log-TFL model has fatter tails so that the CVaRs are higher than that under the lognormal model. The difference in CVaRs for the two distribution models increases from the conservative to the aggressive portfolio because CVaR not only increases with the fat tail, but also with the portfolio's volatility.

From the risk management point of view, the results shown in Table 3 are important. The lognormal model can underestimate the CVaR or expected tail loss by as much as 5.6% for the aggressive portfolio, and thus it can mislead the risk budgeting process.

Impact of Fat Tails on Wealth Accumulation

We ran two sets of Monte Carlo simulations to study the impact of fat tails on portfolio's wealth accumulation: one assumes a lognormal distribution, and the other assumes a log-TLF distribution. Each simulation contains 10,000 simulated 30 year return scenarios. We also included results for the deterministic 6% growth for reference.

The simulated results are similar for the three portfolios, so we only report the results for the moderate portfolio. The wealth accumulation results for the moderate portfolio are shown in Table 4 and Chart 4. Both log-TLF and lognormal distributions have almost the same wealth at the 50-th percentile, but the difference in wealth at the 1-st percentile is significant. This is intuitive since the log-TLF distribution has a fatter tail and thus a larger downside risk or tail risk.

Table 4: Wealth Accumulation for the Moderate Portfolio Under both Log-TLF and Lognormal Models

Year	**Log-TLF**			**Lognormal**			
	1st	5th	50th	1st	5-th	50th	Constant 6%
	$1	$1	$1	$1	$1	$1	$1
1	0.725	0.856	1.072	0.799	0.868	1.069	1.060
2	0.702	0.838	1.156	0.753	0.846	1.147	1.124
3	0.690	0.835	1.239	0.732	0.850	1.231	1.191
4	0.688	0.851	1.332	0.722	0.861	1.321	1.262
5	0.687	0.865	1.433	0.717	0.868	1.415	1.338
6	0.672	0.884	1.539	0.732	0.890	1.516	1.419
7	0.679	0.901	1.645	0.733	0.918	1.625	1.504
8	0.695	0.935	1.775	0.743	0.947	1.747	1.594
9	0.697	0.952	1.911	0.752	0.982	1.876	1.689
10	0.726	0.986	2.046	0.750	1.014	2.013	1.791
11	0.736	1.022	2.202	0.765	1.057	2.164	1.898
12	0.731	1.073	2.375	0.797	1.108	2.322	2.012
13	0.772	1.116	2.547	0.809	1.137	2.506	2.133
14	0.799	1.157	2.740	0.838	1.184	2.697	2.261
15	0.810	1.206	2.962	0.839	1.232	2.900	2.397
16	0.856	1.260	3.189	0.884	1.280	3.093	2.540
17	0.880	1.311	3.423	0.924	1.336	3.349	2.693
18	0.923	1.380	3.702	0.967	1.386	3.593	2.854
19	0.962	1.457	3.942	1.008	1.457	3.867	3.026
20	0.984	1.522	4.235	1.022	1.521	4.159	3.207
21	1.009	1.595	4.546	1.053	1.569	4.478	3.400
22	1.010	1.686	4.903	1.125	1.670	4.795	3.604
23	1.070	1.755	5.249	1.143	1.742	5.133	3.820
24	1.118	1.818	5.661	1.211	1.836	5.498	4.049
25	1.144	1.897	6.091	1.275	1.940	5.924	4.292
26	1.193	2.004	6.510	1.285	2.055	6.321	4.549
27	1.235	2.130	7.044	1.363	2.121	6.830	4.822
28	1.259	2.208	7.596	1.399	2.248	7.321	5.112
29	1.320	2.304	8.160	1.479	2.410	7.928	5.418
30	1.363	2.431	8.782	1.527	2.487	8.475	5.743

Chart 4: Wealth Accumulation for the Moderate Portfolio with Lognormal and Log-TLF Distributions

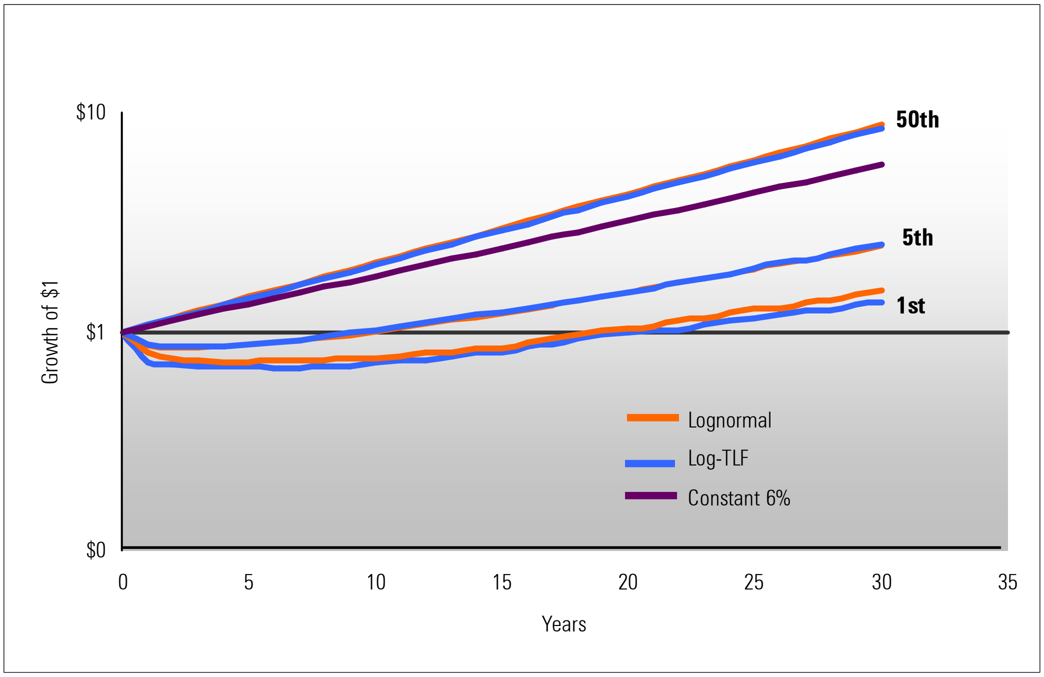

The first implication from Table 4 is that, at the 1-st percentile, the moderate portfolio can lose 27.5% of the total value in one year under the log-TLF model, and lose 20.1% under the lognormal model. In other words, in one out of 100 years, a moderate portfolio can lose 27.5% of total value in one year, and a lognormal model can underestimate that one-year loss by 7.4%.

Put slightly differently, for a moderate portfolio, our Monte Carlo simulations show that it takes about 40 years to lose 20% for the log-TLF model, and it takes about 100 years to lose 20% for the lognormal model. To test this with empirical data, we construct a simple moderate portfolio, 60% S&P 500 and 40% BarCap Aggregate Bond. [1] The ten worst performances for this moderate portfolio are shown in Table 5. It can be seen that the constructed moderate portfolio lost more than 20% in three years: 1931, 1937, and 2008. Thus the likelihood to lose 20% for a moderate portfolio is about three times in 83 years. The estimate from the log-TLF model (two times in 80 years) is much closer to the empirical observation than that from the lognormal model (one time in 100 years).

Table 5.: Worst 10 Annual Returns for the 60% Stocks and 40% Bonds Portfolio from 1926 to 2008

	Moderate Portfolio
Dec 1931	-26.93%
Dec 1937	-20.39%
Dec 2008	-20.10%
Dec 1974	-13.61%
Dec 1930	-12.25%
Dec 2002	-9.16%
Dec 1973	-6.95%
Dec 1941	-6.75%
Dec 1969	-5.40%
Dec 1940	-4.68%

Source: Morningstar Encorr

[1] The BarCap Aggregate Bond is backfilled with U.S. intermediate government bond from 1926 to 1975.

The second implication from Table 4 is that, at the 1-st percentile, the moderate portfolio can lose up to 33% in the first six years for the log-TLF distribution model, while the highest loss is 28% in the first six-years for the lognormal distribution model. These results have important implications for investors who are six years away from retirement: one can potentially lose one third of the total wealth or value in a moderate portfolio with 1% probability. Therefore a principal hedge against this downside risk seems prudent. Such kind of hedge can be implemented by portfolio insurance products, for example, an appropriately priced equity linked certificate deposit (ELCD) with a maturity of six years. Others include recent innovative guaranteed riders such as guaranteed minimum withdrawal benefits (GMWBs).

Conclusions

We show that the log-TLF model is more appropriate in estimating the downside risk than an alternative lognormal distribution model, as evidenced by the fact that the log-TLF model fits the entire distribution of historical monthly returns very well. Our estimates show that the lognormal model underestimates the *annualized* CVaR by 7.86% for S&P 500, and 1.66% for U.S. long-term government bond.

The fat tails have further impact on a portfolio's downside risk and wealth accumulation. In general, diversified portfolio's annualized CVaRs under the log-TLF distribution model are 3.5% ~ 5.6% higher than that under the lognormal distribution model. As a result, the lognormal model can mislead the risk budgeting process.

Finally, our Monte Carlo simulations indicate that one can lose one third of the total value in six years with 1% probability in a moderate portfolio. Therefore a principal hedge against this downside risk seems prudent for a pre-retirement investor.

References

Bachelier, Louis, 1900. Théorie de la Spéculation. Doctoral dissertation. Annales Scientifiques de l'Ecole Normale Supéieure (ii) 17, 21-86. Translation: Cootner 1964. Cootner, Paul H., ed. 1964. The Random Character of Stock Market Prices. Cambridge , MA: MIT Press.

Blattberg, R.C., and N.J. Gonedes. (1974). "A comparison of the stable and student distributions as statistical models for stock prices." Journal of Business, 47, 244-280.

Clark, P.K. (1973). "A subordinated stochastic process model with finite variance for speculative prices." Econometrica 41, 135–155.

Fama, E. F., 1965. "The Behavior of Stock-Market Prices." The Journal of Business, Vol. 38, No. 1, pp. 34-105.

Kaplan, P.D., (2009). "Déja Vu All Over Again", Morningstar Advisor, Fed/March, 29-33.

Lévy, P. 1925. Calcul des probabilités (Gauthier-Villars, Paris).

Mandelbrot, B. (1963). "The variation of certain speculative prices." Journal of Business 36, 392–417.

Mantegna, R.N., H.E. Stanley, (1999). "An Introduction to Econophysics: Correlations and Complexity in Finance." Cambridge University Press, Cambridge.

Rockafellar, R.T. and S.Uryasev. (2000). "Optimization of Conditional Value-At-Risk." The Journal of Risk, Vol. 2, No. 3, 2000, 21-41.

Xiong, J.X., (2009). "Using Truncated Lévy Flight to Estimate Downside Risk." Working paper.

About Ibbotson

Ibbotson Associates is a leading independent provider of asset allocation, manager selection, and portfolio construction services. The company leverages its innovative academic research to create customized investment advisory solutions that help investors meet their goals.

For more information, contact:
Ibbotson Associates
22 West Washington Street
Chicago, Illinois 60602
312 696-6700
312 696-6701 fax
www.ibbotson.com