Using truncated Lévy flight to estimate downside risk

Received (in revised form): 2nd February, 2010 and is available online at http://www.henrystewart.com/jrmfi.aspx

James X. Xiong
is a senior research consultant at Ibbotson Associates, a Morningstar company, in Chicago, Illinois.

Ibbotson Associates, 22 West Washington Street, Chicago, IL 60602, USA
Tel: +1 (312) 384 3764; Fax: +1 (312) 696 6701; E-mail: jxiong@ibbotson.com

Abstract It is well known that the normal distribution model fails to describe the fat tails of markets. The Lévy stable distribution model, meanwhile, has fat tails but leads to an infinite variance, thus complicating risk estimation. This study introduces truncated Lévy flight (TLF) — a better distribution model that has fat tails, finite variance, and more importantly, scaling properties. The paper uses TLF to estimate the downside risk of a variety of asset classes. The downside risk measure used is the conditional value-at-risk (CVaR). The study shows that the lognormal model can underestimate the monthly CVaR by 2.27 per cent for the S&P 500 Index, 0.48 per cent for the US Long-Term Government Bond Index, and 1.16 per cent for the MSCI UK Equity Index. Moreover, the paper extends a univariate TLF model to a multivariate TLF model to study the impact of fat tails on portfolios' downside risk and wealth accumulation.

Keywords: *truncated Lévy flight, downside risk, conditional value-at-risk, fat tail*

INTRODUCTION

The most common model of asset returns is assumed to be normally or Gaussian distributed.[1] 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-order approximation of the behaviour 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 an asset return that is three standard deviations below its arithmetic mean (a three-sigma event) has a probability of only ~0.13 per cent, ie once every 1,000 times. For example, from January 1926 to April 2009, the S&P 500 Total Return Index had a monthly mean return of 0.91 per cent and a monthly standard deviation of 5.55 per cent. Here, a negative three-sigma event would be a return lower than −15.74 per cent. During this time period, there were ten monthly returns worse than −15.74 per cent, as shown in Table 1. This implies that the probability of a three-sigma event is 1 per cent rather than 0.13 per cent, or eight times greater than one would expect under a normal distribution. Hence, a normal distribution fails to describe the 'fat' or 'heavy' tails of the stock market.

Many statistical models have been put forth to account for the heavy tails.

Table 1: The ten worst monthly returns for the S&P 500 (January 1926 to June 2009)

Date	S&P 500 (%)
Sep 1931	−29.73
Mar 1938	−24.87
May 1940	−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.

Well-known examples are Mandelbrot's Lévy stable hypothesis,[2] the Student's *t*-distribution,[3] and the mixture-of-Gaussian distributions hypothesis.[4] The latter two models possess finite variance and fat tails, but they are not stable, which implies that their shapes change at different time horizons and that distributions at different time horizons do not obey scaling relations.

An alternative is the Lévy stable distribution model.[5] In 1963, Mandelbrot modelled cotton prices with a Lévy stable process,[2] a finding later supported by Fama in 1965.[6] A Lévy stable distribution model has fat tails and obeys scaling properties, but it has an infinite variance that conflicts with empirical observations that the return variance is finite. The infinite variance would also complicate the task of risk estimation.

In the context of logarithm of asset returns (plus 1), the corresponding models are the lognormal and log-stable models. Figure 1 illustrates the log-stable and lognormal distributions in fitting the distribution of the monthly S&P 500 returns (see also Martin, Rachev and Siboulet;[7] and Kaplan[8]). The vertical axis of Figure 1 is in log scale with a base of 10, and this helps to view the tails of the distribution more clearly. It is clear that the lognormal distribution fails to fit the return distribution below −15 per cent (the abovementioned three-sigma event). The log-stable distribution fits the tail well, but it extends far beyond the historical maximum loss or gain with non-negligible probabilities, which eventually results in an infinite variance. In other words, the tail for the log-stable distribution is too fat.

What is needed here is a model with distribution somewhere between the

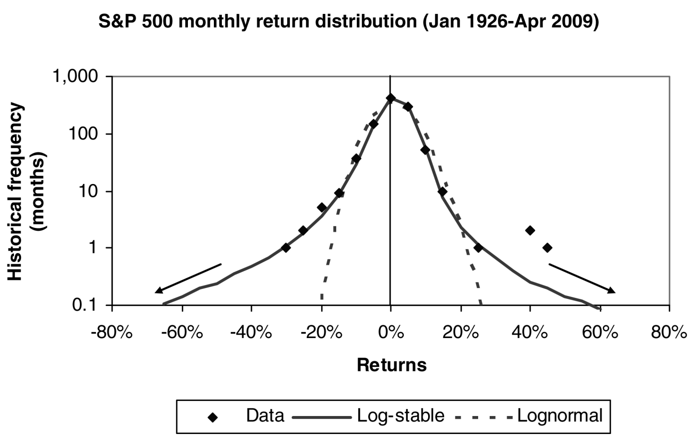

Figure 1: The historical distributions of S&P 500 monthly returns fitted by the log-stable and lognormal models

normal and stable distribution — that is, with an appropriately fat, but finite tail. The model in question is the truncated Lévy flight (TLF) model. In addition to having finite variance and fat tails, the model obeys Lévy stable scaling relations at small time intervals and converges slowly to Gaussian distribution at very large time intervals, which is consistent with empirical observations. The TLF model is discussed in detail below.

The return distribution models are critical to estimates of downside risk. The lognormal distribution model has a 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 one would expect to be exceeded with a given level of probability 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 VaR and losses exceeding the VaR. CVaR is also called the 'expected tail loss'. Previous studies (eg Rockafellar and Uryasev[9]) have shown that CVaR has more attractive properties than VaR; for example, CVaR is a coherent measure of risk, as proved by Pflug.[10] The present paper therefore uses CVaR as a measure of downside risk.

In the rest of the paper, the study first shows that a univariate TLF model fits monthly return distributions for a variety of asset classes represented by the S&P 500, Ibbotson US Long-Term Government Bond Index, and the Morgan Stanley Capital International (MSCI) UK Equity Index very well. To estimate the downside risk for a portfolio more accurately, the univariate TLF model is then extended to a multivariate TLF model.

TRUNCATED LÉVY FLIGHT

A TLF process is a stochastic process with finite variance and fat tails. The probability density function (PDF) of a simple TLF process is defined as:

$$P(x) = 0, \quad x < -l;$$
$$P(x) = P_{Levy}(x), \quad -l \le x \le l;$$
$$P(x) = 0, \quad x > l$$

where $P_{Levy}(x)$ is the PDF of return x for a Lévy stable distribution and l is the cutoff length for the truncation.

Mantegna and Stanley[11,12] show that, for a small time interval (eg a minute), the TLF distribution approximates a Lévy stable distribution with Lévy stable scaling; while for a significantly large but finite time interval (eg a year), the TLF distribution slowly converges to a Gaussian distribution. In other words, the TLF undergoes a crossover from a Lévy stable distribution to a Gaussian distribution as the time interval increases. This crossover is consistent with an independent empirical study of the distribution of daily, weekly and monthly returns for which a progressive convergence to a Gaussian process is deemed to be observed.[13]

Unfortunately, large time interval (such as annual) data are very limited. There are likely to be fewer than 100 annual data points for a market index, which makes the test of convergence to a Gaussian distribution statistically difficult.

Previous studies have mostly focused on high-frequency data such as daily return data. Here, the aim is to fit the monthly data by the log-TLF model because many investors have a relatively long investment horizon and portfolios are often rebalanced monthly.

Table 2 compares the log-TLF model with the lognormal model in fitting the return distributions for three asset classes represented by the S&P 500 (both daily and monthly), Ibbotson US Long-Term Government Bond Index (monthly), and MSCI UK Equity Index (monthly). All monthly data are from Morningstar Encorr. Daily data for the S&P 500, starting from 1950, are available from Yahoo.com.

The implementation of a univariate TLF model is shown in Appendix A. The lognormal model requires only two parameters — mean and standard deviation — while the log-TLF model needs five parameters — α, β, γ, δ and l. The first four are associated with the Lévy stable distribution as follows: α determines the tail weight or the distribution's kurtosis with $0 < \alpha \leq 2$; β determines the distribution's skewness; γ is a scale parameter; and δ is a location parameter. Meanwhile, l is the cutoff length introduced for the truncation.

The estimates of the five parameters are shown in Table 3. In Table 3, the cutoff length is selected by fitting the monthly CVaR to the historical monthly CVaR as shown in Table 2. The cutoff length shown in Table 3 is normalised. One can consider a normalised cutoff length of 6 as a six-sigma event. The estimated parameters α for daily and monthly S&P 500 are close to each other, which indicates the scaling property of the TLF model.

It can be seen that the tail measurements, such as CVaR and kurtosis, are much closer to historical values in the log-TLF model for all three asset classes. The lognormal model underestimates the *monthly* CVaR by 2.27 per cent for the S&P 500, by 0.48 per cent for the US Long-Term Government Bond Index, and by 1.16 per cent for MSCI UK Equity Index.

To observe the entire return distribution spectrum, Figures 2 and 3 compare the log-TLF model with the lognormal model in fitting the return distributions for S&P 500 daily and monthly returns, respectively. The far left data point in Figure 2 corresponds to a

Table 2: Statistics summary for historical returns, as well as simulated returns with both log-TLF and lognormal models

	Mean	SD	Skewness	Kurtosis	CVaR
S&P 500 daily					
Historical (1950–2009)	0.031%	0.96%	−0.69	26.4	−2.2%
Log-TLF	0.031%	0.96%	−0.50	9.3	−2.4%
Lognormal	0.031%	0.96%	0.02	3.0	−1.9%
S&P 500 monthly					
Historical (1926–2009)	0.91%	5.55%	0.36	12.49	−12.29%
Log-TLF	0.91%	5.55%	0.52	11.32	−12.29%
Lognormal	0.91%	5.55%	0.17	3.05	−10.02%
US Long-Term Govt Bond Index monthly					
Historical (1926–2009)	0.48%	2.37%	0.64	8.61	−4.80%
Log-TLF	0.48%	2.37%	0.17	5.98	−4.80%
Lognormal	0.48%	2.37%	0.07	3.00	−4.32%
MSCI UK Equity Index monthly					
Historical (1970–2009)	1.08%	5.86%	1.29	18.37	−11.61%
Log-TLF	1.08%	5.86%	0.80	11.56	−12.09%
Lognormal	1.08%	5.86%	0.17	3.06	−10.45%

SD, standard deviation; CVaR, conditional value-at-risk

Table 3: Parameter estimates with the log-TLF model for S&P 500 (daily and monthly), US Long-Term Government Bond Index (monthly) and MSCI UK Equity Index (monthly) returns

Log-TLF	α	β	γ	δ	Cutoff length
S&P 500 daily	1.50	−0.32	0.0041	0.0013	9.5
S&P 500 monthly	1.40	−0.12	0.024	0.010	6.3
US Long-Term Govt Index monthly	1.65	0.00	0.014	0.004	5.2
MSCI UK Equity Index monthly	1.43	0.00	0.026	0.009	6.3

daily loss of about 20 per cent on 19th October, 1987, 'Black Monday'.

Figure 4 compares the log-TLF model with the lognormal model in fitting the historical return distributions for the US Long-Term Government Bond Index. Figure 5 compares the log-TLF model with the lognormal model in fitting the historical return distributions for the MSCI UK Equity Index. It can be seen that the log-TLF model provides an excellent fit for the S&P 500, US Long-Term Government Bond and MSCI UK Equity indices in all aspects:

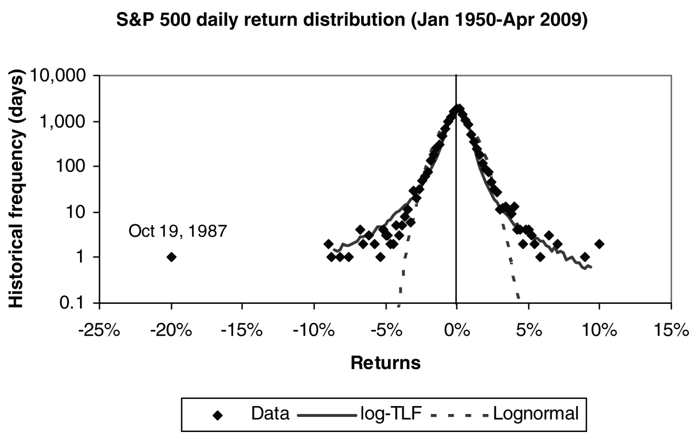

Figure 2: The historical distributions of S&P 500 daily returns fitted by the log-TLF and lognormal models

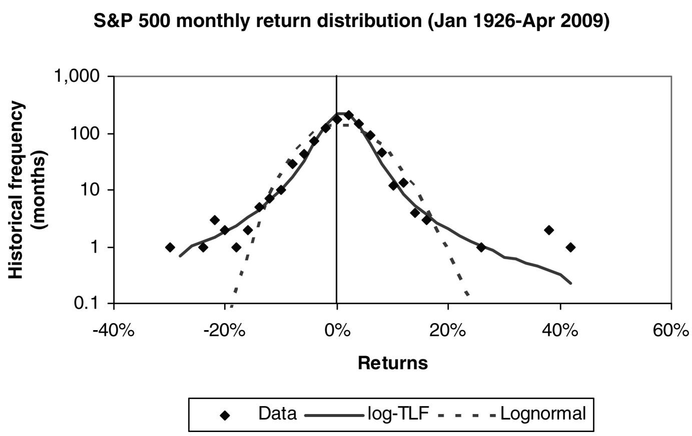

Figure 3: The historical distributions of S&P 500 monthly returns fitted by the log-TLF and lognormal models

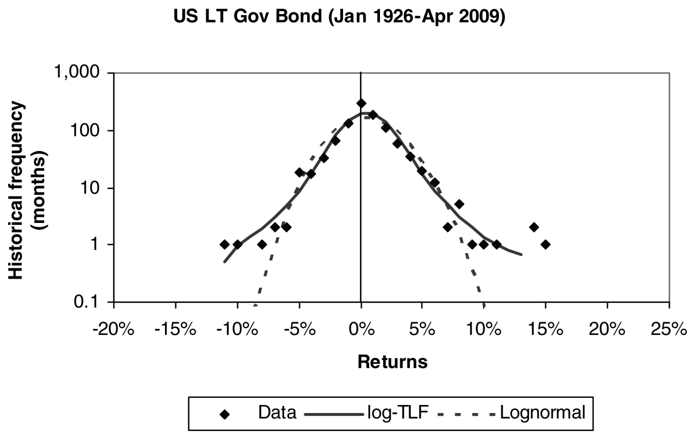

Figure 4: The historical distributions of US Long-Term Government Bond Index monthly returns fitted by the log-TLF and lognormal models

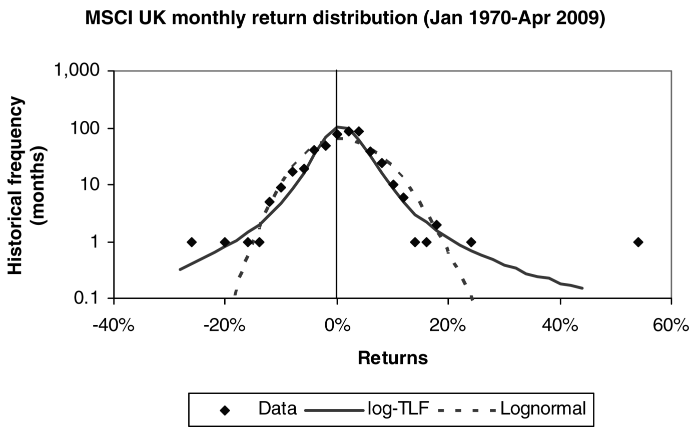

Figure 5: The semi-log historical distributions of MSCI UK Equity Index monthly returns fitted by the log-TLF and lognormal models

centre, tails, minimum and maximum. The only exception is that the Black Monday in Figure 2 is not well fitted by the log-TLF model.

The study further tested other equity and bond indices, and obtained similar results. This indicates the robustness of the log-TLF model, and it is universal among a variety of asset classes.

In summary, previous studies showed that the TLF model did a good job of describing both the asymptotic return distributions measured at different high-frequency time horizons and their scaling properties (self-similarities). Here, it is shown that the log-TLF model is not only capable of fitting the distribution of daily returns, but also monthly returns, and not only fitting stocks, but also bonds. The log-TLF model is clearly superior to the lognormal model in describing the tail distribution. This is critical as many financial applications rely on the accuracy of the distribution model.

On the other hand, like the normal or stable distribution model, the TLF model assumes that returns are independent and identically distributed (i.i.d.) and the volatilities of returns are constant. Therefore, it cannot describe the time-dependent volatility observed in market data.

IMPACT OF FAT TAILS ON PORTFOLIO DOWNSIDE RISK

Armed with a better distribution model, it is possible to test two common applications. The first is to estimate downside risk for a portfolio, and the second is to study the impact of fat tails on wealth accumulation. In both applications, it is necessary to simulate multivariate asset returns with fat tails.

The univariate TLF model is extended to a multivariate TLF model; the detailed methodology is shown in Appendix A. Note that the disadvantage of the multivariate stable or TLF distribution is that it requires the same tail index (α) for each marginal or individual distribution. In a typical portfolio, stocks and bonds tend to have slightly different tail indices as shown in Table 3. Nonetheless, stocks dominate bonds in the risk spectrum, therefore in most cases assuming the same tail index in a portfolio is a reasonable approximation.

Six simple asset classes are used, and three hypothetical portfolios built: conservative (40 per cent stocks/60 per cent bonds), moderate (60 per cent stocks/40 per cent bonds), and aggressive (80 per cent stocks/20 per cent bonds). Portfolio asset allocations and capital market assumptions are shown in Appendix B.

As this study focuses only on the impacts of fat tails, it is assumed that all of the asset classes are approximately symmetrically distributed. In other words, the information from skewness is ignored. A large sample of multivariate distributed returns (1,000,000) is generated for the six asset classes, and the statistics for the three portfolios calculated. The statistics are summarised in Table 3. Again, the portfolio return statistics are compared under both log-TLF and lognormal distribution models.

Table 4 shows that CVaRs under the log-TLF distribution model are 3.5∼5.6 per cent 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 under the lognormal model. The difference in CVaRs for the two distribution models increases from the conservative to the aggressive portfolio because CVaR is not only increased 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 per cent for the aggressive portfolio, and can thus skew the risk-budgeting process.

IMPACT OF FAT TAILS ON WEALTH ACCUMULATION

To study the impact of fat tails on wealth accumulation, the study ran two sets of Monte Carlo simulations: one assuming a

Table 4: CVaR estimates for the three portfolios

Portfolios	Mean (%)	SD (%)	CVaR (log-TLF) (%)	CVaR (lognormal) (%)	CVaR diff. (%)
40/60 (conservative)	6.8	10.4	15.3	11.8	3.5
60/40 (moderate)	8.4	14.9	21.7	17.2	4.5
80/20 (aggressive)	10.0	19.5	28.7	23.1	5.6

SD, standard deviation; CVaR, conditional value-at-risk lognormal distribution, the other a log-TLF distribution. Each simulation contained 10,000 simulated 30-year return scenarios. For reference, the study also included results for the deterministic 6 per cent growth. Appendix B presents the capital market assumptions.

The simulated results are similar for the three portfolios, so the paper reports only the results for the moderate portfolio. The wealth accumulation results for the moderate portfolio are shown in Table 5 and Figure 6. Both log-TLF and lognormal distributions have almost the same wealth at the 50th percentile, but the difference in wealth at the 1st percentile is significant. This is intuitive as the log-TLF distribution has a fatter tail and thus a larger downside risk or tail risk.

The first implication from Table 5 is that, at the 1st percentile, the moderate portfolio can lose 27.5 per cent of the total value in one year under the log-TLF model, and 20.1 per cent under the lognormal model. In other words,

Table 5: Wealth accumulation for the moderate portfolio under both log-TLF and lognormal models

	Log-TLF			Lognormal			
Year	1st	5th	50th	1st	5th	50th	Constant 6%
0	$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

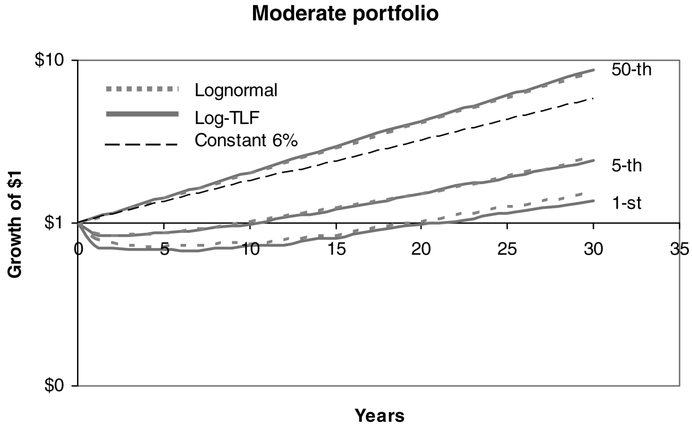

Figure 6: Wealth accumulation for the moderate portfolio with lognormal and log-TLF distributions

in one out of 100 years, a moderate portfolio can lose 27.5 per cent of total value in one year, and a lognormal model can underestimate that one-year loss by 7.4 percentage points.

Put slightly differently, for a moderate portfolio, the Monte Carlo simulations show that it takes about 40 years to see a 20 per cent annual loss for the log-TLF model, and it takes about 100 years to see a 20 per cent annual loss for the lognormal model. To test this with empirical data, a simple moderate portfolio is constructed — 60 per cent S&P 500 and 40 per cent BarCap Aggregate Bond. The BarCap Aggregate Bond Index is backfilled with US Intermediate Government Bond Index data from 1926 to 1975. The ten worst performances for this moderate portfolio are shown in Table 6. It can be seen that the moderate portfolio lost more than 20 per cent in three years: 1931, 1937 and 2008. Thus, the likelihood of losing 20 per cent for a moderate portfolio is about three times in 83 years. The estimate from the log-TLF model (twice in 80 years) is much closer to the empirical observation than that from the lognormal model (once in 100 years).

Table 6: The ten worst annual returns for the 60% S&P 500 and 40% BarCap aggregate bond portfolio from 1926 to 2008

Date	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.

The second implication from Table 5 is that, at the 1st percentile, the moderate portfolio can lose up to 33 per cent in the first six years for the log-TLF distribution model, while the highest loss is 28 per cent in the first six years for the lognormal distribution model. These results have important implications for investors who are six years away from retirement: there is a 1 per cent probability of losing one-third of the total wealth in a moderate portfolio. Therefore, a principal hedge against this downside risk seems prudent. An insurance product like an appropriately priced equity-linked certificate of deposit with a maturity of six years would be one example. Others could include guaranteed riders such as guaranteed minimum withdrawal benefits.

CONCLUSIONS

A large number of previous studies have shown that the TLF model does a good job of describing the asymptotic return distributions measured at high frequencies as well as their scaling properties (self-similarities). An important characteristic of the TLF model is that, for a small time interval (eg a minute), the TLF distribution approximates a Lévy stable distribution with Lévy stable scaling; while for significantly large but finite time intervals (eg a year), the TLF distribution slowly converges to a Gaussian distribution.

The log-TLF model is more appropriate for estimating the downside risk than an alternative lognormal distribution model, as demonstrated by the fact that the log-TLF model fits the entire distribution of historical returns very well for a variety of asset classes. The lognormal distribution model, however, does not. The estimates show that the lognormal model underestimates the *monthly* CVaR by 2.27 per cent for the S&P 500, 0.48 per cent for the US Long-Term Government Bond Index, and 1.16 per cent for the MSCI UK Equity Index.

The fat tails affect a portfolio's downside risk and wealth accumulation expectations. In general, a portfolio's annualised CVaRs under the log-TLF distribution model are 3.5∼5.6 per cent higher than under the lognormal distribution model. As a result, the lognormal model's lower risk estimates can skew the risk-budgeting process. Finally, the Monte Carlo simulations using a log-TLF distribution model indicate that there is a 1 per cent probability of losing one-third of the total value of a moderate portfolio in six years. Using a lognormal distribution, there is a 1 per cent probability of losing closer to a quarter of the value.

ACKNOWLEDGMENTS

The author thanks Peng Chen, Thomas Idzorek, Paul Kaplan and Alexa Auerbach at Morningstar for their helpful comments.

References

1 Bachelier, L. (1900) 'Théorie de la Spéculation', *Annales Scientifiques de l'École Normale Supérieure*, Vol. 17, No. 3, pp. 21–86. Translated by Paul H. Cootner (1964) 'The Random Character of Stock Market Prices', MIT Press, Cambridge, MA.
2 Mandelbrot, B. (1963) 'The variation of certain speculative prices', *Journal of Business*, Vol. 36, No. 4, pp. 392–417.
3 Blattberg, R. C. and Gonedes, N. J. (1974) 'A comparison of the stable and student distributions as statistical models for stock prices', *Journal of Business*, Vol. 47, No. 2, pp. 244–280.
4 Clark, P. K. (1973) 'A subordinated stochastic process model with finite variance for speculative prices', *Econometrica*, Vol. 41, No. 1, pp. 135–155.
5 Lévy, P. (1925) 'Calcul des probabilités', Gauthier-Villars, Paris.
6 Fama, E. F. (1965) 'The behavior of stock-market prices', *Journal of Business*, Vol. 38, No. 1, pp. 34–105.
7 Martin, R. D., Rachev, S. and Siboulet, F. (2003) 'Phi-alpha optimal portfolios & extreme risk management', *Wilmott Magazine of Finance*, Vol. 2003, No. 6, November, pp. 70–83.

8 Kaplan, P. D. (2009) 'Déjà vu all over again', *Morningstar Advisor*, Feb/March, pp. 29–33.
9 Rockafellar, R. T. and Uryasev, S. (2000) 'Optimization of conditional value-at-risk', *Journal of Risk*, Vol. 2, No. 3, pp. 21–41.
10 Pflug, G. C. (2000) 'Some remarks on the value-at-risk and the conditional value-at-risk', in Uryasev, S. (ed.) 'Probabilistic Constrained Optimization: Methodology and Applications', Kluwer Academic Publishers, Dordrecht, pp. 272–281.
11 Mantegna, R. N. and Stanley, H. E. (1994) 'Stochastic process with ultraslow convergence to a Gaussian: The truncated Levy flight', *Physical Review Letters*, Vol. 73, No. 22, pp. 2946–2949.
12 Mantegna, R. N. and Stanley, H. E. (1999) 'An Introduction to Econophysics: Correlations and Complexity in Finance', Cambridge University Press, Cambridge.
13 Akgiray, V. and Booth, G. G. (1988) 'The stable-law model of stock returns', *Journal of Business and Economic Statistics*, Vol. 6, No. 1, pp. 51–57.

APPENDIX A: UNIVARIATE OR MULTIVARIATE TLF DISTRIBUTION

First employ John Nolan's software Stable 5.1 to implement the univariate or multivariate Lévy stable distribution model, and then generate univariate or multivariate return numbers. For details, see: http://academic2.american.edu/~jpnolan/stable/stable.html. Secondly, apply a truncation method on these generated returns so that the return series follows a TLF model. These truncated returns are then used in distribution analyses and CVaR estimates, as well as Monte Carlo simulations.

The univariate returns are generated by the stable distribution functions provided by the software, and the main inputs to the function stablernd() are α, β, γ and δ. The α indicates the heaviness of the tails; the β is the skewness; the γ corresponds to volatility; and the δ corresponds to an expected return.

The truncation is simply implemented, for example, by truncating return (x) which is beyond 9.5-sigma for daily returns or 6.3-sigma for monthly returns.

For multivariate returns, six asset classes are used (see Appendix B). Assume that the return distributions are approximately symmetric. The multivariate returns are generated by the multivariate stable functions mvstableelliptical() and mvstablernd() in the software package, and the main inputs to the function are a positive definite covariation matrix, a location vector, and an α. The covariation matrix is similar to a variance-covariance matrix. The location vector corresponds to an expected return vector. The α indicates the heaviness of the tails.

APPENDIX B: CAPITAL MARKET ASSUMPTIONS

Table A1: Expected return and standard deviations for the six asset classes

	Exp. ret. (%)	SD (%)
Russell 1000 Growth	8.78	23.55
Russell 1000 Value	10.38	18.75
Russell 2000 Growth	11.03	33.59
Russell 2000 Value	16.03	27.70
BarCap US Agg Bond	4.36	6.79
BarCap Govt 1–3 Yr	3.06	4.15

SD, standard deviation

Table A2: Correlation matrix

	LG	LV	SG	SV	Aggr. bond	Short bond
LG	1.00	0.83	0.86	0.74	0.21	0.10
LV	0.83	1.00	0.72	0.84	0.26	0.15
SG	0.86	0.72	1.00	0.87	0.11	0.02
SV	0.74	0.84	0.87	1.00	0.19	0.09
Aggr. bond	0.21	0.26	0.11	0.19	1.00	0.91
Short bond	0.10	0.15	0.02	0.09	0.91	1.00

Table A3: Hypothetical asset class model portfolios

	Conservative (%)	Moderate (%)	Aggressive (%)
LG	10	15	20
LV	10	15	20
SG	10	15	20
SV	10	15	20
Aggr. bond	30	20	10
Short bond	30	20	10