Internal vs. External Risk Measures: How Captial Requires Differ in Practice

Martin Eling und Luisa Tibiletti

Preprint Series: 2009-17

Fakultät für Mathematik und Wirtschaftswissenschaften
UNIVERSITÄT ULM

Internal vs. External Risk Measures: How Capital Requirements Differ in Practice

Martin Eling and Luisa Tibiletti*

Abstract: We compare capital requirements derived by tail conditional expectation (TCE) with those derived by tail conditional median (TCM) and find that there is no clear-cut relationship between these two measures in empirical data. Our results highlight the relevance of TCM as a robust alternative to TCE, especially for regulatory control.

JEL Classification: G10, G11, G23, G29

Keywords: Risk measures, Tail Conditional Expectation, Tail Conditional Median, Value-at-Risk, Robust Statistics

* Martin Eling (martin.eling@uni-ulm.de), Institute of Insurance Science, Ulm University, Germany. Luisa Tibiletti (luisa.tibiletti@unito.it), Department of Statistics and Mathematics, University of Torino, Italy.

1. Introduction

Value-at-risk (VaR) or tail conditional expectation (TCE)? Since the publication of Artzner et al. [4], a great deal of literature has appeared that discusses the drawbacks of VaR as compared to the benefits of TCE or conditional value-at-risk (CVaR). Specifically, Artzner et al. [4] define a set of axioms and call risk measures coherent if they satisfy these axioms. One of the axioms is subadditivity, which basically means that "a merger does not create extra risk" (Artzner et al. [4]). The authors show that VaR does not always satisfy the subadditivity axiom and thus conclude that VaR is not suitable for risk management. An alternative that satisfies the subadditivity axiom and is flexible enough to handle optimization procedures is CVaR. A very recent stream of literature focuses on the pros and cons of VaR and TCE, in an attempt to discover which is best.

Only recently the question of how to choose the best objective-oriented risk measure has arisen in literature (see Heyde et al. [10]). The first question that needs to be answered is: For what "audience" is the risk measure intended—equity shareholders, regulatory/legal agencies, or internal management? There is no reason to believe that one, single risk measure will satisfy the needs of these different parties. In particular, some risk measures may be suitable for internal risk management but not for external regulatory agencies, and vice versa. For internal risk control, for example, coherent and convex risk measures are preferable because of their subadditive properties. However, for external risk measures, a different set of properties might be more appropriate, including consistency in implementation, which means robustness. In such a case, the so-called natural risk statistics, including the tail conditional median (TCM) as a special case, display more robustness than tail conditional expectation (see Cont et al. [7]).

The choice of risk measure is of great practical importance: It determines the minimum reserve for the computation of margin requirements in financial trading, insurance risk premiums, regulatory deposit requirements, among many others. The level of capital requirement is crucial to all profitability ratios that relate company profit to the underlying capital basis. A different risk measure may lead to lower (higher) capital requirement and thus to a higher (lower) profitability of the firm. Thus, a key question for practitioners is whether, given that the confidence level α is fixed, the choice of either TCE or TCM will lead to different capital requirements? And, if so, which measure will result in the most restrictive capital requirements?

The aim of this paper is to answer these questions. We thus empirically analyze a number of different risk positions involving different asset classes and different time horizons. We also consider simulated data on distributions commonly used in finance as well as popular equity return models (ARCH/GARCH, jump diffusion process). The degree of fatness in the tail of a distribution, as well as the availability of data (and the resulting estimation risk), are two major factors when comparing TCE and TCM. There is a clear-cut relationship between TCE and TCM in theory (requirements from TCE are larger than those from TCM); however, there is no such relationship in the empirical data. If there are only few data available for estimating, TCM is the more robust choice.

These findings support the Heyde et al. [10] proposal to implement objective-oriented risk measures. In fact, choosing the most appropriate risk measure based on what it is expected to represent is not only correct as a normative approach, but should also be done in practice, since internal and external measures lead to different results.

The remainder of the paper is organized as follows. Section 2 discusses the properties desirable for internal and external measures. Section 3 presents numerical results on the TCE and TCM mismatch. Section 4 concludes.

2. Objectives of Risk Measures: Internal and External Risk Measures

Firms and regulators have different perspectives on risk management. Jarrow and Purnanandam [11] term these different perspectives as “capital budgeting” and “capital determination,” respectively. The firm’s goal is to achieve a portfolio allocation that maximizes its risk/return tradeoff, subject to any regulatory capital requirements. In contrast, taking the firm portfolio allocation as fixed, the regulator’s goal is to determine how much capital the firm must have in order to limit the consequences of losses within a given time. Heyde et al. [10] stress this point from an axiomatic viewpoint, focusing on the different axioms that are desirable to each involved party. In their scheme, risk measures tailored to firm’s goals are called “internal” measures, and those that are more in line with the regulator’s goal are called “external” measures. For a deeper discussion, see Heyde et al. [10].

2.1 Internal Risk Measures: Tail Conditional Expectation

In the early 1990s, due to a number of spectacular company bankruptcies attributed to the inappropriate use of derivatives and a lack of sufficient internal controls, a new method of measuring risk called value-at-risk (VaR) was developed. Originally VaR was intended to measure the risks in derivatives markets, but it has become widely used to measure all kinds of financial risks, primarily market and credit risks, and even as a tool for active internal risk management.

Artzner et al. [4] point out VaR´s lack subadditivity. The authors introduce a new family of subadditive measures called “coherent risk measures.” One of these is the tail conditional expectation (TCE). More precisely, the TCE at level $\alpha \in (0\%,\ 100\%)$ is defined by:

$$TCE_\alpha = \text{mean of the } \alpha\text{-tail distribution of the loss } X.$$

If the distribution of *X* is continuous, then

$$TCE_\alpha(X) = E\left[X \,\middle|\, X \geq VaR_\alpha(X)\right].$$

For discrete distributions, if one appropriately defines quantiles for discrete random variables (see Acerbi and Tasche [1]), then TCE_α is the sample tail conditional expectation.

It is well known that expectation is a data-sensitive central measure and, consequently so is TCE. It follows that TCE is strongly influenced by model assumptions regarding the heaviness of tail distributions. Its appropriateness for use as an external risk measure by regulators is questionable. In fact, each firm has developed its own financial model and tail heaviness has become a subjective issue (see Heyde et al. [10], pp. 12 ff.). However, even if TCE is not an appropriate regulatory or external risk measure, it has the potential to be an excellent internal risk measure. Since TCE is also a convex risk measure, it is widely used in stochastic convex optimization for determining optimal asset allocation.

2.2 External Risk Measure: Tail Conditional Median

Heyde et al. [10] discuss desirable axioms for an external risk measure. Subadditivity is relaxed only to comonotonic data. Comonotonic subadditivity is consistent with the prospect theory of risk in psychology. The new data-based class of risk measures introduced is called “natural risk statistics” and a representation theorem is given (see Ahmed et al. [3] for alternative proofs). A special case of natural risk measures is provided by tail conditional median (TCM). More precisely, the TCM at level α is defined by:

$$TCM_\alpha = \text{median of the } \alpha\text{-tail distribution of the loss } X,$$

or

$$TCM_\alpha(X) = median\left[X \,\middle|\, X \geq VaR_\alpha(X)\right].$$

This representation holds for both discrete and continuous distributions. If *X* is continuous, TCM can be rewritten in terms of VaR, since

$$TCM_{\alpha}\left(X\right)=VaR_{\frac{1+\alpha}{2}}\left(X\right).$$

As mentioned by Heyde et al. [10], and contrary to conventional wisdom, the above equation makes it evident that VaR, at a higher level, can also incorporate tail information. For example, if one wants a central measure of the loss beyond the 95% level, VaR at 97.5% is a possibility because it gives the tail conditional median at 95% level. Note that this holds only for continuous distributions and not in the discrete case (which is the case in real world risk management using empirical data). However, also in the discrete case, VaR at 97.5% can be used as an approximation of TCM at 95%. As a median, TCM is a robust central measure and thus particularly suitable as an external risk measure. The empirical results on the robustness of conditional median presented by Ogryczak and Zawadzki [13] are further confirmation of the measure's theoretical attractiveness.

3. Numerical Analysis

3.1. Data

In this section we analyze the properties of tail conditional median as compared to those of tail conditional expectation. In the first step, we consider empirical return data on stocks and hedge funds. In the second step, we use simulated data and analyze tail distributions commonly used in finance as well as popular equity return models such as ARCH/GARCH and jump diffusion processes. To calibrate these models, we again use the empirical return data and maximum likelihood estimation.

From a practical point of view, the most important thing to discover is whether these different risk measures lead to different results. An answer of "yes" will have important implications for risk management. If the answer comes up "no," the discussion becomes academic. However, as we will show in the following, the TCE vs. TCM debate turns out to be important for both practitioners and academics. We consider three data sources that incorporate both different return distribution characteristics and various time horizons.

(1) We consider returns of the 500 stocks contained in the S&P 500 and the S&P 500 index itself. We have return data available from the Datastream database for the period January 1990 to December 2004. We consider these data both on a daily and on a monthly basis.

(2) We analyze 1,374 mutual funds with monthly returns from January 1996 to December 2005. We choose a group of equity-oriented mutual funds previously analyzed in another context (Eling [9]) so as to improve comparability with the stock data presented under (1). The data originate from the Datastream database. The consideration of mutual funds is different from the

analysis of stocks in two aspects. First, mutual funds are portfolios of stocks and thus more diversified than the individual stocks. Second, using mutual fund data provides a natural way to incorporate transaction costs into our analysis. The mutual fund data are net of all fees, while the stock data do not incorporate any transaction costs.

(3) We also look at hedge fund data provided by the Center for International Securities and Derivatives Markets (CISDM). Our database contains 205 hedge funds reporting monthly returns, again net of all fees, for the period of January 1996 to December 2005. The consideration of hedge fund data provides further insights because the returns of hedge funds are very skewed and have much heavier tails than traditional investments. Hedge funds might therefore be the asset class with the most extreme tails in financial markets.

Table 1 contains descriptive statistics on the return distributions of the stocks, mutual funds, and hedge funds. The table sets out mean, median, standard deviation, minimum, and maximum of the first four moments of the return distribution (mean value, standard deviation, skewness, and excess kurtosis). For example, the standard deviation in Row 5 means that across the 500 stocks, the standard deviation of the returns has a mean of 10.60% (second column in Row 5) with a standard deviation of 4.57% (fourth column in Row 5). Note that all data in Table 1 are monthly numbers.

Insert Table 1 here

The individual stocks generate the highest returns on average, but these are accompanied by a very high risk in terms of standard deviation. The returns of most of the individual stocks are positively skewed and have positive excess kurtosis. On the basis of the Jarque-Bera test, the assumption of normally distributed stock returns must be rejected for 56.83% at the 1% significance level. The diversified mutual funds provide a lower return with a lower risk. The lower return might be driven by the transaction costs that are considered within the mutual fund returns, but not for the individual stocks. The lower standard deviation illustrates the risk reduction through portfolio diversification. As indicated, the returns of hedge funds are more skewed and have fatter tails than do those of the mutual funds. It is thus not surprising to see that the Jarque-Bera rejection rate is far higher for hedge funds than it is for mutual funds.

3.2 Empirical Results

(1) S&P 500 Index and 500 Stocks

We split the analysis of the S&P 500 data into two steps. In the first step, we consider only the S&P 500 index, which aggregates the performance of all 500 stocks; this analysis is comparable to the analysis presented in Heyde et al. [10], p. 25. In the second step, the 500 stocks contained in the index are analyzed on a disaggregated level; here, we present average values and other statistics about the 500 individual stocks.

Our analysis of the S&P 500 index covers 3,784 daily index returns in the period January 1990 to December 2004. Table 2 shows the tail conditional expectation (TCE) and tail conditional mean (TCM) for different confidence levels. To highlight the difference between TCE and TCM, we present the difference between TCE and TCM in column 4 and the relative difference between TCE and TCM in column 5 (defined as (TCE-TCM)/TCE).

Insert Table 2 here

We confirm Heyde et al.'s [10], p. 25, results concerning the S&P 500 daily losses. In almost all cases, TCE is larger than TCM, denoting the presence of fat tails. The differences can be quite significant; for example, at a confidence level of 90%, TCE is 11.23% higher than TCM. Note, however, that at a confidence level of 99.9%, TCM is larger than TCE. This switching behavior illustrates that the magnitude of fatness is less for extreme catastrophic events. The findings in Table 2 are the only ones in this paper that are based on a single index; all the tables that follow present average values for a large number of stocks or funds.

To achieve a broader view of the differences between TCE and TCM, we now consider the 500 stocks that are listed in the S&P 500 index (3,784 daily returns in the period from January 1990 to December 2004). Table 3 contains for different confidence levels α the corresponding values of TCE and TCM, their difference, and their relative difference. The numbers presented in Table 3 are arithmetic averages across 500 TCE and TCM values, i.e., we first calculate the TCE and TCM for each of the 500 stocks and then determine the average across the 500 values of TCE and TCM.

Insert Table 3 here

Comparing results of Table 3 with those of Table 2, we see that TCE is always larger than TCM. The relative difference is greater than that revealed for the single index; it is about 12% for confidence levels above 95%. It thus seems that the TCE always leads to higher capital requirements than the TCM.

Figure 1 shows the frequency distribution of (TCE-TCM)/TCE for the 500 stocks in the S&P. The horizontal axis shows the relative difference between TCE and TCM and the vertical axis shows the number of funds. Results for three confidence levels (90%, 95%, 99%) are presented. Beneath the figure are some descriptive statistics on the distribution (min, 25% quantile, median, 75% quantile, max, mean, range (=max-min), and percentage of negative values). The mean value corresponds to the (TCE-TCM)/TCE value presented in Table 3.

Insert Figure 1 here

The frequency distribution illustrates that the difference between TCE and TCM is positive in almost 100% of the cases. While at a confidence level of 95%, TCE is always equal to or larger than TCM, at a confidence level of 90% we find one case where the difference is negative. Again, these finding indicate that in the case of stocks, TCE nearly always leads to higher capital requirements than does TCM.

We now turn a look at monthly data, which is important because for certain risk positions there are no daily data available (e.g., in case of hedge funds). In Table 4 we consider monthly data for the 500 stocks in the S&P. For this purpose, we transformed the 3,784 daily returns into 180 monthly returns.

Insert Table 4 here

Again, TCE is larger than TCM, although the difference between them is relatively low for confidence levels higher than 95%. The perfect matching that occurs at confidence levels higher than 99% is not surprising because we consider the lowest 1% values within 180 monthly returns, meaning that we look at a very small number of data points and thus mean and median are similar.

Figure 2 shows again the frequency distribution of (TCE-TCM)/TCE. The horizontal axis presents the relative difference between TCE and TCM and the vertical axis the number of funds.

Insert Figure 2 here

An important aspect of Figure 2 is the much higher number of negative values for (TCE-TCM)/TCE than were seen in Figure 1. At a confidence level of 90%, this finding would indicate that for 6.20% of the 500 stocks, the capital requirement derived by way of TCM would be larger than those derived by way of TCE.

(2) Mutual Funds

In Table 5, we consider data from 1,374 mutual funds reporting monthly net of fee returns in the period from January 1996 to December 2005. The values presented are again arithmetic averages across 1,374 TCE and TCM values. We first calculate the TCE and TCM for each of the 1,374 funds using 120 monthly observations and then determine the average across the 1,374 values of TCE and TCM.

Insert Table 5 here

Again, TCE is larger than TCM and, also again, the difference between the two is negligible at confidence levels greater than 95%. Figure 3 shows the frequency distribution of (TCE-TCM)/TCE.

We see an even greater number of negative values for the relative difference than it was the case in Figures 1 and 2. The range increases, e.g., for α = 95%, the range is 71.95% compared to 28.50% with the stocks on a daily basis (Figure 1) and 54.06% with stocks on a monthly basis (Figure 2).

Insert Figure 3 here

(3) Hedge Funds

In Table 6, we consider 205 hedge funds reporting monthly net of fee returns in the period from January 1996 to December 2005.

Insert Table 6 here

Table 6 reveals that the differences between TCE and TCM are greater in regard to hedge funds than they are with regard to either stocks or mutual funds. At a confidence level of 95%, TCE is 13.05% higher than TCM; for mutual funds, this difference was only 6.22%. Again, the perfect congruence of results at confidence levels higher than 98.5% is due to purely technical reasons because we are only able to consider 120 monthly returns. Although it would be helpful to extend the time series to more than 120 data points, doing so is not possible as no such data are available. Hedge funds publish return information only once every month and historical hedge fund returns (prior to 1996) suffer from backfilling bias (see Eling [9]), among other problems. In short, there simply are no other hedge fund data available for risk management purposes.

Again, we look at the data frequency in order to capture the variability of the values around their means. Figure 4 shows the frequency distribution of (TCE-TCM)/TCE for the 205 hedge funds. The horizontal axis shows the relative difference between TCE and TCM and the vertical axis shows the number of funds.

Insert Figure 4 here

We find that, on average, the TCE leads to higher capital requirements than the TCM; however, in about 7.84% of the cases TCM is higher than TCE (for α = 90%). Although the hedge funds returns analyzed are not normally distributed and exhibit heavy tails, the distribution of (TCE-TCM)/TCE is not too different in the case of hedge funds that was observed for mutual funds.

3.3 Data Fitting and Simulation Results

What causes the observed empirical differences between TCE and TCM? To answer this question, we now turn to simulated data and consider distributions commonly used in finance as well as popular equity return models. As we will show, the degree of fatness in the tail of a distribution, as well as the availability of data (and the resulting estimation risk), both have a major influence when comparing TCE and TCM.

(1) Distributions Commonly Used in Finance

In this section, we analyze six distributions commonly used in finance to discover which of the two risk measures is the largest. The six distributions are: the normal, the skew normal (see Adcock [2]), the student t (see Kole et al. [12]), the skew student t (see Theodossiou [14]), the normal inverse Gaussian (NIG; see Barndorff-Nielsen [5]), and the hyperbolic (see Eberlein et al. [8]). In contrast to the normal distribution the alternative distributions are able to account for skewness (skew normal), kurtosis (student t), or even for both skewness and kurtosis (skew student t, NIG, hyperbolic) in returns. Some of these distributions are related to each other, e.g., the student t, the normal inverse Gaussian, and the hyperbolic all belong to the class of generalized hyperbolic distributions.

We proceed in three steps. In Step 1, we determine parameters of these six distributions for the 3,784 daily S&P 500 index returns using maximum likelihood estimation. In Step 2, we conduct a simulation study based on 500,000 random numbers from these distributions and calculate TCE and TCM to discover which one is the largest. In Step 3, we vary the simulation settings so as to glean more insight into differences between TCE and TCM, e.g., by using different degrees of skewness and kurtosis for a given distributional assumptions. Here, we also vary the number of simulated returns to see to what extent the empirical results are driven by estimation risk. An alternative to the simulation study conducted here might be to find closed-form solutions for TCE and TCM under the various distributions. However, to our knowledge, there are no closed-form solutions for TCE and TCM for most of the distributions and deriving same is beyond the scope of this paper (closed-form solutions for TCE for normal and skew normal are presented in Vernic [15] and Bolance et al. [6]). We thus rely on simulation. All simulations have been implemented using the R package gyhp and sn.

The results of Steps 1 and 2 are presented in Table 7. Column 1 presents the distributions, columns 2 to 4 the TCE (for 99%, 95% and 90%), columns 5 to 7 the TCM, and columns 8 to 10 the relative difference between these two measures ((TCE-TCM)/TCE).[1]

Insert Table 7 here

[1] The log-likelihood value of the maximum likelihood estimation illustrates that the distributions have different goodness of fit and thus different quality to describe the data. However, the aim of this paper is to illustrate how TCE and TCM behave under various distributions and not to answer the question which one of the distributions fits best the empirical data. This is the reason why we do not remove one of the distributions based on goodness of fit considerations. All estimation results are available upon request.

In each simulated distribution, TCE is higher than TCM, thus indicating a clear relationship between the two. At 99% confidence level the normal distribution (skewness and kurtosis of zero) results in a relative difference of 3.40%; for the skew normal, the relative difference is only slightly higher (3.48%). However, for the skew student t distribution, which exhibits fat tails, the relative difference is 14.93%. It thus appears that the relative difference between TCE and TCM is greater with fatter tails. Furthermore, the Table shows that the lower the confidence level, the higher is the relative difference between TCE and TCM.

To illustrate this relationship in more detail, Figure 5 shows the relative difference ((TCE-TCM)/TCE) for the skew student t with varying degrees of freedom (between 2 and 10 degrees of freedom, the degrees of freedom in the maximum likelihood estimation of the S&P 500 index was 3.39). The lower the degree of freedom, the fatter the tail of the skew student t distribution. With infinitely high degrees of freedom, the skew student t converges toward the skew normal distribution. With a skewness of zero and infinitely high degrees of freedom, the skew student t converges toward a normal distribution.

Insert Figure 5 here

Again, TCE is always larger than TCM. The lower the degree of freedom, the fatter the tail of the distribution and the greater the relative difference between TCE and TCM. For example, with two degrees of freedom and at a confidence level of 99%, TCE is 10.04% and TCM is 7.06%, resulting in a relative difference of 29.66%. For 10 degrees of freedom, the relative difference is only 5.94%. The analysis thus confirms that the relative difference between TCE and TCM positively depends on the degree of fatness. We also again see that the lower the confidence level, the higher the relative difference: with 10 degrees of freedom and at a 90% confidence level, the relative difference is 9.53%, whereas it is only 5.03% at a 99.90% confidence level. We thus conclude that the farther we go into the tail of a distribution, the lower is the relative difference between the two risk measures.

The simulation results reveal a clear relationship between TCE and TCM. The relative difference might depend on the degree of fatness in the tail of the distribution, but TCE is always higher than TCM. In the empirical analysis, however, we found no such clear-cut relationship between the two risk measures. A possible reason might be estimation risk caused by an insufficient number of empirical observations. To test this hypothesis, we now vary the number of simulated returns. In a simulation setting, the estimation risk could be quantified by the standard error of the simulated mean. Figure 6 shows the simulation results for the skew student t distribution with 500,000 (dark line), 10,000 (dashed line), and 1,000 simulated returns (light line).

Insert Figure 6 here

The dark line is relatively smooth, but a kink can be observed at very high confidence levels. Even with 500,000 simulations, considering TCE and TCM at a 99.90% confidence level means that the mean and median are compared for only 500 simulated values, which may not be too robust, as indicated by the kink in Figure 6. The effect is much more severe for the case of 10,000 simulated returns, i.e., at a 99.90% confidence level, a comparison of mean and median is based on merely 10 values. Substantially more variation is thus observed with the dashed line. For the case of 1,000 simulated returns (light line) and at a confidence level of 99.90%, this means calculating mean and median based on one value, and thus at this confidence level the difference has to be zero (see Figure 6, light blue line at 99.90%). For lower confidence levels, variations in the simulation results are much more extreme and we find some instances where TCE is lower than TCM. Note that in the empirical part of this paper, we considered 3,784 returns with daily data and 120 returns with monthly data. It is thus not surprising to find no clear-cut relationship in empirical data, especially when monthly data are considered.

The estimation results thus broadly depend on the nature of empirical data used to estimate the risk measures. If only a few observations are used to estimate TCE and TCM, estimation risk can result in a situation where TCE < TCM. However, in theory and with a sufficiently high number of observations, there is a clear-cut relationship between TCE and TCM and this relationship is TCE > TCM. But this is not necessarily true in practice, where a low number of returns are used for estimation; an example are hedge funds, for which only monthly data are available. Practitioners thus face substantial estimation risk, which may be a strong argument for using TCM instead of TCE. TCM is the more robust measure, i.e., TCM is likely to result in better estimates than TCE because it exhibits lower estimation risk. TCE is not very robust and is sensitive to model assumptions and outliers (see Heyde et al. [10]).

Yamai and Yoshiba (2005) compare estimation risk of expected shortfall (i.e., TCE) and value at risk under a generalized Pareto distribution. We adopt their methodology to analyze estimation risk of TCE and TCM. We evaluate estimation risk by obtaining 10,000 estimates of the two risk measures. To obtain each estimate, we run Monte Carlo simulations with a sample size of 10,000 random numbers, assuming that the underlying loss will follow the normal, skew-normal, student t, or skew student t distribution; we consider TCE and TCM at the 99% confidence level. We iterate this procedure 10,000 times, thus obtaining 10,000 estimates of TCE and TCM. We then calculate the average value, the standard deviation (i.e., the standard error of the simulated mean), and the 5% and 95% confidence levels of the estimates. Panel A of Table 8 summarizes the results. Panel B of

Table 8 presents the results when only a sample size of 1,000 random numbers is used for estimation.

Insert Table 8 here

For normally distributed data, there is not much difference between TCE and TCM. For the fat-tailed skew student t distribution, however, the standard error of TCE is substantially higher than the standard error of TCM. For example, with the skew student t distribution the standard deviation of TCM is 0.17%, whereas it is 0.25% with TCE. The standard deviation increases when a smaller sample size of only 1,000 random numbers (Panel B) is considered, but overall the results for TCE and TCM are very comparable.

In Figure 7, we again vary the degrees of freedom of the skew student t to discover to what extent estimation risk depends on the fatness of the distribution. The figure shows the estimation risk for TCE and TCM, here given by the relative standard error of the simulation, i.e., the standard error is divided by the mean. We also include value at risk so that our results can be compared with those of Yamai and Yoshiba [16].

Insert Figure 7 here

Estimation risk is much higher at low degrees of freedom than it is at higher degrees of freedom. We thus conclude that if returns are normally distributed, estimation risk has no substantial influence in the comparison between TCE and TCM. However, in a world of non-normal returns with significant fat tails, differences in estimation risk could be highly significant, with TCM estimation risk lower than that of TCE; TCM thus provides a robust alternative to TCE.

(2) ARCH, GARCH, and Jump Diffusion Models

We complement our analysis by considering popular equity return models, such as the ARCH/GARCH models (under different assumptions regarding the distribution of innovations) and the jump diffusion model. We first estimate the parameters of the different models based on the S&P 500 index data and then compare TCE and TCM under the estimated models (with 500,000 simulated returns).[2] We again vary our assumptions with regard to the number of simulated returns to see whether the above results also hold in the context of equity return models (see Figure 9).

---

[2] Again, statistical tests could indicate the goodness of fit of the various models and distributions considered here. But our aim is to illustrate the TCE and TCM behavior under different models that have been used in literature and not to find the model that fits best. This is the reason why we do not remove one of the models based on goodness of fit considerations.

Figure 8 shows (TCE-TCM)/TCE for ARCH (1) and GARCH (1,1) processes calibrated to S&P 500 data under different distributions (normal, skew normal, GED, skew GED, student t, skew student t). All calculations were made with the R package fGarch.

Insert Figure 8 here

Again, we find that TCE is higher than TCM. The fatter the tail, the higher the relative difference between TCE and TCM. For example, looking at the left side of Figure 8, the difference ranges between 4% and 8% for the normal and skew normal distributions at confidence levels from 99.9% to 90%. For the student and skew student, however, the relative difference is between 12% and 16%. Figure 8 also confirms that the lower the confidence level, the higher the relative difference between TCE and TCM.

Figure 9 illustrates the results for a jump diffusion process with Bernoulli distributed jumps; drift and volatility of the process were calibrated to the S&P 500 index. The data fitting and simulations were conducted with the R package EMJumpDiffusion. We present results for 1,000, 10,000, and 500,000 simulated returns.

Insert Figure 9 here

The results presented in Figure 9 are very comparable to the results presented in Figure 6, which confirms the above findings.

4. Conclusion

If the decision as to which risk measure is the most appropriate is based on axiomatical grounds, the choice will be guided by the purpose of risk management, i.e., different methods will be better depending on if it is internal or external risk management that is desired. Two options in this context are the tail conditional expectation (TCE) and the tail conditional median (TCM). A relevant question posed by practitioners is whether the choice of risk measure has an effect on the resulting capital requirements. The aim of this paper was to answer this question. Empirical investigations on possible mismatching between TCE and TCM were performed for different asset classes (stocks and the S&P 500 index; mutual funds; hedge funds) and different time horizons (daily and monthly data).

Heyde et al. [10], p. 23, based on one set of data (the S&P 500 index), conclude that the differences between TCE and TCM are highly significant, both theoretically and practically. We performed a broader theoretical and empirical analysis; our results (1) confirm the Heyde et al. findings and (2) emphasize the relevance of TCM. We find that TCE and TCM can lead to very different capital requirements. On average, and at standard confidence levels, TCE is about 10% larger than TCM. However, there is no clear-cut relationship between TCE and TCM; depending on the tails and the number of data points used in the estimation, we find that TCM can be higher than TCE. In theory, capital requirements derived by TCE are larger than those derived by TCM, but this is not necessarily true when insufficient amounts of real data are employed. Estimation risk is thus an important consideration when comparing TCE and TCM and, as we show, TCM might be the more robust choice in this context. The findings support the proposal to use objective-oriented risk measures and underline the relevance of TCM as a robust natural risk statistic.

References

[1] Acerbi, C., and Tasche, D. (2002). On the Coherence of Expected Shortfall. *Journal of Banking and Finance*, 26, 1487–1503.

[2] Adcock, C. J. (2007). Extensions of Stein's lemma for the skew-normal distribution. *Commun. Statist. – Theory & Methods*, 36, 1661–1671.

[3] Ahmed, S., Filipovic, D., and Svindland, G. (2008). A Note on Natural Risk Statistics. *Operations Research Letters*, 36, 662–664.

[4] Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent Measures of Risk. *Mathematical Finance* 9, 203–228.

[5] Barndorff-Nielsen, O. E. (1997). Normal Inverse Gaussian Distributions and the Modelling of Stock Returns. *Scandinavian Journal of Statistics*, 24, 1–13.

[6] Bolance, C., Guillen, M., Pelican, E., and Vernic, R. (2008): Skewed bivariate models and nonparametric estimation for the CTE risk measure. *Insurance: Mathematics and Economics*, 43, 386–393.

[7] Cont, R., Deguest, R., and Scandolo, G. (2007). Robustness and Sensitivity Analysis of Risk Measurement Procedures. Financial Engineering Report No. 2007-06, Center for Financial Engineering, Columbia University.

[8] Eberlein, E, Keller, U., and Prause, K. (1998). New Insights into Smile, Mispricing and Value at Risk: the Hyperbolic Model. *Journal of Business*, 71, 371–405.

[9] Eling, M. (2008). Does the Measure Matter in the Mutual Fund Industry? *Financial Analysts Journal* 64 (3), 54–66.

[10] Heyde, C. C., Kou, S., and Peng, X. H. (2007). What Is a Good Risk Measure: Bridging the Gaps Between Data, Coherent Risk Measures, and Insurance Risk Measures. Financial Engineering Report 2007-09, Center of Financial Engineering, Columbia University.

[11] Jarrow, R. A., and Purnanandam, A. K. (2005). A Generalized Coherent Risk Measure: The Firm's Perspective. *Finance Research Letters*, 2, 23–29.

[12] Kole, E., Koedijk, K., and Verbeek, M. (2007). Selecting copulas for risk management. *Journal of Banking & Finance*, 31, 2405–2423.

[13] Ogryczak, W., and Zawadzki, M. (2002). Conditional Median: A Parametric Solution Concept for Location Problems. *Annals of Operations Research*, 110, 167–181.

[14] Theodossiou, P. (1998). Financial Data and the Skewed Generalized T Distribution. *Management Science*, 44, 1650–1661.

[15] Vernic, R. (2006) Multivariate skew-normal distributions with applications in insurance. *Insurance: Mathematics and Economics*, 38, 413–426.

[16] Yamai, Y., and Yoshiba, T. (2005): Value-at-risk versus expected shortfall: A practical perspective. *Journal of Banking & Finance*, 29, 997–1015.

	Cross-Sectional Analysis (across funds)				
Time-Series Analysis	Mean	Median	Standard Deviation	Minimum	Maximum
A. 500 stocks from the S&P 500 (monthly data, January 1990 to December 2004)					
JB rejection: 56.83% (66.67%) at 1% (5%) level					
Mean value (%)	1.65	1.37	1.09	-0.61	6.77
Standard deviation (%)	10.60	9.39	4.57	2.65	28.97
Skewness	0.21	0.17	0.67	-2.43	7.03
Excess kurtosis	2.22	1.19	4.56	-0.42	65.90
B. 1,314 mutual funds (monthly data, January 1996 to December 2005)					
JB rejection: 19.84% (26.73%) at 1% (5%) level					
Mean value (%)	0.53	0.49	1.19	–9.52	9.79
Standard deviation (%)	4.70	4.50	2.43	0.06	29.31
Skewness	–0.29	–0.32	0.76	–9.50	9.38
Excess kurtosis	0.76	0.11	4.35	–7.19	100.83
C. 205 hedge funds (monthly data, January 1996 to December 2005)					
JB rejection: 80.00% (84.39%) at 1% (5%) level					
Mean value (%)	1.05	0.97	0.54	-0.52	4.36
Standard deviation (%)	4.14	3.46	3.14	0.23	25.49
Skewness	-0.07	0.00	1.49	-6.40	7.87
Excess kurtosis	6.14	2.59	10.41	-1.33	81.90

Table 1: Descriptive statistics for stocks, mutual funds, and hedge funds

Alpha	TCE	TCM	TCE-TCM	(TCE-TCM)/TCE
99.90%	0.0610	0.0631	-0.0022	-3.58%
99.50%	0.0406	0.0364	0.0042	10.38%
99.00%	0.0346	0.0306	0.0039	11.36%
98.50%	0.0315	0.0285	0.0030	9.54%
98.00%	0.0294	0.0267	0.0028	9.42%
97.50%	0.0279	0.0254	0.0026	9.16%
97.00%	0.0266	0.0242	0.0024	9.06%
96.50%	0.0255	0.0233	0.0022	8.75%
96.00%	0.0246	0.0223	0.0022	9.15%
95.50%	0.0238	0.0217	0.0021	8.84%
95.00%	0.0231	0.0210	0.0021	9.11%
90.00%	0.0183	0.0163	0.0021	11.23%
85.00%	0.0154	0.0137	0.0017	11.00%
80.00%	0.0133	0.0111	0.0023	16.88%
70.00%	0.0104	0.0082	0.0022	21.24%
60.00%	0.0083	0.0062	0.0021	25.22%
50.00%	0.0067	0.0045	0.0022	32.76%

Table 2: TCE and TCM for S&P 500 index (daily data)

Alpha	TCE	TCM	TCE-TCM	(TCE-TCM)/TCE
99.90%	0.1536	0.1455	0.0081	5.12%
99.50%	0.1011	0.0894	0.0117	10.99%
99.00%	0.0833	0.0730	0.0104	11.94%
98.50%	0.0745	0.0652	0.0093	12.11%
98.00%	0.0685	0.0600	0.0085	12.15%
97.50%	0.0643	0.0562	0.0080	12.21%
97.00%	0.0608	0.0531	0.0077	12.33%
96.50%	0.0581	0.0507	0.0074	12.48%
96.00%	0.0557	0.0485	0.0071	12.55%
95.50%	0.0536	0.0467	0.0069	12.70%
95.00%	0.0518	0.0451	0.0067	12.85%
90.00%	0.0406	0.0348	0.0058	14.14%
85.00%	0.0343	0.0290	0.0053	15.55%
80.00%	0.0300	0.0249	0.0051	16.87%
70.00%	0.0239	0.0192	0.0047	19.78%
60.00%	0.0193	0.0149	0.0045	23.27%
50.00%	0.0122	0.0076	0.0046	42.92%

Table 3: TCE and TCM for 500 stocks from the S&P 500 (daily data)

Alpha	TCE	TCM	TCE-TCM	(TCE-TCM)/TCE
99.90%	0.3017	0.3017	0.0000	0.00%
99.50%	0.3017	0.3017	0.0000	0.00%
99.00%	0.2754	0.2754	0.0000	0.00%
98.50%	0.2582	0.2500	0.0083	3.24%
98.00%	0.2442	0.2355	0.0087	3.67%
97.50%	0.2334	0.2224	0.0110	4.93%
97.00%	0.2245	0.2136	0.0109	4.99%
96.50%	0.2164	0.2047	0.0118	5.55%
96.00%	0.2090	0.1962	0.0128	6.08%
95.50%	0.2031	0.1895	0.0135	6.66%
95.00%	0.2012	0.1874	0.0138	6.82%
90.00%	0.1628	0.1465	0.0163	10.00%
85.00%	0.1397	0.1229	0.0168	12.04%
80.00%	0.1228	0.1061	0.0167	13.71%
70.00%	0.0978	0.0816	0.0161	16.59%
60.00%	0.0785	0.0630	0.0155	20.18%
50.00%	0.0623	0.0472	0.0151	25.12%

Table 4: TCE and TCM for 500 stocks from the S&P 500 (monthly data)

Alpha	TCE	TCM	TCE-TCM	(TCE-TCM)/TCE
99.90%	0.1647	0.1647	0.0000	0.00%
99.50%	0.1647	0.1647	0.0000	0.00%
99.00%	0.1457	0.1457	0.0000	0.00%
98.50%	0.1457	0.1457	0.0000	0.00%
98.00%	0.1344	0.1267	0.0077	5.74%
97.50%	0.1344	0.1267	0.0077	5.74%
97.00%	0.1262	0.1194	0.0069	5.44%
96.50%	0.1195	0.1120	0.0075	6.29%
96.00%	0.1195	0.1120	0.0075	6.29%
95.50%	0.1139	0.1068	0.0071	6.22%
95.00%	0.1139	0.1068	0.0071	6.22%
90.00%	0.0909	0.0826	0.0083	9.19%
85.00%	0.0769	0.0673	0.0096	12.47%
80.00%	0.0668	0.0569	0.0099	14.76%
70.00%	0.0521	0.0418	0.0103	19.77%
60.00%	0.0324	0.0223	0.0101	31.11%
50.00%	0.0121	0.0015	0.0106	87.29%

Table 5: TCE and TCM for 1,374 mutual funds (monthly data)

Alpha	TCE	TCM	TCE-TCM	(TCE-TCM)/TCE
99.90%	0.1330	0.1330	0.0000	0.00%
99.50%	0.1330	0.1330	0.0000	0.00%
99.00%	0.1110	0.1110	0.0000	0.00%
98.50%	0.1110	0.1110	0.0000	0.00%
98.00%	0.0980	0.0889	0.0091	9.32%
97.50%	0.0980	0.0889	0.0091	9.32%
97.00%	0.0893	0.0805	0.0087	9.79%
96.50%	0.0829	0.0724	0.0105	12.68%
96.00%	0.0829	0.0724	0.0105	12.68%
95.50%	0.0778	0.0676	0.0101	13.05%
95.00%	0.0778	0.0676	0.0101	13.05%
90.00%	0.0590	0.0499	0.0091	15.38%
85.00%	0.0487	0.0396	0.0091	18.65%
80.00%	0.0415	0.0334	0.0082	19.66%
70.00%	0.0315	0.0236	0.0079	24.97%
60.00%	0.0242	0.0169	0.0073	30.24%
50.00%	0.0182	0.0111	0.0070	38.70%

Table 6: TCE and TCM for 205 hedge funds (monthly data)

Distribution	TCE			TCM			(TCE-TCM)/TCE		
Confidence level	99%	95%	90%	99%	95%	90%	99%	95%	90%
Empirical	3.46%	2.30%	1.83%	3.07%	2.09%	1.63%	11.33%	9.28%	11.13%
Normal	2.66%	2.05%	1.74%	2.57%	1.94%	1.63%	3.40%	5.16%	6.42%
Skew normal	2.74%	2.09%	1.77%	2.64%	1.98%	1.65%	3.48%	5.30%	6.73%
Student t	4.17%	2.41%	1.84%	3.57%	2.04%	1.53%	16.85%	18.27%	20.26%
Skew student	4.31%	2.45%	1.86%	3.66%	2.06%	1.53%	14.93%	16.05%	17.71%
NIG	3.22%	2.24%	1.81%	3.02%	2.06%	1.62%	6.68%	8.88%	11.51%
Hyperbolic	3.67%	2.54%	2.04%	3.45%	2.32%	1.82%	6.10%	9.45%	12.02%

Table 7: Parameter estimation and TCE and TCM for various tail distributions

	Normal		Skew normal		Student t		Skew student		NIG		Hyperbolic	
	TCE	TCM	TCE	TCM	TCE	TCM	TCE	TCM	TCE	TCM	TCE	TCM
Panel A: 10,000 simulation runs with 10,000 random numbers												
Mean	2.66%	2.57%	2.74%	2.64%	4.13%	3.56%	4.27%	3.66%	3.21%	3.02%	3.67%	-3.46%
St. Deviation	0.05%	0.05%	0.05%	0.05%	0.23%	0.16%	0.25%	0.17%	0.09%	0.09%	0.10%	0.10%
5% Quantile	2.74%	2.65%	2.82%	2.73%	4.53%	3.85%	4.71%	3.96%	2.70%	2.86%	2.66%	-2.83%
95% Quantile	2.59%	2.49%	2.66%	2.56%	3.78%	3.31%	3.90%	3.39%	3.36%	3.17%	3.84%	-3.62%
Panel B: 10,000 simulation runs with 1,000 random numbers												
Mean	2.64%	2.57%	2.72%	2.64%	4.09%	3.60%	4.23%	3.70%	3.18%	3.02%	3.64%	3.47%
St. Deviation	0.15%	0.15%	0.16%	0.16%	0.73%	0.51%	0.77%	0.55%	0.27%	0.26%	0.31%	0.31%
5% Quantile	2.89%	2.83%	2.98%	2.92%	5.35%	4.50%	5.59%	4.70%	8.49%	8.71%	8.49%	8.86%
95% Quantile	2.41%	2.33%	2.48%	2.39%	3.16%	2.88%	3.24%	2.93%	3.64%	3.48%	4.18%	4.01%

Table 8: Estimation risk of TCE and TCM (99% confidence level)

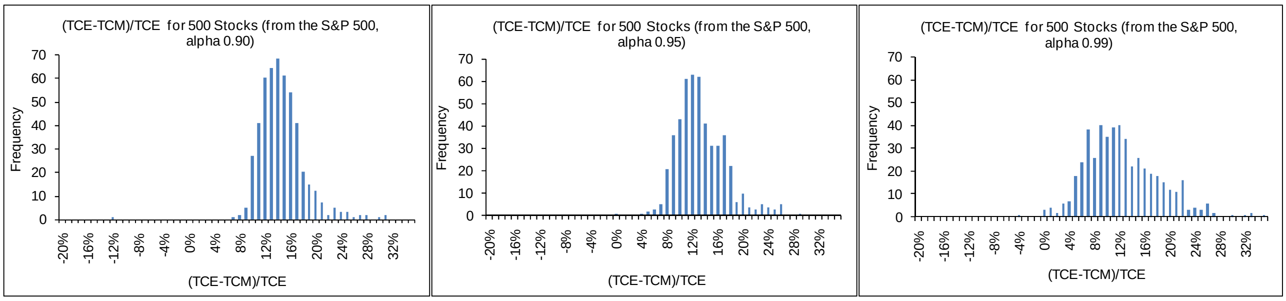

Alpha	Min	25%	50% (Median)	75%	Max	Mean	Range	% negative
90.00%	-12.04%	11.83%	13.72%	15.90%	30.39%	14.14%	42.43%	0.20%
95.00%	0.00%	10.26%	12.24%	15.19%	28.50%	12.85%	28.50%	0.00%
99.00%	-4.66%	7.89%	11.13%	15.53%	34.19%	11.94%	38.86%	0.20%

Figure 1: Frequency distribution of (TCE-TCM)/TCE (500 stocks, daily data)

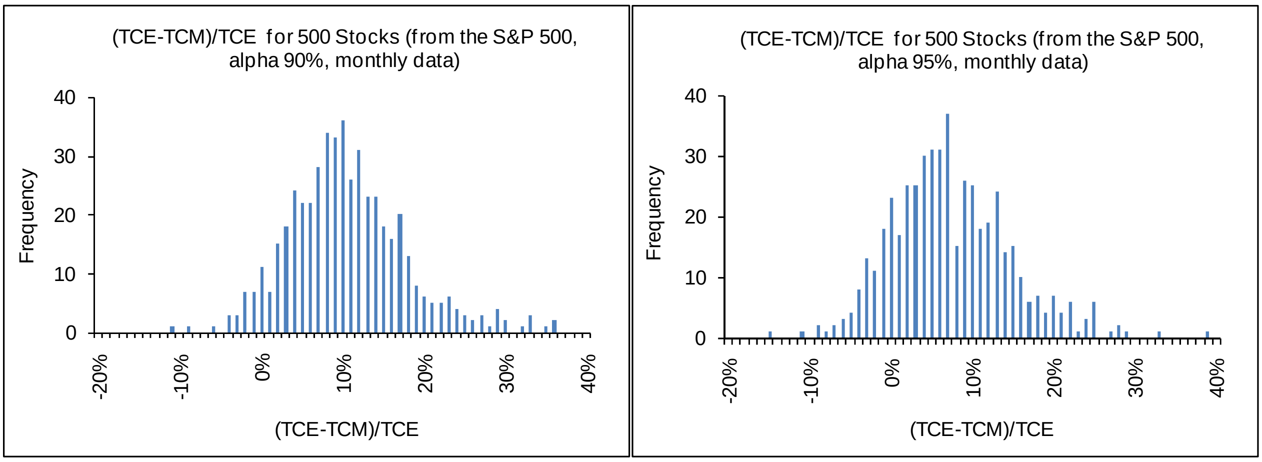

Alpha	Min	25%	50% (Median)	75%	Max	Mean	Range	% negative
90.00%	-11.83%	5.29%	9.21%	13.93%	35.61%	10.00%	47.44%	6.20%
95.00%	-15.75%	1.88%	6.06%	11.37%	38.31%	6.82%	54.06%	16.00%

Figure 2: Frequency distribution of (TCE-TCM)/TCE (500 stocks, monthly data)

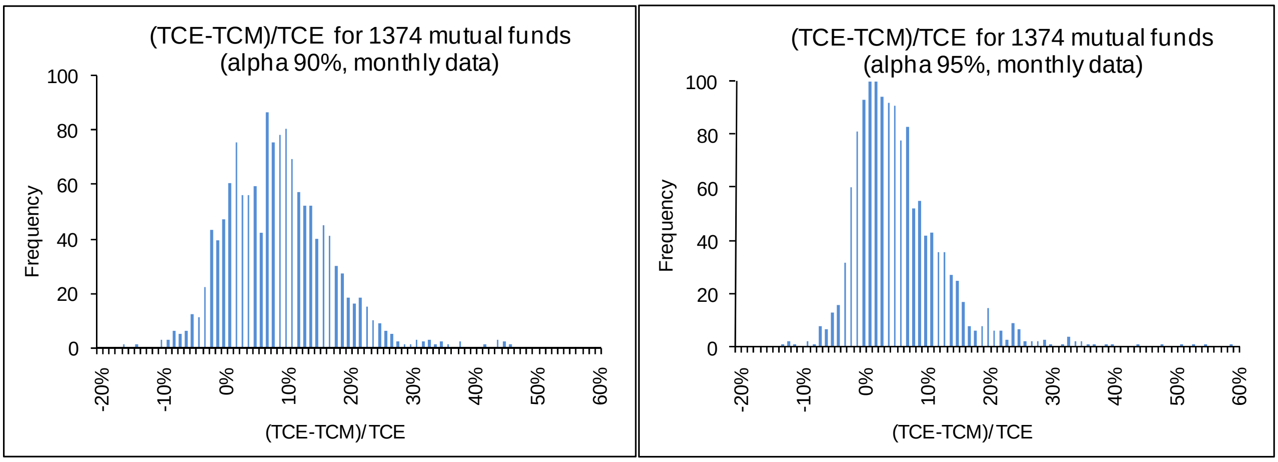

Alpha	Min	25%	50% (Median)	75%	Max	Mean	Range	% negative
90.00%	-16.10%	2.25%	7.91%	13.07%	45.26%	8.34%	61.36%	14.13%
95.00%	-13.05%	0.20%	3.76%	8.32%	58.90%	5.25%	71.95%	22.56%

Figure 3: Frequency distribution of (TCE-TCM)/TCE (1,374 mutual funds, monthly data)

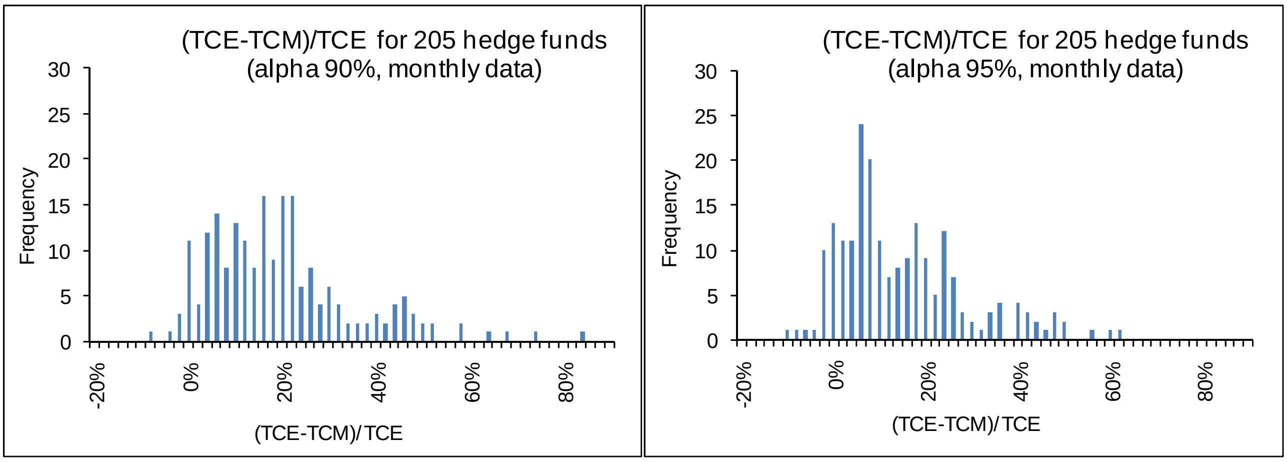

Alpha	Min	25%	50% (Median)	75%	Max	Mean	Range	% negative
90.00%	-9.59%	7.07%	16.03%	24.93%	82.22%	18.64%	91.81%	7.84%
95.00%	-11.37%	4.12%	9.73%	21.75%	61.97%	13.79%	73.34%	13.17%

Figure 4: Frequency distribution of (TCE-TCM)/TCE (205 hedge funds, monthly data)

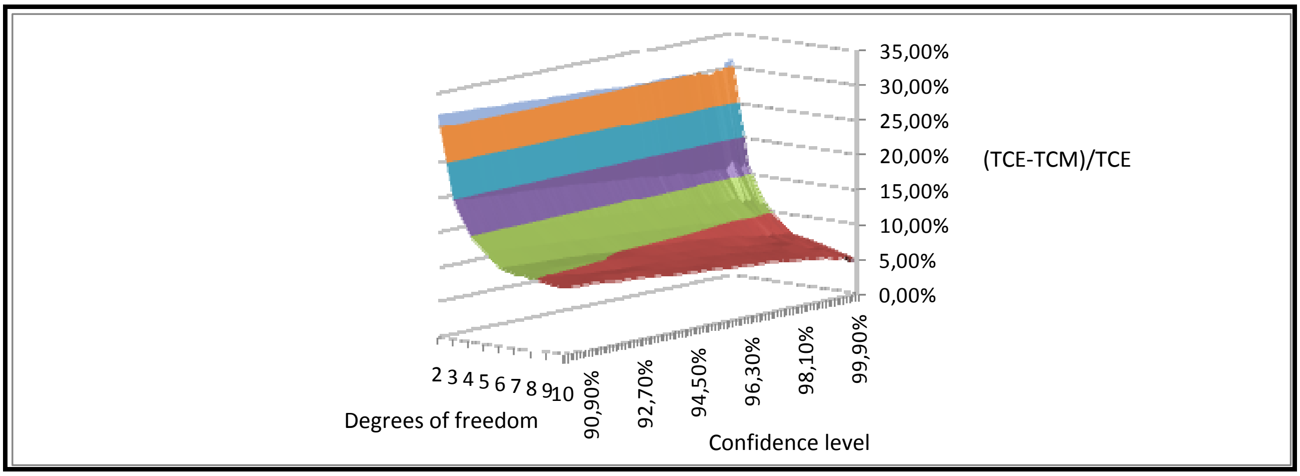

Figure 5: (TCE-TCM)/TCE for the skew student distribution and varying degrees of freedom

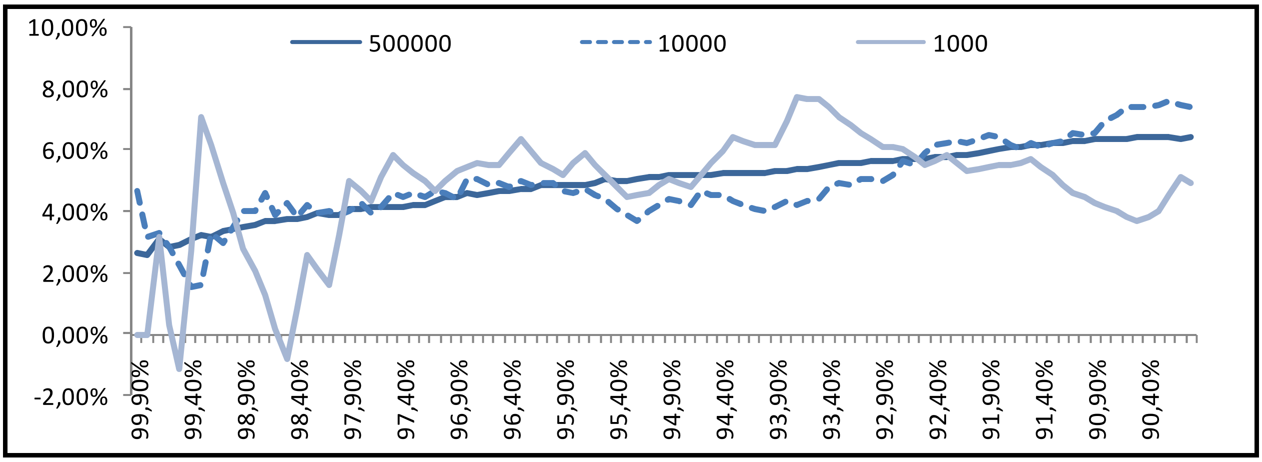

Figure 6: (TCE-TCM)/TCE for the skew student distribution and varying number of simulated returns

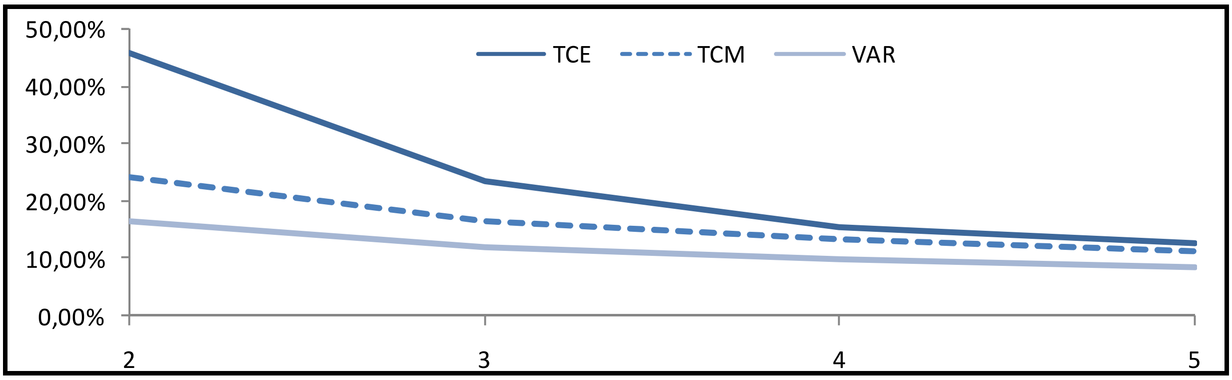

Figure 7: Relative standard deviation of estimates (99% confidence level)

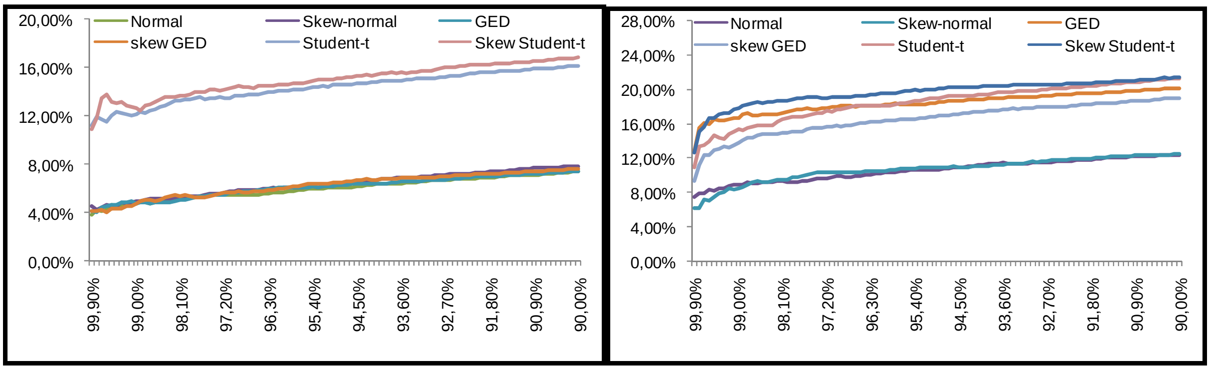

Figure 8: (TCE-TCM)/TCE for ARCH (1) process (left) and GARCH (1,1) process (right) with different distribution of innovations

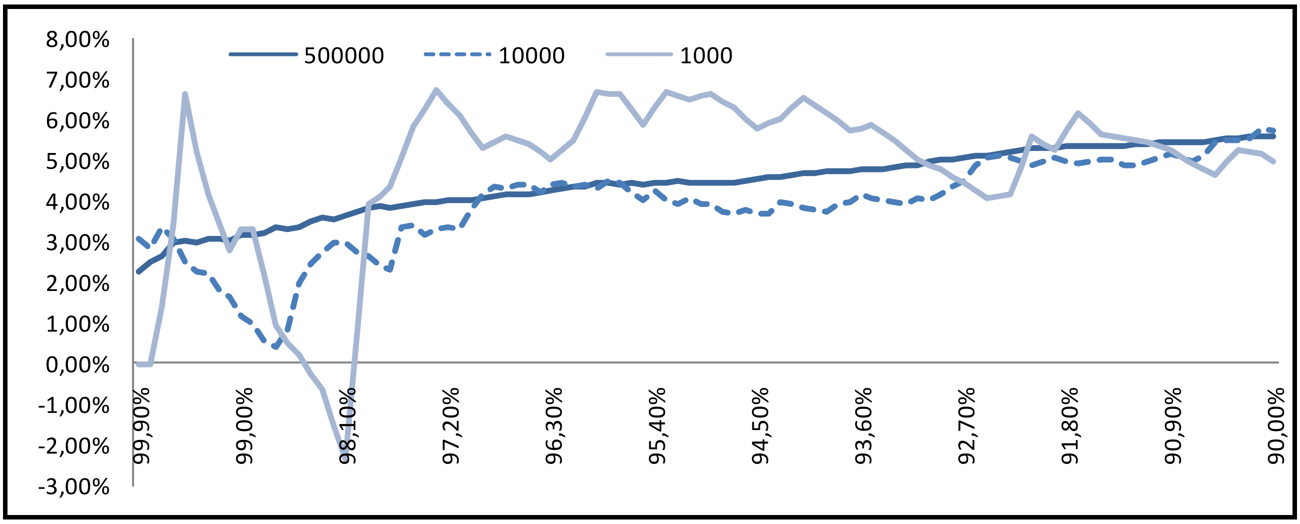

Figure 9: Jump diffusion process for 1,000, 10,000, and 500,000 simulated returns