Robust Properties of Stock Return Tails

Blake LeBaron *

September 2008

Revised: April 2009

Abstract

This paper explores the tail features of daily stock returns. Recently developed versions of the Hill estimator are used to measure the extreme positive and negative returns for a small set of individual daily stocks covering the period of 1926 through 2004. The findings report many of the accepted stylized facts about stock returns. Scaling exponents are reliably near 3, and generally stable over time and across positive and negative tails. A simple measure of tail behavior, the Gaussian crossing point, is introduced which gives further information on tail behavior including some intriguing results suggesting that positive tails may be fatter than negative ones.

JEL Classification: G12, C10, C32

Keywords: Tail exponents, Fat tails, Volatility

*International Business School, Brandeis University Blake LeBaron Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a research associate at the National Bureau of Economic Research.

1 Introduction

For almost half a century it has been known that the distributions of relatively high frequency asset returns do not follow a Gaussian distribution.[1] The descriptive, and technical term "fat tails" has made it into the common language of finance. It refers to the fact that there are more extreme returns, both large and small, than would be predicted by a normal distribution. The exact causes for these features are still not completely understood, even though the features have made their way to becoming a financial "stylized fact". This paper is concerned with the general robustness of the shape of distribution tails.

Given that stock returns are not normal, the properties of the unconditional distribution are important, and distribution tails are particularly interesting. Measures of tail properties are important for several reasons. First, for researchers calibrating to these features, they give them a quantitative target which is more challenging, and interesting than simply getting nonnormal return distributions. Second, the shape of the tail parameter gives us important information about the existence of higher moments in return series. Unstable, or nonexistent, higher moments can cause problems for estimating other parameters, or various measures of risk. Third, estimates of tail shape can be used for better risk estimation, since they provide information on tail probabilities. Finally, tail shape also can connect various risk measures as in expected tail loss and VaR.[2]

Many studies of financial returns provide only rudimentary estimates of tail shapes. Usually a single statistic, kurtosis, is reported without any standard errors. Very few papers use kurtosis in any more detailed fashion than as evidence against the standard normal distribution. Also, it may be unstable, or meaningless, if fourth moments don't exist. Fortunately, there are other measures of tail behavior which are more informative, and don't suffer these problems. Specifically, the large area of extreme value theory, provides a theoretical and empirical backing for different estimates of tail behavior. This body of theory shows that tail behavior may be very similar across many different physical and social processes, and can be well described by a power law, and a scaling exponent.[3] Unfortunately, estimating these exponents is still a tricky process, and has kept this from becoming a common tool. Recently, new methods have been developed which make estimating tail exponents more reliable. Specifically, the bias adjusted estimator developed in Huisman, Koedijk, Kool & Palm (2001) will be used extensively.

[1] The classic papers on this are Mandelbrot (1963) and Fama (1965).

[2] See (Chistoffersen 2003) for an example.

[3] It is important to note that it is not necessary that the tails follow a power law. This paper avoids the general issue of whether there is, or isn't a true power law in the data. This issue still seems very difficult to empirically test. Brock (1996) points out some of the problems in trying to empirically test for the presence of power laws. Clauset, Shalizi & Newman (2009) is a recent paper involved in testing the reliability of power laws across many diverse data sets. For a recent survey of power laws and their potential mechanisms in economics and finance see Farmer & Geanakoplos (2008).

Section 2 introduces several of the tail estimators used in the paper, and provides monte-carlo analysis. Section 3 applies these estimators to the stock return series, and section 4 concludes.

2 Tail estimators

2.1 OLS Hill Estimator

Extreme value theory provides a strong justification for the estimation of certain tail parameters. It states that both the distributions of maximums and conditional tails drawn from all densities will fall into certain classes of limiting distributions.[4] This paper concentrates on the peaks over threshold application of the theory. This deals with conditional probabilities such that a value exceeds a given threshold, u. This implies that given we are far enough out in the tail, stock returns will follow,

$$P(r_t < x | r_t < u) = k|x|^{-\alpha}. \tag{1}$$

The parameter α, the tail exponent, is critical, since it determines the rate at which the probability density drops off as one moves out into the tail.[5] This above probability is applied to the left tail, but could equally be applied to the right tail as well. α is referred to as the tail exponent. Often it is connected to $\gamma = 1/\alpha$ which is known as the shape parameter.

It is clear that the above relationship would give a power law scaling relationship in a log/log transformed plot of extreme returns, and log probabilities with slope $-\alpha$. Knowing α will reveal detailed information on tail probabilities, and therefore a precise risk measure out in the tail. Also, it provides information about the existence of higher moments.[6]

Unfortunately, estimating α is not an easy task, although there are many estimators in the exteme value literature. The most common of these is the Hill (1975) estimator, which is given by,

$$\hat{\gamma}_{n_k,n} = \frac{1}{n_k} \sum_{i=1}^{n_k} (\log x_{(n-i+1)} - \log x_{(n-n_k)}) \tag{2}$$

[4]The literature on extreme value theory is extensive. Books such as Embrechts, Kluppelberg & Mikosch (1997) and de Haan & Ferreira (2006) offer good surveys of the field. A short recent introduction is also contained in in Alfarano & Lux (Forthcoming). Recent applications to stock returns are Longin (2005) and Brooks, Clare, Molle & Persand (2005). Earlier return estimates are contained in Longin (1996),Loretan & Phillips (1994), and Lux (1996).

[5]This assumes that the family of extreme value distributions falls into the Frechet type. It is possible that there is no power-law in the tail if the distributions were of the Gumbel type. This paper is taking the position of explicitly assuming there is a power law tail of the Frechet type.

[6]All moments $E(r_t^m)$ will not exist for $m \geq \alpha$. An interesting question is what the small sample properties are for moments which technically exist, $m < \alpha$, but with m near the α boundary.

where $x_{(i)}$ are the order statics for the series, $x_{(n)} < x_{(n-1)}, \ldots, x_{(1)}$. The sample size is n, and $n_k < n$ refers to the number of entries in the tail being considered. Also, let k be the tail fraction where $n_k = kn$. [7] The crucial parameter here is n_k, which is the number of entries in the tail to be used. The Hill estimator is easy to implement, but it depends critically on n_k. Unfortunately, in practice it turns out to be very sensitive to the choice of n_k, and there is a large literature that tries to find an "optimal" value of n_k for a given sample.[8] That the choice of n_k, or the tail scaling region, is important should be intuitive. Chose n_k too large, and you have moved away from the extreme value scaling region of the distribution. The theory provides no information about the shape as one moves far into the distribution, and the power law relationship becomes irrelevant. It is well known that the Hill estimator is biased in this situation, and the bias increases as one moves deeper into the distribution. This pushes one to move out into the extreme tail, but now one faces the problem of choosing n_k too small. Now the problem is one of variance in the estimate, since the sample sizes are approaching zero. This is a classic bias/variance trade off problem.[9]

An important result from research on Hill estimators is that asymptotically, the bias can be shown to be a function of k. Specifically,

$$E(\hat{\gamma}) = \gamma + (k)^{\rho} \tag{3}$$

where ρ is another parameter to be estimated.[10] In an important paper, Huisman et al. (2001), assume that $\rho = 1$, giving the approximate bias as a linear function of k. In small sample cases close to financial data this appears to be a reasonable approximation. They then recommend the following strategy. Obtain Hill estimates of γ for several different values of k, $\hat{\gamma}(k_j)$. Then regress these on k as in,

$$\hat{\gamma}(k_j) = a + bk_j + e_j. \tag{4}$$

The linear term, b will estimate the bias, and a will give a bias corrected estimate of γ.[11] The tail exponent estimate would be given by $\hat{\alpha} = 1/a$.

[7] This would refer to the left tail, but it could be reversed to the right tail.

[8] The most automatic methods involve the use of the subsample bootstrap. Good examples of this are Hall (1990), Danielson, de Haan, Peng & de Vries (2001), Dacorogna, Muller, Pictet & de Vries (1995), Draisma, de Haan & Peng (1999), and Gomes & Oliveira (2001).

[9] Most procedures for finding an optimal n_k are based on searching for the optimal estimator in mean squared error, MSE, around the true value. This gives an explicit weighting for bias and variance.

[10] See Hall (1990) and Dacorogna et al. (1995) for derivations.

[11] A related line of research attempts to estimate the bias by estimating the parameter ρ, and then using a bias corrected version of the Hill estimator. An example of this approach is Gomes & Pestana (2007) and Caeiro & Gomes (2008).

Another fact about the Hill estimator is that the asymptotic variance is approximately,

$$\sigma^2(\hat{\gamma}) \approx \frac{\gamma^2}{n_k}. \tag{5}$$

Since for a given sample, n, this implies that the errors in the above regression fall in proportion to k. This will then yield some improvement in efficiency by moving to a weighted least squares estimator with a weight of $\sqrt(k_j)$ given to observation j. (Huisman et al. 2001) provide many examples and experiments showing the improved properties of this estimator versus traditional Hill estimators. A few more examples will be provided here, but the basic message is the same. This appears to be a much better method than the simple Hill approach in small samples. This will be referred to as the OLS Hill estimator throughout the paper.[12]

2.2 Gaussian crossing point

Most tail estimators attempt to give a relatively detailed picture of the tail of a distribution. In this section a simpler estimator is introduced which gives up some of the detail of traditional measures, but hopes to gain some in robustness, simplicity and general usability. Most discussion of fat tailed distributions naturally concentrates on the tail and the large amount of density far from the central part of the distribution. This ignores the fact that there is often more density near the mean of a fat tailed distribution than in a normal, just as there is more mass in the tails. Qualitatively, this means that just as there are more very small, and very large stock returns than predicted by a normal distribution, there are also more returns close to the mean.

To be more concrete, consider the following probability,

$$F(x) = P(r_t < x), \tag{6}$$

where $F()$ is the true unconditional probability distribution of the stock return. For simplicity, assume the distribution is symmetric with $E(r_t) = 0$. Also, define another comparison measure based on the Gaussian distribution with the same expected value and variance as the stock returns,

$$F_G(x) = P_G(r_t < x). \tag{7}$$

[12] All implementations will use the weighted least squares approach. Some monte-carlo experiments have been done, and it appears the efficiency gain is not that large, but since this is easy to implement it will be done.

Consider moving out into the left tail from $x = 0$. For most stock return distributions the Gaussian probabilities will drop off slower from 0.5 at first,

$$F(x) < F_G(x). \tag{8}$$

In words, this means that if you consider the tail starting at a relatively small negative return, then you would consider the Gaussian fatter than the true empirical distribution. For fat tailed distributions this must change eventually, and for a small enough x we will have

$$F(x) > F_G(x) \tag{9}$$

reflecting the well known large probability mass out in the tail that characterizes "fat tails". Define the crossing point as, x_c, where $F(x_c) = F_G(x_c) = p_c$. p_c corresponds to the quantile probability at the crossing point. This quantile crossing level is a probability description of where the "fat" part of the distribution begins.

This provides a cruder, possibly more robust, picture of a distribution tail. For example, say one estimates $p_c = 0.02$. This would say that if one were to use Gaussian approximations to the distribution to estimate quantiles or VaR estimates, then if you are interested in tail probabilities greater than 0.02 your risk measures will actually be too conservative, as the estimated losses would be too large. Once you move out beyond 0.02 into the outer tail, your Gaussian risk measures would under estimate the risk. It is clear this measure works equally well on either tail of the distribution, and can give further information on symmetry properties in the tails.

2.3 Monte-carlo experiments

Monte-carlo testing requires putting some weight on a null hypothesis for the tail distribution of stock returns. This has been the subject of debate for many years, and there is no agreement on the subject. However, unconditional tail behavior is increasingly modeled as a Student-t distribution. Figure 1 provides a graphical defense that this is a reasonable family. It plots the probability of $|r_t| > x$, the pooled tail probability, in log/log space for daily General Motors stock returns from 1926-2004. For a power law distribution this would yield a straight line. The stock returns are marked by stars and show a nearly linear relationship. They favorably compare with three student-t distributions with degrees of freedom of 3, 4, and 5.

Figure 2 examines the properties of the Hill estimator applied to a single sample of length $20,000$ drawn from a student-t distribution with three degrees of freedom. The value of the tail index for this distribution is 3. The figure plots the tail index on the y-axis, and the fraction of the tail on the x-axis. The solid line corresponds to the Hill estimator, and is typical of what one sees for tail estimation in financial series. For small values of the tail fraction, k, the estimate is unstable, reflecting the small sample sizes. As k increases the variability diminishes, but the bias begins to take over, and the value drops well away from the true value. The dashed line repeats the same experiment for the OLS Hill estimator.[13] It is much more stable near the true value, and does not appear to show much of a strong bias as k increases.

Figure 3 repeats the last figure using daily GM returns. The Hill estimator is again very unstable for low values of k, then it appears to quickly stabilize, but for slightly larger values of k, the bias takes over, and the estimator drops of dramatically as k increases. A similar experiment is performed with the OLS Hill estimator. In contrast to the traditional Hill this shows a long and stable region for the estimated tail index. For both the simulated and actual data it appears that the OLS Hill estimator gives a much larger region of k values yielding reasonable estimates. This is in contrast to the Hill estimator, where the procedure for selecting the optimal k will need to be very precise.

The next two figures, 4 and 5, perform a monte-carlo simulation with 1000 replications of student-t samples of length $20,000$ with three degrees of freedom. This distribution has a known tail exponent of 3, and therefore a shape parameter of 1/3. Figure 4 estimates the small sample bias for the two estimators for varying tail fractions, k. This is done for the shape parameter, γ, since this is the object of the bias correction. The increase in the bias for the Hill estimator is clear and very dramatic as k increases. Also, for smaller values of k the linear approximation for bias appears to be reasonable, which gives defense to the approximation in equation 4. The OLS Hill estimates appear to be doing a reasonable job of reducing the bias even for large values of k. Figure 5 repeats the experiment for the estimated variance. The solid line shows the Hill estimator while the dashed shows the OLS Hill estimator. Both drop off as k increases, and there is evidence for a slightly lower variance for Hill estimator.

Figure 6 estimates the mean squared error of the shape parameter estimator using the same monte-carlo

[13] For all runs in this paper the OLS Hill is estimated using a range of Hill estimates from 5 to $n_k = nk$ incremented by 1 using equation 4. Regressions all start at $j = 5$, and are run to $j = n_k$ incremented by one, using all possible Hill estimators in the sample range. The experiments are all concerned with varying the endpoint n_k, the largest cutoff point used in the regression. Changes in these parameters have very little impact on the estimated values.

experiment with a student-t distribution with 3 degrees of freedom. The mean squared error is estimated as,

$$MSE = \frac{1}{M}\sum_{i=1}^{M}(\hat{\gamma}_i - \gamma)^2 \tag{10}$$

where $\hat{\gamma}_i$ is the appropriate estimate, and $\gamma = 1/3$ is the true value. The minimum for the Hill estimator is obtained for a very small value of $k = 0.005$. For the OLS Hill the minimum is at $k = 0.21$. The y-axis uses different scales for the MSE, with the OLS Hill on the right axis. The estimated minimum MSE for the Hill estimator is $1.1x10^{-3}$, and for the OLS Hill it is $1.8x10^{-4}$ which corresponds to a reduction in MSE by a factor of 5 by moving to the OLS Hill.

It is interesting to note that this optimal k for the OLS Hill is much farther into the distribution than is usually recommend for optimal Hill estimates. At this point, the bias of the Hill estimator is very large, and swamps the MSE. Optimal levels for k for Hill estimators tend to be less, than 0.05, but this depends on the sample size. This ability to go far into the distribution is what makes the OLS Hill estimator feasible for smaller sample sizes.

Optimal values of k for most tail estimators are often very sensitive in terms of MSE. A small deviation from the optimal value can have disastrous effects in terms of MSE estimates. Also, the optimal k is sensitive to the sample size. Figure 7 shows an estimate of the optimal k using the OLS Hill estimator, and the previous Student-t monte-carlo run. For these runs the sample size is varied from 2000 through $20,000$. The figure shows that the optimal value of k, the solid line, is not very sensitive to the sample size. Beyond the optimal value it also does a kind of robustness check. It reports the values of k for which the MSE level is 25 percent greater than the optimal. In other words, it is asking if you were willing to deal with a slightly less optimal estimate how close to the optimal k do you have to be? These values are given by the dashed lines. It is interesting that there is a very wide range of k values for all samples sizes which would give you a reasonable (in MSE terms) estimator. This is very important evidence that the OLS Hill estimator, unlike the traditional Hill estimator, is not that sensitive to the choice of k.

Evidence presented in this section suggests that using the Huisman et al. (2001) estimator with $k = 0.2$, or $k = 0.21$ will give a relatively stable and reasonable picture for the tail index of stock returns for sample sizes that are reasonably large, 2000 to $20,000$. The next section will apply this to a set of individual daily stock returns.

3 Stock return tails

3.1 Tail Exponent

In this section the tail index estimates are explored, as well as estimates of the Gaussian crossing points. In both cases symmetry in the return distributions are considered along with stability over the long range sample. All return distributions use daily log returns, with dividends, from the CRSP data set from January 1st, 1926 through December 31st, 2004.[14] This gives a full sample size of $N = 21016$ daily returns. Table 1 shows the list of 14 firms and some basic summary statistics. The firms were chosen for their long life, allowing comparisons over the entire period. There are obvious survivorship biases in doing this, so some caution is warranted. Also, there are important questions about whether these firms have changed their character over time. The latter criticism is interesting, but correcting for it would involve replacing various names with others. The final item, VW Index, is the daily value weighted index with dividends from the CRSP dataset. The values in the table are typical for most asset returns. There appears to be some slight evidence for a negative skew, but this is far from universal. The key feature is the excess kurtosis which is is clear for all the firms in the sample. These values are well beyond the value of 3 from a Gaussian distribution.

Table 2 gives estimates for the tail index on the individual returns using the OLS Hill estimator evaluated at $k = 0.15, 0.20, 0.25$ fractions. The OLS estimate uses Hill estimates from 5 to $n_k = nk$ incremented by 1. From the monte-carlos these appear to be near optimal values for k. The table shows that the estimates are generally not sensitive to the choice of k. The last column in the table presents a bootstrap estimate of the standard error of the estimated tail index.[15]

The values in the table provide a uniform picture which agrees with the literature on tail estimation. First, the indices are all comfortably greater than 2, which indicates the existence of means and variances. Second, they are generally in a a region near 3. For moment existence the values suggest that one should be wary of estimated 4th and higher moments, and some caution is probably also warranted for third moments as well. The precision of the estimates is not bad, with standard errors near 0.10. Another interesting feature is that while the estimates are in the same general range, they probably would be statistically different across firms. There is no long range convergence to a universal distribution for this set of stocks which is interesting. Table 3 repeats the estimates for the right tail. Results are similar to the left tail. There is some indication

[14] The time horizon ends at the end of 2004 to keep Sears in the data set.

[15] For this case, the bootstrap is implemented as an independent, identically distributed (IID) null hypothesis where the returns are redrawn with replacement. 1000 bootstrap replications are performed.

of higher values indicating a thinner tail for positive returns.

Table 4 performs a bootstrap simulation replicating a symmetric null hypothesis. The returns are sampled with replacement as in a traditional bootstrap, but they are then multiplied by 1 or -1 with probability of 0.5 on each. In other words the signs are flipped randomly for each return. Under a null hypothesis of symmetry about zero, this will have no impact on the the distribution.[16] The table repeats the tail index for the left and right tails, and then reports the difference. Remember, a larger tail index corresponds to a thinner tail. A negative difference would indicate that the left tail is fatter than the right tail. The table shows that for most of the firms this difference is negative. The final column corresponds to the symmetric bootstrap result. It reports the probability that the observed difference in the data is greater than that from the bootstrapped distribution. A small value indicates that few of the bootstrap samples had a difference as small as that in the data, and that it would not be likely to see the asymmetry from the data in the simulated null hypothesis. The table shows that most of these simulated p-values are actually quite large. Only two are less that 5 percent. One important exception is the value weighted index which shows a significantly fatter right tail, than the left, with a large p-value of 0.98. From the standpoint of tail exponents the evidence for asymmetry is generally not significant for the individual firm returns. One implication of this result is that it allows one to reasonably pool the left and right tails when estimating tail exponents, leading to greater precision.

3.2 Gaussian crossing point

The estimates for the Gaussian crossing points are given in table 5 for both the left and right tails along with the difference and the p-value for the difference under the bootstrap, assuming a symmetric null hypothesis. The values reported are the probability level at the crossing point. Using the previous notation, these values correpond to p_c. For example, for ATT the left tail crossing probability is at $p_c = 0.026$. This shows that the empirical probability distribution function will cross the Gaussian out at the 0.026 quantile level in the left tail. In other words, the tail "fatness" for the distribution begins at the 0.026 quantile level. A similar number is reported for the right tail, of 0.032. This indicates that the part of the distributions where empirical probabilities exceed those from a Gaussian appears sooner as one moves into the right than the left tail. The difference between these two values is significant as reported by the p-value of 0.036. The table shows that 8 of the firms report a p-value on this difference of less than 0.05, indicating a common pattern

[16] For examples of this sort of test see Brock, Hsieh & LeBaron (1991) chapter 3 where it is referred to as a sign scrambling test. Also, see LeBaron & Samanta (2004), Lisi (2007) and Perez-Alonso (2007).

of asymmetry from the left to the right tail. It is interesting that this evidence points in the direction of a fatter right tail which is not something that is believed to be common in financial time series.[17] The VW index again reverses this result, displaying a crossing point which is much larger on the left tail, than on the right. This may correspond to the casual intuition that there are more extreme returns on down days, than on up days.

The actual magnitudes of these estimates is also important. They are all in the range of $0.02 - 0.03$, indicating that the crossing points are generally far out into the tail of the distributions. This means that risk managers using Gaussian distributions will only underestimate risk once they've gone out into the 2 percent tail areas of the return distributions.

3.3 Temporal Stability

The long sample sizes used in this study allow for the examination of the temporal stability of these measures. The table 6 displays estimated left tail exponents using the first and last 5000 days in the sample.[18] The bootstrap is simulated by drawing two 5000 length samples from the entire distribution with replacement. The table shows surprisingly few patterns across the individual stock returns. Some appear to have larger tail exponents (thinner tails) during the first part of the sample, but others are the reverse. The VW index is a major exception showing a strong tendency toward fatter tails in the early part of the sample. This could be due to the fact that the index contains a much larger cross section of firms in the later sample. For the individual firms there is certainly no general trend that fat tails have gone away, and the overall evidence for changes in the tail characteristics over time appears weak.

3.4 Volatility and tails

One of the most reliable and well known features of stock return series is that they are not independent over time, and it is well known that conditional variances, or volatility, are changing through time.[19] Returns should be thought of as following a process given by,

$$r_t = \sigma_t \epsilon_t \tag{11}$$

[17] Corroborating evidence on individual firms is in Jiang & Yao (2007).

[18] These correspond to the periods of approximately 1926-1942 and 1985-2004 in the data.

[19] The origins of this are in Mandelbrot (1963), and more formal modeling begins with Engle (1982) and Bollerslev (1986). There are many surveys of this literature. A recent one is Andersen, Bollerslev, Christoffersen & Diebold (2005).

as a first approximation. Here, ϵ_t is some independent, identically distributed random variable, and σ_t represents the volatility process.

Recently, analysis of volatility has been enhanced by using both intraday data, and high low range information.[20] The long range CRSP returns data provides daily high/low range information as well as returns which allows for a more detailed picture of the dynamics of conditional variances in the returns series. Following Parkinson (1980), volatility on day t will be estimated with

$$\hat{\sigma}_t^2 = \frac{1}{4\log(2)}(\log(h_t) - \log(l_t))^2. \tag{12}$$

This will then be used to construct returns by applying it to a parametric distribution for ϵ_t.

Table 7 takes the daily estimates of volatility as given, draws an independent set of these using an iid bootstrap, and then applies these to a Gaussian distribution for ϵ_t. The experiment is to see if this very simple volatility process can generate tail estimates which are close to those coming from the returns series. If they do, this gives strong evidence that changing conditional variances alone are driving much of the tails, and also that the range estimators have given an accurate picture of volatility. The first three columns of the table correspond to this bootstrap experiment. The column labeled bootstrap mean gives the mean tail estimate from the bootstrap data, and the third column presents the fraction of bootstraps with tail estimates less than the actual return series. The values give a general picture that the simulated series are actually fatter tailed than the original returns. The last 3 columns repeat the same test, but perform a monte-carlo experiment, assuming a parametric Gaussian distribution for the logged volatility, $\log(\sigma_t^2)$. Volatility is now drawn from this distribution with mean and variance estimated separately for each return series. The table repeats the results from the first 3 columns showing a generally fatter tailed distribution.

It is not immediately clear why the range based estimators do not provide a good generator for returns series. There may be many reasons. First, the Gaussian noise assumed in the stochastic volatility process may not be correct. Many studies have replaced this with other fat tailed distributions such as a Student-t distributions. Another problem is that the high/low range estimator is subject to microstructure bias. The high might often be the the current ask, while the low is the current bid. This would bias the range to be more extreme that the true volatility process warrants. This might lead to fatter tails when it is applied to the simulated noise process. Finally, recent studies have begun trying to separate discrete jumps from more continuous aspects of returns. The methodology used here ignores this possibility both in its impact on the

[20]See for example Alizadeh, Brandt & Diebold (2002) and Martens & van Dijk (forthcoming 2008). Early derivations of range based estimators can be found in Garman & Klass (1980) and Parkinson (1980).

volatility estimate, and on the returns generation process. Most techniques to adjust for this use intraday data which is not available over these longer horizons.[21]

4 Conclusions

This paper paper has both a methodological and an empirical message. In terms of methodology, it was shown that the modified Hill estimator developed in Huisman et al. (2001) looks very promising in terms of providing a reliable estimate for tail exponents. Not only does it appear to be more precise in terms of mean squared error, but it also seems to be more forgiving to getting the exact tail scaling region correct. This is probably a very important issue for financial, and macro economic time series where samples are often very short. With an easy to use and reliable tail exponent estimator in hand, researchers should think about adding this to the usual set of descriptive statistics that are often reported for financial time series. It could be argued that the simple tail exponent estimate is at least as reliable as a kurtosis estimate given that the evidence suggests that fourth moments might not exist for most daily return series.

The empirical conclusions about tail properties are also important. The evidence defends a common belief that most return series show tail exponents near 3. This appears generally stable across the set of individual firms, and over the sample. It is very interesting that the vast technological changes that have occurred over the past century have not greatly impacted these general stylized facts. One should still view extremely precise estimates of these values with some caution, since even the estimators used in this paper still show fairly large standard errors.

This paper also introduced a very simple measure of tail behavior, the Gaussian crossing point. This also gives a simple picture of fat tailed behavior that could potentially be useful in terms of risk management and the assessment of VaR measures and extreme market moves. Most estimates using this measure were consistent with those from the tail exponents, confirming the general fatness of the return tails relative to a Gaussian. The crossing estimates of near 0.02 in terms probability distribution levels remind us that the "fatness" in the tails only starts as one moves to the extremes.

Both the tail measures were used to explore asymmetry with mixed results. From the standpoint of the tail exponent, there was little evidence for asymmetry in the return distributions. However, the crossing point evidence pointed to a larger amount of return activity in the extreme right tail, pointing to an unusual number of positive jumps. It disagrees with a general feeling that most asymmetry is in the negative

[21] Examples are Martens & van Dijk (forthcoming 2008) and Andersen, Bollerslev, Frederiksen & Nielsen (2007) and Corsi, Pirino & Reno (2008).

directions, as in large crashes. Curiously, the index generated significant results, which varied across the two tail measures. For the more traditional tail index it presented strong evidence for a fatter *right* tail. All of these results need some further exploration, but they emphasize that the common assumption about fatter left tails it not strongly supported in the data, and may depend critically on the type of tail measured in use.

Finally, some simple explorations were performed to understand how changing volatility might be involved. Simple volatility processes built from range based estimators generated returns series with even fatter tails than the stock returns. It is not clear exactly what is at work here, but a possible explanation is that jumps are not being correctly taken into account. Doing this correctly might require modifying identified jump days as in Jiang & Yao (2007), or moving to high frequency data as in Andersen et al. (2007).

This research is part of a continuing project to try to document how reliable our stylized facts are in financial markets. It is clear that not only are "fat tails" consistent across time and firms, but they also give us information about the general properties of extreme return distributions which are useful both to practitioners interested in risk management, as well as theorists interested in uncovering their economic causes.

References

Alfarano, S. & Lux, T. (Forthcoming), Extreme value theory as a theoretical background for power law behavior, *in* 'Power Laws in the Social Sciences: Discovering Compexity and Non-equilibrium Dynamics in the Social Universe'.

Alizadeh, S., Brandt, M. W. & Diebold, F. X. (2002), 'Range-based estimation of stochastic volatility models', *Journal of Finance* **57**, 1047–1091.

Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2005), Volatility forecasting, Technical report, NBER.

Andersen, T. G., Bollerslev, T., Frederiksen, P. H. & Nielsen, M. O. (2007), Continuous-time models, realized volatilties, and testable distributional implications for daily stock returns, Technical report, University of Aarhus.

Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **21**, 307–328.

Brock, W. A. (1996), 'Scaling in economics: A reader's guide', *Industrial and Corporate Change* **8**, 409–446.

Brock, W. A., Hsieh, D. & LeBaron, B. (1991), *Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence*, MIT Press, Cambridge, MA.

Brooks, C., Clare, A. D., Molle, J. W. D. & Persand, G. (2005), 'A comparison of extreme value theory approaches for determining value at risk', *Journal of Empirical Finance* **12**, 339–352.

Caeiro, F. & Gomes, M. I. (2008), 'Minimum-variance reduced-bias tail index and high quantile estimation', *REVSTAT: Statistical Journal* **6**, 1–20.

Chistoffersen, P. F. (2003), *Elements of Financial Risk Management*, Academic Press.

Clauset, A., Shalizi, C. R. & Newman, M. E. J. (2009), Power-law distributions in empirical data, Technical report, University of Michigan, Ann Arbor, Michigan.

Corsi, F., Pirino, D. & Reno, R. (2008), Volatility forecasting: The jumps do matter, Technical report, Universita di Siena.

Dacorogna, M., Muller, U. A., Pictet, O. V. & de Vries, C. G. (1995), The distribution of extremal foreign exchange rate returns in estremely large data sets, Technical report, Olsen and Associates.

Danielson, J., de Haan, L., Peng, L. & de Vries, C. G. (2001), 'Using a bootstrap method to choose the sample fraction in tail index estimation', *Journal of Multivariate Analysis* **76**, 226–248.

de Haan, L. & Ferreira, A. (2006), *Extreme Value Theory: An Introduction*, Springer Verlag.

Draisma, G., de Haan, L. & Peng, L. (1999), 'A bootstrap-based method to achieve optimality in estimating the extreme-value index', *Extremes* **2**, 367–404.

Embrechts, P., Kluppelberg, C. & Mikosch, T. (1997), *Modeling Extremal Events for Insurance and Finance*, Springer Verlag.

Engle, R. F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.

Fama, E. F. (1965), 'The behavior of stock market prices', *Journal of Business* **38**, 34–105.

Farmer, J. D. & Geanakoplos, J. (2008), Power laws in economics and elsewhere, Technical report, Santa Fe Institute, Santa Fe, NM.

Garman, M. B. & Klass, M. J. (1980), 'On the estimation of security price volatilities from historical data', *Journal of Business* **53**, 67–78.

Gomes, M. I. & Oliveira, O. (2001), 'The bootstrap methodology in statistics of extremes - Choice of the optimal sample fraction', *Extremes* **4**, 331–358.

Gomes, M. I. & Pestana, D. (2007), 'A sturdy reduced-bias extreme quantile (var) estimator', *Journal of the American Statistical Association* **102**, 280–292.

Hall, P. (1990), 'Using the bootstrap to estimate mean squared error and select smoothing parameter in nonparametric problems', *Journal of Multivariate Analysis* **32**, 177–203.

Hill, B. (1975), 'A simple general approach to inference about the tail of a distribution', *The Annals of Mathematical Statitics* **3**, 1163–1174.

Huisman, R., Koedijk, K. G., Kool, C. J. M. & Palm, F. (2001), 'Tail-index estimates in small samples', *Journal of Business and Economic Statistics* **19**, 208–216.

Jiang, G. J. & Yao, T. (2007), Stock price jumps and cross-sectional return predictability, Technical report, Eller College of Management, University of Arizona.

LeBaron, B. & Samanta, R. (2004), Extreme value theory and fat tails in equity markets, Technical report, International Business School, Brandeis University, Waltham, MA.

Lisi, F. (2007), 'Testing asymmetry in financial time series', *Quantitative Finance* **7**, 687–696.

Longin, F. (1996), 'The asymptotic distribution of extreme stock returns', *Journal of Business* **63**, 383–408.

Longin, F. (2005), 'The choice of the distribution of asset returns:How extreme value theory can help?', *Journal of Banking and Finance* **29**, 1017–1035.

Loretan, M. & Phillips, P. C. B. (1994), 'Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with some applications to several financial series', *Journal of Empirical Finance* **1**(2), 211–248.

Lux, T. (1996), 'The limiting extremal behavior of speculative returns: An analysis of intra-daily data from the frankfurt stock exchange', *Applied Financial Economics* **11**, 299–315.

Mandelbrot, B. B. (1963), 'The variation of certain speculative prices', *Journal of Business* **36**, 394–419.

Martens, M. & van Dijk, D. (forthcoming 2008), 'Measuring volatility with the realized range', *Journal of Econometrics* .

Parkinson, M. (1980), 'The extreme value method for estimating the variance of the rate of return', *Journal of Business* **53**, 61–65.

Perez-Alonso, A. (2007), 'A bootstrap approach to test the conditional symmetry in time series models', *Computational Statistics and Data Analysis* **51**, 3484–3504.

Table 1: *Summary Statistics*

	Mean	Std	Skewness	Kurtosis
ATT	0.028	1.314	-0.105	26.692
Coke	0.048	1.412	-0.300	17.931
DuPont	0.040	1.572	-0.202	14.947
Edison	0.034	1.468	-1.096	38.294
Exxon	0.044	1.394	-0.127	17.031
GE	0.043	1.716	0.044	13.071
GM	0.040	1.770	0.193	13.792
Goodrich	0.027	2.449	0.478	21.443
IBM	0.052	1.542	-0.256	15.807
Kansas Southern	0.047	3.008	0.239	23.325
Navistar	0.012	2.427	0.232	21.604
Sears	0.037	1.872	-0.258	30.755
US Steel	0.024	1.867	-0.210	27.686
Woolworth	0.024	1.841	0.347	12.914
VW Index	0.041	1.029	-0.132	20.920

Summary statistics: All values correspond to daily logged returns from January 1st, 1926 through December 31st, 2004. The sample size is 21,016. Mean and Std. are reported as percentages. Kurtosis is the raw kurtosis value which would be equal to 3 for a Gaussian distribution.

Table 2: *Tail Exponents: Left Tail*

	k=0.15	k=0.25	k=0.20	Std(k=0.20)
ATT	3.217	3.576	3.380	0.145
Coke	3.412	3.709	3.528	0.119
DuPont	3.388	3.530	3.367	0.126
Ed	3.048	3.189	2.915	0.115
Exxon	3.608	3.914	3.636	0.151
GE	3.281	3.495	3.311	0.136
GM	3.313	3.685	3.448	0.131
Goodrich	2.793	2.921	2.821	0.110
IBM	3.244	3.467	3.257	0.129
Kansas Southern	2.831	2.792	2.828	0.102
Navistar	3.396	3.375	3.297	0.159
Sears	3.193	3.370	3.203	0.148
US Steel	3.653	3.803	3.691	0.154
Woolworth	3.213	3.472	3.275	0.107
VW Index	3.034	3.398	3.151	0.102

Estimated tail exponents from the left tail of log daily returns, for individual equity firms, full sample. Estimator is the OLS bias adjusted Hill estimator. k is the fraction of the left tail of the returns distribution used. Bootstrap standard errors are from n draws of the an returns with replacement. Standard errors are estimated for $k = 0.20$.

Table 3: *Tail exponents: Right Tail*

Name	k=0.15	k=0.25	k=0.20	Std(k=0.20)
ATT	3.320	3.732	3.581	0.160
Coke	3.760	4.126	3.920	0.162
DuPont	3.554	3.911	3.652	0.145
Ed	3.262	3.233	2.999	0.112
Exxon	3.534	3.846	3.596	0.160
GE	3.330	3.604	3.448	0.158
GM	3.518	3.577	3.507	0.139
Goodrich	2.770	3.090	2.905	0.089
IBM	3.363	3.675	3.495	0.144
Kansas Southern	3.135	3.203	3.151	0.138
Navistar	3.025	3.253	3.076	0.131
Sears	3.295	3.436	3.237	0.151
US Steel	3.659	4.045	3.795	0.189
Woolworth	3.526	3.473	3.562	0.122
VW Index	2.800	2.908	2.824	0.107

Estimated tail exponents from the right tail of log daily returns, for individual equity firms, full sample. Estimator is the OLS bias adjusted Hill estimator. k is the fraction of the left tail of the returns distribution used. Bootstrap standard errors are from n draws of the an returns with replacement. Standard errors are estimated for $k = 0.20$.

Table 4: *Tail Exponents: Difference*

	Left	Right	Difference	Prob(Difference$> D^*$)
ATT	3.380	3.581	-0.201	0.196
Coke	3.528	3.920	-0.392	0.032
DuPont	3.367	3.652	-0.286	0.084
Ed	2.915	2.999	-0.083	0.292
Exxon	3.636	3.596	0.039	0.596
GE	3.311	3.448	-0.137	0.252
GM	3.448	3.507	-0.059	0.320
Goodrich	2.821	2.905	-0.084	0.292
IBM	3.257	3.495	-0.238	0.104
Kansas Southern	2.828	3.151	-0.323	0.036
Navistar	3.297	3.076	0.221	0.856
Sears	3.203	3.237	-0.034	0.448
US Steel	3.691	3.795	-0.105	0.292
Woolworth	3.275	3.562	-0.286	0.068
VW Index	3.151	2.824	0.327	0.982

Tail exponent symmetry tests. Difference is the left - right tail difference. The final column displays the fraction of bootstrap differences less than the actual data from a sign scrambling bootstrap assuming a null hypothesis of symmetry in the return distribution.

Table 5: *Right and Left Gaussian Crossing Points and Differences*

	$p_{c,l}$	$\sigma(p_{c,l})$	$p_{c,r}$	$\sigma(p_{c,r})$	$p_{c,l} - p_{c,r}$	Prob(Difference> D^*)
ATT	0.026	0.003	0.032	0.003	-0.006	0.055
Coke	0.023	0.003	0.038	0.004	-0.015	0.002
DuPont	0.023	0.003	0.033	0.003	-0.010	0.010
Edison	0.027	0.003	0.029	0.003	-0.002	0.287
Exxon	0.025	0.003	0.028	0.003	-0.003	0.216
GE	0.025	0.002	0.027	0.003	-0.002	0.338
GM	0.022	0.003	0.034	0.003	-0.012	0.000
Goodrich	0.022	0.002	0.025	0.002	-0.003	0.142
IBM	0.024	0.003	0.032	0.004	-0.008	0.031
Kansas Southern	0.027	0.002	0.031	0.003	-0.004	0.123
Navistar	0.025	0.003	0.029	0.003	-0.004	0.119
Sears	0.022	0.003	0.028	0.003	-0.006	0.031
US Steel	0.019	0.004	0.028	0.004	-0.009	0.045
Woolworth	0.024	0.003	0.038	0.003	-0.014	0.000
VW Index	0.030	0.002	0.018	0.002	0.012	1.000

Gaussian crossing point estimates: Values represent probability distribution levels at which the empirical cumulative distribution crosses the corresponding Gaussian. Difference is the left - right tail difference. The final column displays the fraction of bootstrap differences less than the actual data from a sign scrambling bootstrap assuming a null hypothesis of symmetry in the return distribution.

Table 6: *First and Last 5000 Daily Returns*

	Early	Late	Difference	Prob(Difference> D^*
ATT	4.059	3.449	0.610	0.910
Coke	3.178	3.667	-0.490	0.126
DuPont	3.125	3.995	-0.870	0.016
Edison	3.752	3.564	0.187	0.722
Exxon	3.579	3.814	-0.234	0.270
GE	3.537	3.499	0.038	0.532
GM	3.460	3.973	-0.513	0.092
Goodrich	3.118	3.273	-0.155	0.302
IBM	2.908	3.197	-0.289	0.232
Kansas Southern	3.613	3.263	0.350	0.887
Navistar	2.925	4.072	-1.148	0.005
Sears	3.168	3.483	-0.315	0.217
US Steel	3.661	5.143	-1.481	0.001
Woolworth	3.206	3.268	-0.062	0.414
VW Index	3.338	4.364	-1.026	0.001

Tail exponent stability tests: Difference is the difference between the first and last 5000 daily returns in the sample. The final column displays the fraction of bootstrap differences less than the actual data from an IID bootstrap.

Table 7: *Volatility Comparisons: Left tail exponent*

	Return Data	Bootstrap mean	Prob(Tail$>D^*$)	Monte-carlo mean	Prob(Tail$>D^*$)
ATT	3.380	2.990	1.000	2.279	1.000
Coke	3.528	3.359	0.860	2.849	1.000
DuPont	3.366	3.177	0.980	2.852	1.000
Edison	2.916	2.898	0.620	3.129	0.060
Exxon	3.636	3.218	1.000	3.434	0.980
GE	3.311	3.071	0.940	3.525	0.060
GM	3.448	3.127	0.960	3.285	0.860
Goodrich	2.822	2.726	0.860	3.070	0.000
IBM	3.257	3.311	0.380	2.867	1.000
Kansas Southern	2.828	2.181	1.000	2.471	1.000
Navistar	3.297	2.726	1.000	2.796	1.000
Sears	3.203	2.977	0.980	2.954	0.980
US Steel	3.690	3.284	1.000	3.551	0.820
Woolworth	3.275	2.943	1.000	2.926	0.980

Comparisons with simple stochastic volatility: Bootstrap tests correspond to resampled volatility estimates from daily high/low range information. Monte-carlo results correspond to volatilities drawn from a log normal distribution fit to each individual return volatility series. Bootstrap means, and the fraction of bootstraps with smaller tail exponents than the actual are given in the columns labeled, Prob(Tail$>D^*$).

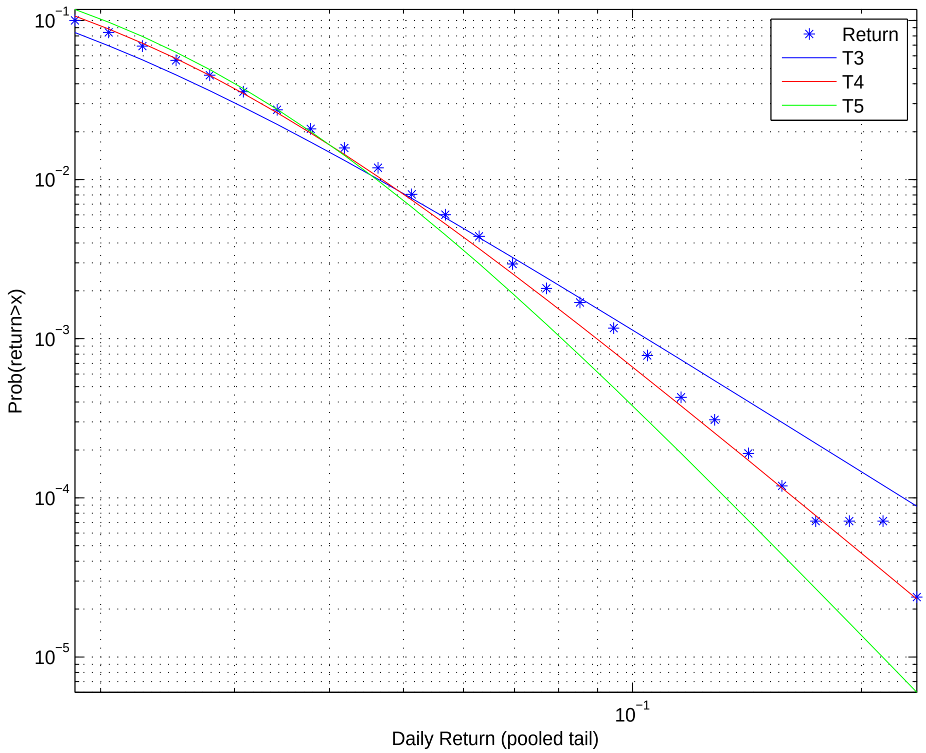

Figure 1: *GM Returns: Pooled tail and student-t*

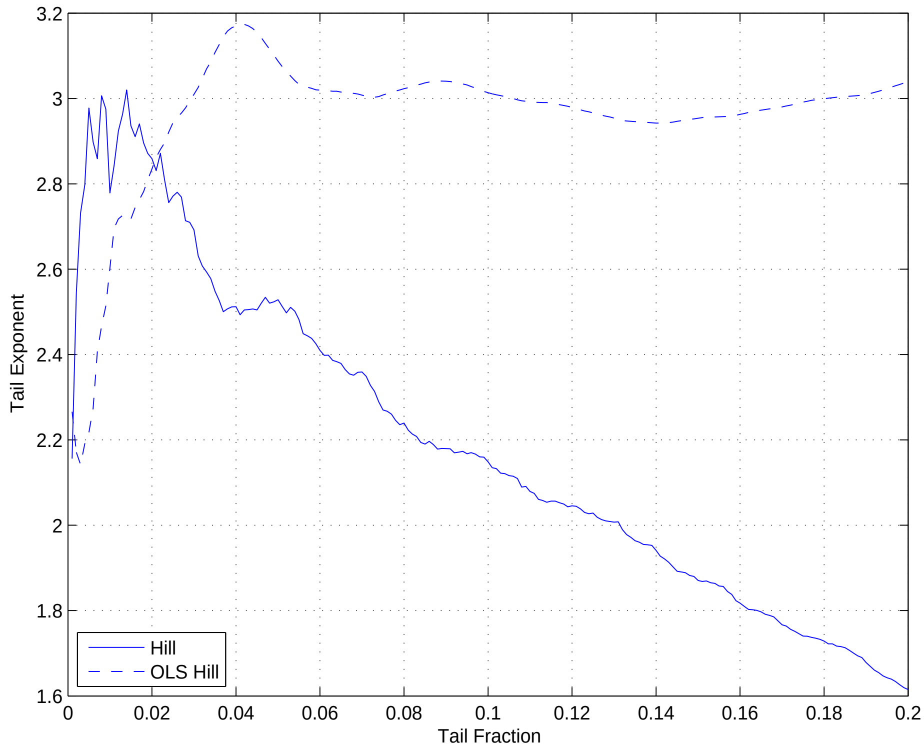

Figure 2: *Hill and Hill OLS: Student-t(3)*

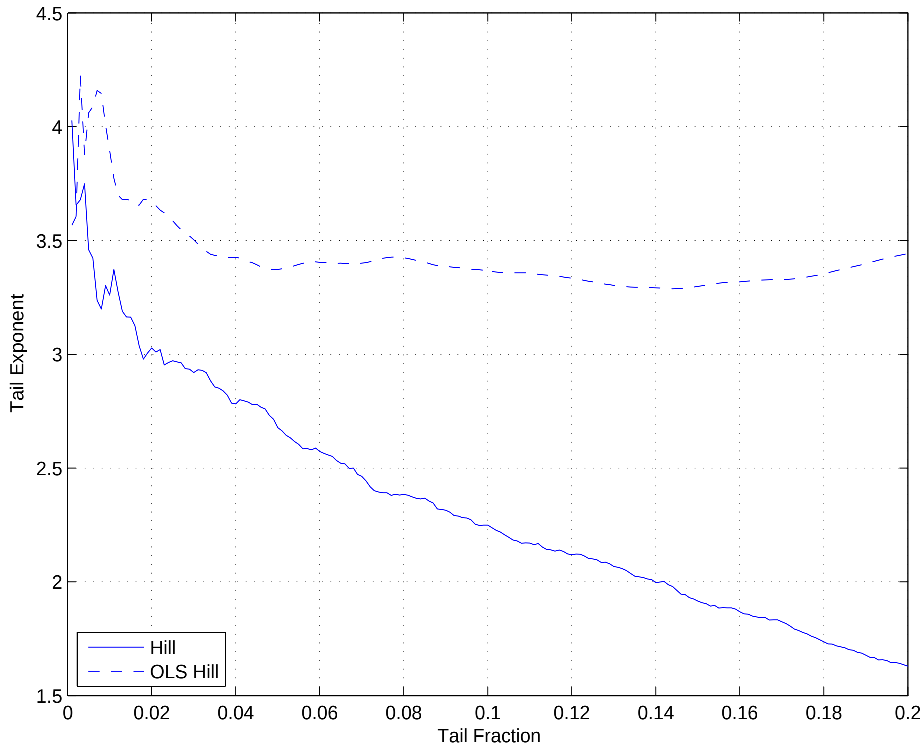

Figure 3: *General Motors Daily Return (Left Tail): Hill and OLS-Hill Tail Index Estimates*

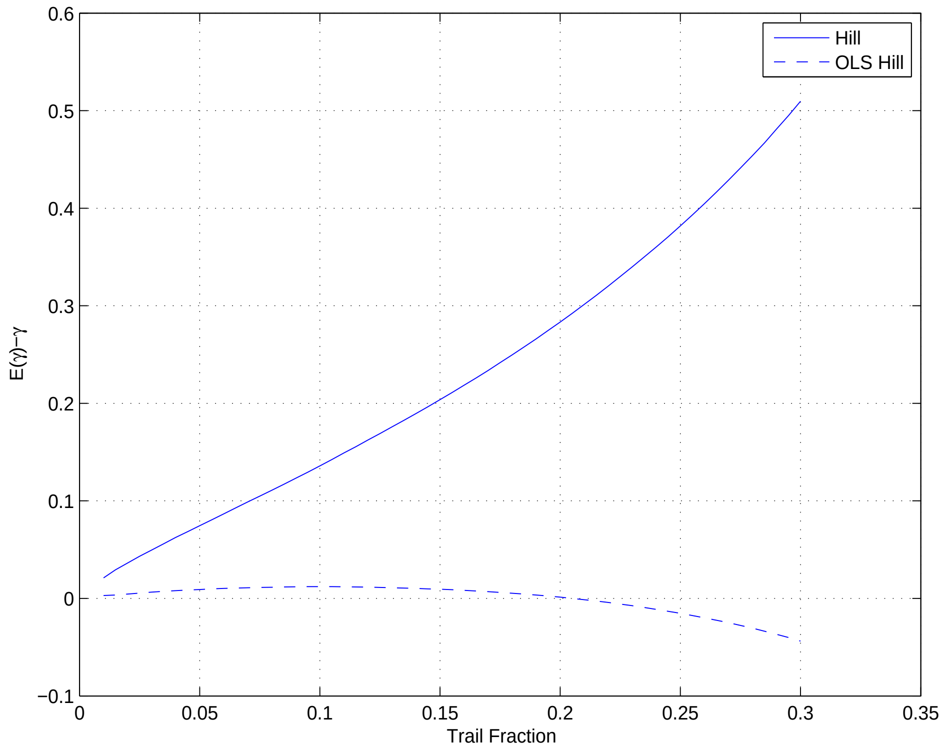

Figure 4: *Hill and OLS Hill Bias Estimates: Student-t(3)*

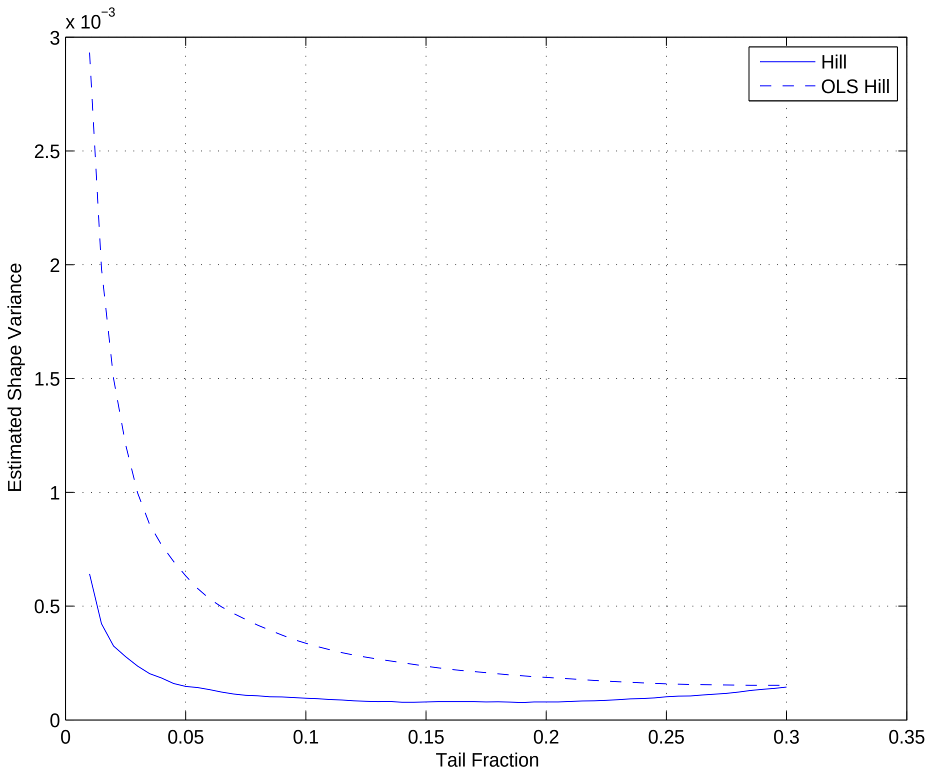

Figure 5: *Hill and OLS Hill Variance Estimates: Student-t(3)*

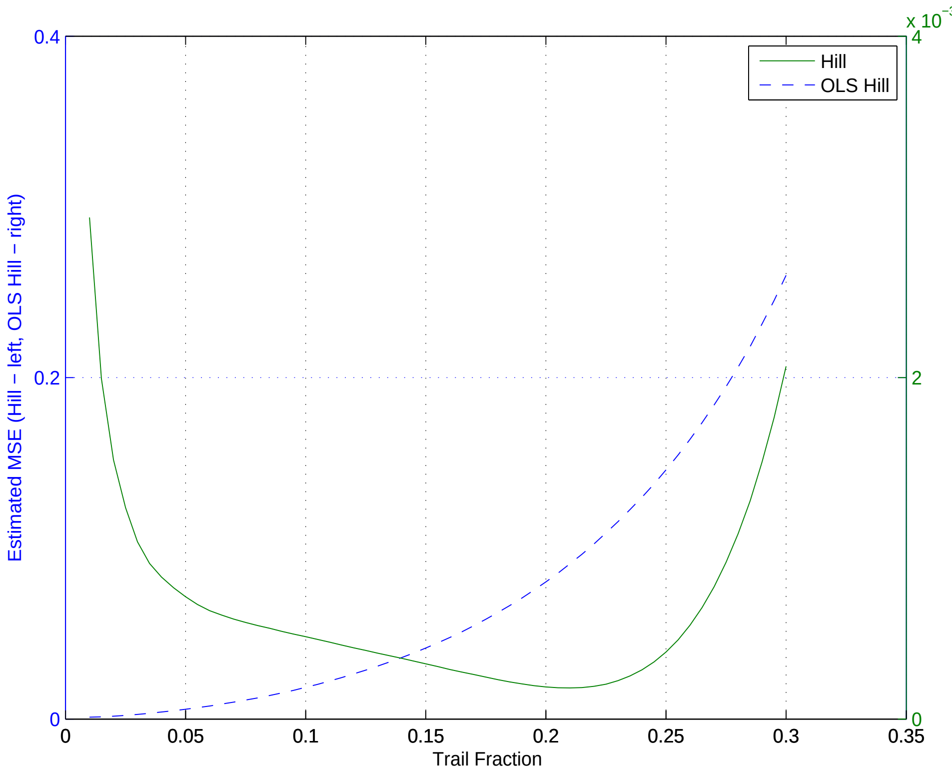

Figure 6: *Hill and OLS Hill MSE: Student-t(3)*
Notes: MSE for Hill and OLS Hill estimators. Min(Hill) $= 1.1e - 3$ and Min(OLS Hill) $= 1.8e - 4$..

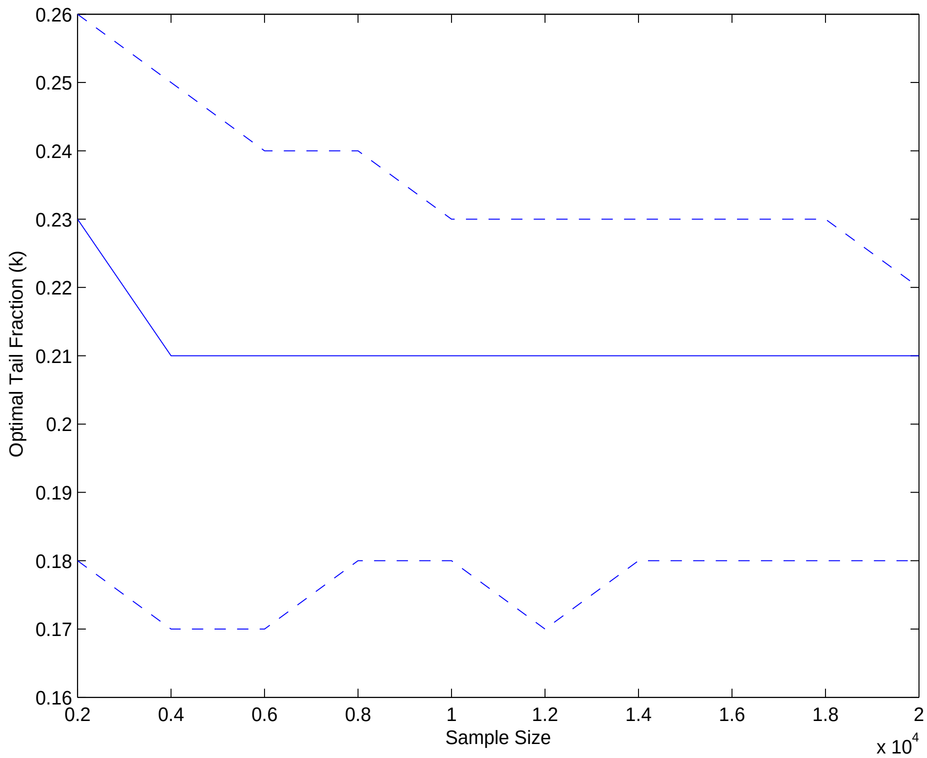

Figure 7: *Monte-carlo (Student-t(3)) MSE Optimal Tail Fraction and* 25 *Percent Bands*

Notes: Optimal tail fractions for varying sample sizes.
k values within the bands yield estimators with a MSE
within 25 percent of the optimum.