10-2020

TERES: Tail Event Risk Expectile Shortfall

Andrija MIHOCI
German University in Cairo

Wolfgang Karl HARDLE
Singapore Management University

Cathy Yi-Hsuan CHEN
University of Glasgow

Citation

MIHOCI, Andrija; HARDLE, Wolfgang Karl; and CHEN, Cathy Yi-Hsuan. TERES: Tail Event Risk Expectile Shortfall. (2020). *Quantitative Finance*. 21, (3), 449-460.
Available at: https://ink.library.smu.edu.sg/skbi/8

Published in Quantitative Finance, 2020, 21 (3), 449-460. DOI: 10.1080/14697688.2020.1786151

TERES: Tail Event Risk Expectile Shortfall

ANDRIJA MIHOCI*†‡, WOLFGANG KARL HÄRDLE§¶‖**†† and CATHY YI-HSUAN CHEN§‡‡

†German University in Cairo, Berlin, Germany
‡Berlin School of Economics and Law, Berlin, Germany
§Humboldt-Universität zu Berlin, Berlin, Germany
¶Xiamen University, Xiamen, People's Republic of China
‖Singapore Management University, Singapore, Singapore
**Charles University, Praha, Czech Republic
††National Chiao Tung University, Hsinchu, Taiwan
‡‡University of Glasgow, Glasgow, UK

We propose a generalized risk measure for expectile-based expected shortfall estimation. The generalization is designed with a mixture of Gaussian and Laplace densities. Our plug-in estimator is derived from an analytic relationship between expectiles and expected shortfall. We investigate the sensitivity and robustness of the expected shortfall to the underlying mixture parameter specification and the risk level. Empirical results from the US, German and UK stock markets and for selected NASDAQ blue chip companies indicate that expected shortfall can be successfully estimated using the proposed method on a monthly, weekly, daily and intra-day basis using a 1-year or 1-day time horizon across different risk levels.

Keywords: Expected shortfall; Expectiles; Value at risk; Risk management; Tail risk

JEL Classification: C13, G10, G31

1. Introduction

Understanding financial risk plays an important role in risk management and statistics. Researchers and practitioners aim to control the risk in changing tail dependence because portfolios and tail event risk depend upon the tail structure. In our paper, we rely on expectile-based risk management and focus on modeling the expected shortfall, which is a tail risk measure that is defined as the expected value of a random variable below a given threshold; for example, an expectile or quantile level.

Expected shortfall is a coherent risk measure that takes diversification and risk aggregation effects into account, see Artzner *et al.* (1998), Acerbi *et al.* (2001), Delbaen (2002), and Acerbi (2007) who discuss and quantify expected shortfall risk exposure. Although the subadditivity of expected shortfall has academic support, coherent risk measures have been criticized; see for example Cont *et al.* (2010), Wang (2016). From the perspective of coherent risk measures, expected shortfall is, nevertheless, superior to quantile-based measures. Researchers and practitioners may find challenges in robust estimation and prediction backtesting, due to the elicitability property (Ziegel 2016). Therefore, we emphasize here the clearly advantageous distributional robustness property of expected shortfall in the context of Huber (1964). More specifically, we provide evidence of expected shortfall robustness superiority relative to the underlying distributional assumptions by considering an expectile-based normal–Laplace contaminated mixture in the framework of the robustness analysis of Huber (1964).

In practice, expected shortfall successfully captures tail structures; see for example McNeil *et al.* (2015), Franke *et al.* (2019). Its unique combination of desirable properties (Kusuoka 2001) has been recognized in portfolio selection (Bassett *et al.* 2004) and risk management. Recently, expected shortfall has been recommended by the Basel Committee on Banking Supervision (2013) for internal use and there is supportive evidence that it allows fund managers to comply with investor preferences more accurately when compared to value at risk or variance-based risk measures (Koenker 2005).

This paper proposes a method for estimating expected shortfall and compares the results with value at risk for generalized tail events such as expectiles and quantiles. We think of these two tail events because of the distinct attributes that

*Corresponding author. Email: Andrija.Mihoci.Dr@gmail.com

they possess. For instance, quantiles are elicitable and not coherent, whereas expectiles are both elicitable and coherent. M-quantiles, conceptually analogous to expectiles, can be considered to be one of the admissible tail events. Breckling and Chambers (1988) piloted the invention of M-quantiles. Otto-Sobotka *et al.* (2019) further generalize them with a consideration of heteroscedasticity. The resulting risk measures, when given selected thresholds with known statistical properties, can change in terms of smoothness or monotonicity. For a given tail distribution, expected shortfall can be considered in terms of the concept of generalized tail event (Breckling and Chambers 1988), which helps us to analyze the expected shortfall variation stemming from a change of the underlying distributional framework. For example, the effect of 'lengthening or shortening the distribution tail' on expectile driven expected shortfall can be studied by employing a normal–Laplace mixture. Our approach essentially quantifies the consequences of potential model misspecifications and biased risk indicators in risk management which may result from the lack of tail-relevant observations, see for example McNeil and Frey (2000).

Concerning expectile-based expected shortfall properties, we focus on the effects of varying tail structures. A common argument for the use of value at risk is its low degree of sensitivity with respect to the tail structure change. Stahl *et al.* (2012) advocate the usage of a more appropriate robustness notion and argue that, considering risk management applications, expected shortfall represents a robust risk measure. However, a Gaussian framework may only be suitable as long as the interest is on modeling the distributional center of a financial return time series; see for example Hull and White (1998). Most importantly, for a downside tail risk measure (Acerbi and Tasche 2002), it is necessary to be aware of the implications of deviations from the theoretical normal model on the resulting risk measures. Tail Event Risk Expectile Shortfall (TERES) makes use of expectiles and the normal–Laplace contaminated mixture (Huber 1964), to gain a better understanding of the tail risk sensitivity of expected shortfall.

The research questions of our study are therefore as follows: How to utilize expectiles in modeling expected shortfall based on tail information? What are the key advantages of employing TERES using expectiles and the normal–Laplace mixture setup in practice? How sensitive are the results obtained by the proposed framework relative to quantile-based approaches?

In our paper, we suggest using expectiles and a normal–Laplace mixture to integrate tail information in risk modeling. Furthermore, the statistical properties of tail events can be transferred to modeling expected shortfall as an analytic function of expectiles. The key advantages of TERES include the smoothness of the results, more robust findings compared to value at risk figures when changing the (extreme) risk level and advances in the tail risk assessment of data from a relatively short time series. We also demonstrate that TERES achieves a tailor-made balance between tail sensitivity and robustness. Finally, the proposed framework is successfully employed for the risk management of financial time series on a broader set of data frequencies, such as on a monthly, weekly, daily and intra-day level.

The paper is structured as follows. After considering tail event risk examples in section 2, the TERES framework is introduced in section 3. Empirical results concerning expected shortfall analysis are provided in section 4. Finally, section 5 concludes. The reader can find the R-codes tailored for this study in the quantlets platform www.quantlet.de.

2. Tail event risk examples

Financial markets are characterized by time changing risk structures. Therefore, the uncertainty measurement associated with extreme tail events has to be carefully treated. In this section, we briefly describe the data series that will be used in risk modeling throughout the rest of this paper. We focus on three examples, namely: equity risk exposure assessment, high-frequency data risk modeling and portfolio allocation. For each example, we discuss the tail risk characteristics, display associated series dynamics and offer recommendations for risk management. A comprehensive risk assessment for the introduced time series using our TERES framework is conducted in the empirical part, see section 4.

Example 1 *Equity risk exposure* Equity returns exhibit non-stationarity in the variance and higher order moments. Figure 1 displays the daily DAX 30 (2540 trading days), FTSE 100 (2557 trading days) and S&P 500 (2518 trading days) index return series from 1 January 2007 to 31 December 2016. Even after time-dynamic mean–variance standardization using a GARCH$(1, 1)$ model by Engle (1982) and Bollerslev (1986) heavy tails still prevail. Therefore, for the sake of brevity, the results are not reported here and are instead available from the authors upon request.

It is evident that the center, certainly not the tail, of the data distribution may be well modeled using a conditional Gaussian setup. Therefore, in practice it is reasonable to choose to model the risk of the observed return series using a normal–Laplace mixture.

Example 2 *High-frequency finance* Intraday return time series share similar features to those observed on stock indices. Consider the mid-quote price returns of 16 largest NASDAQ companies from three sectors, namely: technology (eight companies), consumer services (four companies) and health care (four companies), see table 1. Because the 'Brexit' referendum result has had a significant influence on stock market movements, we correspondingly focus on the order book activities on 27 June 2016 (when the S&P 500 was at its lowest level after the vote) and 30 June 2016 (the upward movement of the S&P 500 series). A comprehensive description and statistical analysis of the limit order book data is provided by Mihoci (2017).

Similar to the equity example, the time series suggests a modeling approach that does not solely depend upon a Gaussian framework or upon a fixed probability level (i.e. modeling risk using value at risk).

Example 3 *Portfolio allocation* A successful asset allocation takes the portfolio tail structure into account. Consider, for

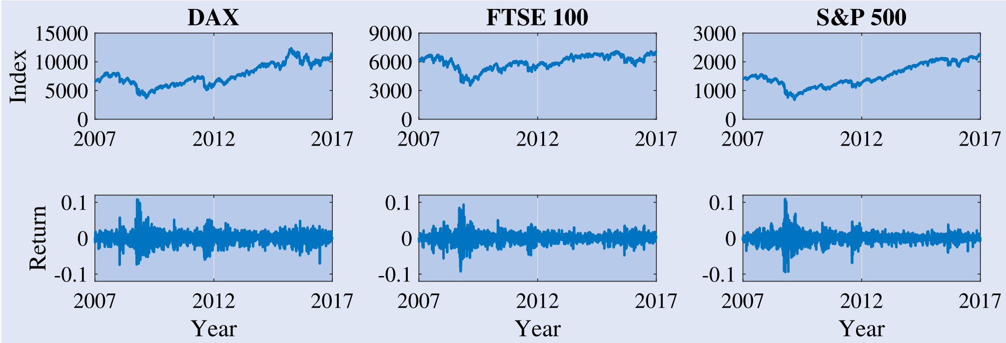

Figure 1. Time series of daily DAX, FTSE 100 and S&P 500 values and returns from 1 January 2007 to 31 December 2016.

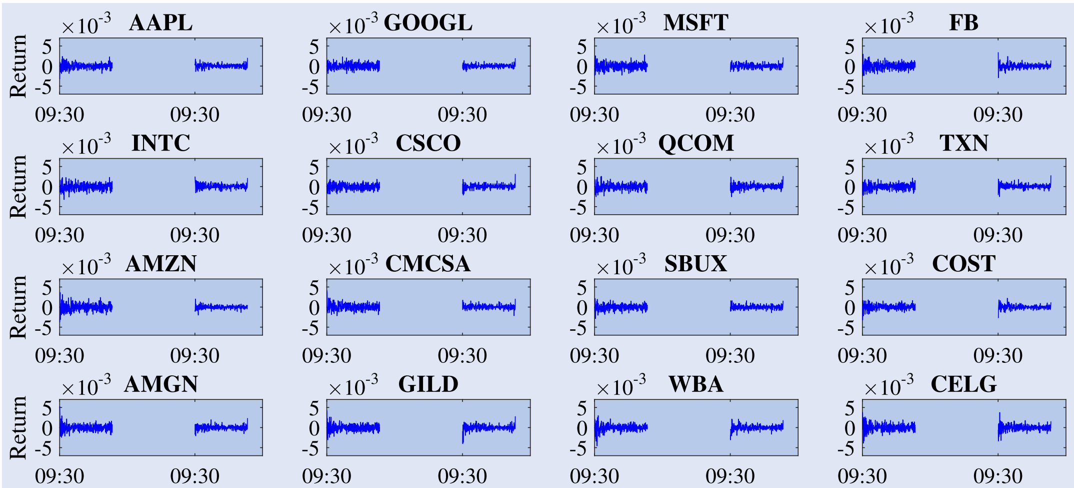

Figure 2. Time series of intra-day mid-quote price returns for selected NASDAQ companies on 27 and 30 June 2016. On each day, we utilize data at every minute between 9:31 and 16:00 (390 observations).

example, the Tail Event Driven ASset allocation (TEDAS) framework that was developed by Härdle *et al.* (2015). Generally speaking, TEDAS enables us to broaden the classical portfolio allocation by using quantile regression between the index series and portfolio constituents at selected fixed tail levels (e.g. 5%, 15%, 25%, 35% and 50%). Consequently, TEDAS allows us to construct tail event driven portfolios. The returns of two TEDAS applications that significantly outperform the underlying index series' performance are displayed in figure 3:

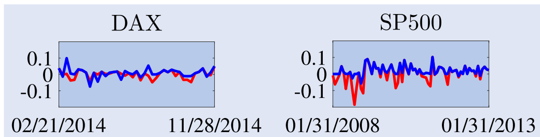

Figure 3. Returns of two TEDAS applications (blue) and their respective benchmarks (red). Left: Weekly German stock portfolios vs. DAX 30 index (41 trading weeks), right: monthly mutual funds portfolio returns vs. S&P 500 index (73 trading months).

(i) German equity market: weekly portfolios are formed out of a pool of 125 selected small and mid cap stocks for the span from 21 February 2014 to 28 November 2014 (41 trading weeks),

(ii) Worldwide mutual funds market: monthly portfolios are considered from 1 February 2008 to 31 January 2013 (73 trading months); here 583 mutual funds have been selected.

Table 1. Selected NASDAQ companies.

Industry: Technology (8)		Industry: Consumer Services (4) and Health Care (4)	
Apple Inc.	AAPL	Amazon.com, Inc.	AMZN
Alphabet Inc.	GOOGL	Comcast Corporation	CMCSA
Microsoft Corporation	MSFT	Starbucks Corporation	SBUX
Facebook, Inc.	FB	Costco Wholesale Corporation	COST
Intel Corporation	INTC	Amgen Inc.	AMGN
Cisco Systems, Inc.	CSCO	Gilead Sciences, Inc.	GILD
QUALCOMM Incorporated	QCOM	Walgreens Boots Alliance, Inc.	WBA
Texas Instruments Inc.	TXN	Celgene Corporation	CELG

For further details, we refer to Härdle *et al.* (2015). As frequently encountered in risk management, both cases result in relatively short time series due to weekly or monthly rebalancing schemes. However, in practice, a substantial amount of information about the underlying risk processes is demanded. This effect is intensified by the risk assessment requirement to accurately analyze the extreme outer tail. As explained later on, TERES offers a suitable framework for evaluating (infrequently re-constructed) tail event driven allocation strategies.

3. Tail Event Risk Expectile based Shortfall

The expected shortfall of a random variable Y with cumulative distribution function $F(\cdot)$ and probability density function $f(\cdot)$ is defined as

$$s_\eta = \mathsf{E}[Y \mid Y < \eta], \tag{1}$$

where η represents a given threshold. Expected shortfall is a coherent downside risk measure and with the threshold selected as the τ-th quantile (a common choice in practice) belongs to the class of spectral risk measures (Overbeck 2004).

The integration of tail information into expected shortfall is in our paper conducted through expectiles and a normal–Laplace mixture. In this section, we first discuss the concept of generalized tail events that is used to express the expected shortfall as a function of the expectile and tail information. Thereafter, we model the expectile risk level τ as a function of a given underlying distributional specification. Finally, based on the proposed mixture model, we discuss our tail event driven risk modeling framework.

3.1. *Expected shortfall and expectiles*

The expectile and the quantile can be seen as two potentially different parameters of a location model for a random variable Y with constant θ and innovation ε; that is $Y = \theta + \varepsilon$. Given the check function

$$\rho_{\tau,\gamma}(u) = |\tau - \mathbf{I}\{u < 0\}|\,|u|^\gamma \tag{2}$$

the τ-th quantile and the τ-th expectile, respectively, are identified as location parameter θ:

$$q_\tau = \arg\min_\theta \mathsf{E}\left[\rho_{\tau,1}(Y - \theta)\right] \tag{3}$$

$$e_\tau = \arg\min_\theta \mathsf{E}\left[\rho_{\tau,2}(Y - \theta)\right]. \tag{4}$$

In terms of a location model, quantiles take solely the sign of residuals into account, while expectiles consider their magnitude. As evident from risk management practice, expectiles additionally do not necessarily lead to substantially increased margin requirements relative to quantile-based approaches.

In risk management practice, expected shortfall estimation utilizes expectiles and quantiles figures. For this purpose, consider the expectile e_{w_τ} that equals the τth quantile; that is $q_\tau = e_{w_\tau}$. Relating the expectile and the quantile by their corresponding levels w_τ and τ allows us to recover an indicator of the average distance of the tail probability mass from the quantile threshold q_τ. Knowledge of the expectile level w_τ then enables us to express the expected shortfall as follows:

$$s_{q_\tau} = s_{e_{w_\tau}} = e_{w_\tau} + \frac{e_{w_\tau} - \mathsf{E}[Y]}{1 - 2w_\tau}\frac{w_\tau}{\tau} \tag{5}$$

see Taylor (2008). The resulting expected shortfall specification s_{q_τ} can be seen as a function of the expectile level, weighted by the ratio $\frac{w_\tau}{\tau}(1 - 2w_\tau)^{-1}$.

Our proposed TERES framework, which is based on a normal–Laplace mixture, relies on equation (5). In a nonparametric context, the distribution function of the data would be used in expected shortfall estimation. In a recent study, Schulze Waltrup *et al.* (2015) employed a spline-based estimation technique. As an incremental advantage of our framework, beyond the nonparametric case, we confirm whether 'lengthening the distributional tail' leads to an increase in the considered ratio. We additionally employ the TERES framework as an instrument for comparison of the expected shortfall and value at risk properties.

3.2. *Expectile–quantile transformation and tail structures*

The characteristics of w_τ in the context of expected shortfall are rarely studied. If the underlying distribution $F(\cdot)$ is known, then the expectile–quantile transformation is given by

$$w_\tau = \frac{LPM\,(q_\tau) - q_\tau \tau}{2\,\{LPM\,(q_\tau) - q_\tau\tau\} + q_\tau - \mathsf{E}[Y]} \tag{6}$$

where

$$LPM\,(u) = \int_{-\infty}^{u} yf\,(y)\,\mathrm{d}y \tag{7}$$

denotes the lower partial moment. Expression (6) holds for any $F(\cdot)$ due to the one-to-one mapping between expectiles and quantiles (Jones 1994). The mass distance indicator $\frac{w_\tau}{\tau}$ quantifies the average distance of the probability mass of $f(\cdot)$ from the τ th quantile, which in TERES is used in tail focused inference.

As mentioned previously, the expectile–quantile relationship (6) depends on the underlying distributional specification. To illustrate the concept, consider the following four cases: (a) uniform, (b) standard normal, (c) Laplace and (d) stable distributions; as noted above, this paper focuses on the normal–Laplace mixture density.

(a) Let Y denote a standard uniform distributed variable $Y \sim U[0, 1]$, $f(y) = \mathbf{I}\{0 \le y \le 1\}$ with $q_\tau = \tau$. The expectile-quantile transformation level equals

$$w_\tau = \frac{\tau^2}{2\left(\tau^2 - \tau + 0.5\right)} \tag{8}$$

Graphical inspection of (8) reveals that in this case the τ-expectile is closer to the distributional center than the τ-quantile.

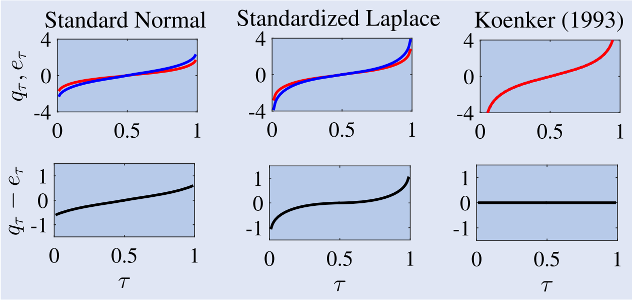

Figure 4. Top: Quantile (blue) and expectile (red) values for different τ level. Bottom: Difference between quantiles and expectiles. The three cases include: standard normal, standardized Laplace and the cdf given in (10).

(b) Risk underestimation in practice often occurs when a normality assumption is applied in modeling heavy-tailed data. For a standard normal distribution with probability density function $\varphi(\cdot)$, the LPM equals $-\varphi(\cdot)$ and therefore

$$w_\tau = \frac{-\varphi(q_\tau) - q_\tau \tau}{-2\{\varphi(q_\tau) + q_\tau \tau\} + q_\tau - \mathsf{E}[Y]} \tag{9}$$

A closer look at the expectile and quantile values for a given level τ shows that the expectile values are closer to the distributional center as compared to quantiles, see figure 4.

(c) For the slightly heavier-tailed standardized Laplace (i.e. double exponential) distribution the LPM corresponds to

$$LPM\ (u) = \begin{cases} \exp(u)(u-1), & \text{if } u < 0 \\ \exp(-u)(-u-1), & \text{else} \end{cases}$$

which exhibits lower (larger) values for given level τ in the left (right) tail as compared to the standard normal case, see figure 4. On average the mass distance indicator $\frac{w_\tau}{\tau}$ increases as the degree of tail heaviness increases. The conclusion is less clear in the extreme outer tail. Additionally, we may ask whether the τ-expectile is generally closer to the distributional center than the τ-quantile. Given the distribution by Koenker (1993)

$$F(x) = \begin{cases} 0.5 - 0.5\left(1 - \frac{4}{4+x^2}\right)^{0.5}, & \text{if } x < 0 \\ 0.5 + 0.5\left(1 - \frac{4}{4+x^2}\right)^{0.5}, & \text{else} \end{cases} \tag{10}$$

both coincide; that is $e_\tau = q_\tau$ for $\tau \in [0,1]$. In practice, this distribution exhibits considerably heavier tails than typically selected distributions. Thus we conclude that for a realistic degree of tail heaviness, $w_\tau < \tau$ holds for the left distributional tail and for the normal–Laplace mixture employed in this paper.

(d) The class of stable distributions (Mandelbrot 1963, Fama 1965, Cizek *et al.* 2011) is briefly discussed here. The class captures a wide array of probability distribution functions (including normal and Cauchy), as reflected by a stability index $\alpha \in (0,2]$ and a skewness parameter $\beta \in [-1,1]$. Location and scale are controlled by parameters μ and σ. A closed-form solution for the cumulative distribution function does not exist. Therefore, an approximation of the lower tail for this class of distributions is given by Nolan (2015)

$$f(y|\alpha,\beta,\sigma) \approx c_\alpha \sigma^\alpha |\beta - 1|(-y)^{-\alpha-1} \tag{11}$$

$$c_\alpha = \alpha \sin\left(\frac{\pi\alpha}{2}\right)\frac{\Gamma(\alpha)}{\pi} \tag{12}$$

For $\alpha \in [1,2]$, the LPM of stable distributions is approximated by

$$LPM\ (u) \approx \frac{c_\alpha \sigma^\alpha |\beta - 1|}{1-\alpha}(-u)^{1-\alpha}, \quad \text{as } u \to -\infty. \tag{13}$$

Consider the standardized symmetric case ($\sigma = 1, \beta = 0$). The stability index α allows the controlled increase of the degree of tail-heaviness. As α increases the tail becomes lighter and the mass distance indicator is (again) found to decrease. Even more interestingly, the example points out the LPM as the parameter that is most dependent on the selection of the underlying distribution, especially in the extreme outer tail. Consider, for example, $q_\tau = -5$ where the LPM lies within the range [6.23e-19, 0.01] and the corresponding τ in [0.0002, 0.04]. In this situation, the expected shortfall, which depends on the LPM, is certainly less insensitive than the value at risk.

3.3. *Tail event risk modeling framework*

The TERES framework enables us to quantify risk effects of progressing from value at risk to expected shortfall. We will now proceed with an analysis of the connection between the distributional (tail-) information provided by the mass distance index and the expected shortfall.

The non-normality apparent on many time series renders a comparison of the robustness properties of the expected shortfall and the value at risk inside a realistic distributional neighborhood. Even though our focus here is on financial applications, this holds true for most statistical studies.

Normal–Laplace mixture

Here, we will model the tail structure by adopting a mixture of distributions. To analyze the robustness properties of expected shortfall under deviations from a normality assumption, we follow the approach outlined in Huber (1964). The underlying idea is to emphasize the use of a neighborhood for robustness analysis. The aim is to find an estimator that displays good performance and which sticks loosely to an underlying model assumption (e.g. normality) but does rely on the theoretical model to be literally and exactly true.

Consider for this purpose the random variable Y with probability density function

$$f(y|\delta) = (1-\delta)\varphi_{\theta_1}(y) + \delta h_{\theta_2}(y) \tag{14}$$

where φ denotes a normal probability density function with $\theta_1 = (\mu_1, \sigma_1)^\top$ and h another continuous symmetric distribution with parameter set θ_2. The parameter subset for the Laplace distribution is re-parametrized as $\theta_2 = (\mu_2, \sigma_2)^\top$, with

$$h_{\theta_2}(u) = \frac{1}{\sigma_2\sqrt{2}} \exp\left\{-\sqrt{2}\frac{|u-\mu_2|}{\sigma_2}\right\}. \quad (15)$$

Discussion

Now that a suitable distribution h is at hand, we can present the distributional environment (14) in more detail. The most light tailed scenario results for $\delta = 0$, where a normal distribution is always employed. For increasing values of δ the distributional tail is 'lengthened' until at $\delta = 1$ the more heavy tailed Laplace distribution is selected. In the example of a standard power exponential distribution by Rigby and Stasinopoulos (2004), the normal and Laplace cases would correspond to the power parameter equal to 2 and 1, respectively, see section 4. By construction, the normal–Laplace mixture mechanism characterizes the dominance of each mixture component, in the spirit of Huber (1964). If one is, for example, interested in understanding slight departures from the normality assumption, then the focus will be on the neighborhood resulting from δ values that are close to 0.

Obviously, the expectile, quantile and the corresponding expected shortfall will vary within the selected neighborhood. Within a δ-contaminated neighborhood, as in (14), w_τ becomes a function of the mixture parameter δ. In settings with contamination that exhibits heavier tails than under normality, such as Laplacian, this relationship is monotonically decreasing in δ for low risk levels (such as $\tau = 0.05$ or $\tau = 0.10$) if the expected shortfall is calculated strictly based only on the assumed distribution. Our approach, however, makes use of (5) by combining pre-estimated quantiles or expectiles with the relevant expectile level w_τ. This turns out to be specifically tailored to our needs, because we are interested in comparing the robustness properties of expected shortfall and value at risk which originate in the influence of w_τ. Thus we keep the τ-quantile fixed while varying w_τ over the selected distributional neighborhood.

Example

Consider the standardized Laplacian contamination case as defined in (15) and (14) with parameter set $(\mu_1 = 0, \sigma_1 = 1, \mu_2 = 0, \sigma_2 = 1)^\top$. The upper left-hand side figure 5 shows the resulting expected shortfall for levels of contamination δ ranging from 0 (normal) to 1 (Laplace case). To analyze the neighborhood around the Laplace scenario in the outer tail, we provide the expected shortfall for risk level $\tau = 0.005$ in the top right-hand side of figure 5. We note that the strongest variation of the expected shortfall (i.e. the absolutely highest slope) is observed at the normal scenario. This is very reasonable. For example, Huber (1964) points out that 'lengthening' the tail structure while starting from a light tailed case has a more pronounced influence on the estimator than a 'shortening' of tails. Nevertheless, as (5) shows, that knowledge of the quantile could potentially provide a significant reduction in expected shortfall variation. In financial examples, we find this effect to be very strong; as displayed by the empirical applications of TERES in section 4.

Expectile–quantile relationship

Here, we will briefly discuss the sensitivity of the expectile–quantile transformation. Regarding the check function (2) it becomes apparent that the expectile–quantile transformation (6) results as a special case for $\gamma = 2$ and

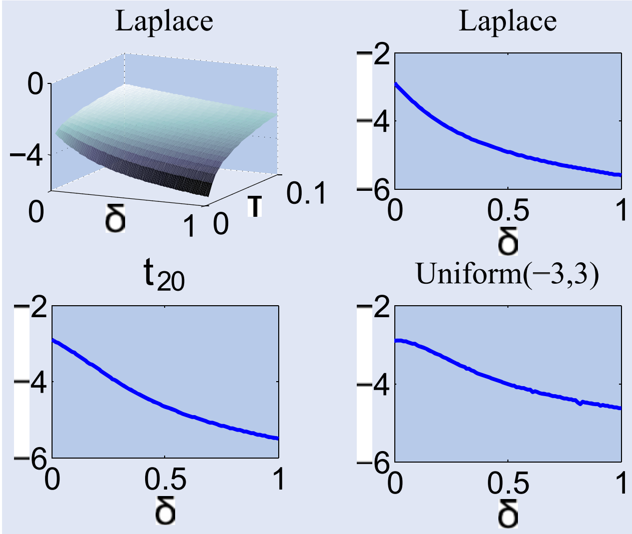

Figure 5. Theoretical s_{q_τ} for mixtures of standard Normal and standardized Laplace (top left). Furthermore, the Laplace, t with 20 df and Uniform$(-3, 3)$ contaminated normal environments of degree δ for the fixed (tail) risk level $\tau = 0.5\%$ are depicted.

the question arises as to how strong the influence of the norm selection is here. In the given framework, the norm is controlled by fixing a $\gamma \geq 1$. Note that

$$w(\alpha, \gamma) = \frac{\int_{-\infty}^{q_\alpha} |y - q_\alpha|^{\gamma-1}\, dF(y)}{\int_{-\infty}^{\infty} |y - q_\alpha|^{\gamma-1}\, dF(y)}, \quad \gamma \geq 1$$

holds for the special case of quantiles $w(\alpha, 1) = \alpha$. The analysis reduces to the consideration of a value $\gamma > 1$ (i.e. two for the expectile case) and we then investigate the convergence rate, with respect to γ of the term $|y - q_\alpha|^{\gamma-1}$ towards one. Generally we observe exponential convergence for $|y - q_\alpha| > 1$ whereas for $|y - q_\alpha| < 1$ root convergence results. For $|y - q_\alpha| > 1 \approx 1$, nearly linear convergence towards the quantile case is observed. Therefore, we conclude that when using standardized data and modest risk levels the framework shows a reasonable degree of stability in the selection of the loss function.

4. Empirical study

In this section, TERES is employed to investigate the properties of expected shortfall given a time series of empirical quantile estimates. In estimation, we utilize the expectation–maximization algorithm (Hartley and Rao 1967, Dempster *et al.* 1977). The essential steps of the algorithm in the case of the normal–Laplace mixture parameter estimation are for convenience summarized in table 2. As introduced in section 2, we employ TERES in the risk assessment of different asset classes. First, we model the risk dynamics of daily stock indices, namely the DAX, the FTSE and the S&P 500. Thereafter a high-frequency example is presented where we analyze NASDAQ data over 2 days following the Brexit referendum.

Table 2. Expectation-maximization (EM) algorithm for the estimation of normal-Laplace mixture parameters for a selected time series.

EM: Normal-Laplace Mixture Estimation Algorithm

* Select a (return) time series of n observations, $r_t, t = 1, \ldots, n$

* Step a: Estimate parameter vectors $\theta_1 = (\mu_1, \sigma_1)^\top$ and $\theta_2 = (\mu_2, \sigma_2)^\top$ using maximum-likelihood method

Here, $\widetilde{\theta}_1 = (\widetilde{\mu}_1, \widetilde{\sigma}_1)^\top$ and $\widetilde{\theta}_2 = (\widetilde{\mu}_2, \widetilde{\sigma}_2)^\top$ with

$$\widetilde{\mu}_1 = n^{-1} \sum_{t=1}^{n} r_t, \widetilde{\sigma}_1 = n^{-1} \sqrt{\sum_{t=1}^{n} (r_t - \widetilde{\mu}_1)^2}, \widetilde{\mu}_2 = n^{-1} \sum_{t=1}^{n} r_t \text{ and } \widetilde{\sigma}_2 = n^{-1} \sqrt{2} \sum_{t=1}^{n} |r_t - \widetilde{\mu}_2|$$

* Step b: For each data point, we calculate two density values. The first (second) value represents the density value evaluated at the data point from the normal (Laplace) distribution given the parameters from Step a, for the latter see expression (15)

$$\widetilde{p}_{1t}^{(0)} = \varphi_{\widetilde{\theta}_1}(r_t) \text{ and } \widetilde{p}_{2t}^{(0)} = h_{\widetilde{\theta}_2}(r_t), t = 1, \ldots, n$$

* Step c.0: E-Step at iteration $i = 0$: Is it more likely that an observation belongs to the normal or the Laplace distribution if the mixture parameter equals $\widetilde{\delta}^{(0)}$?

$$\text{Calculate weights } \widetilde{w}_{1t}^{(0)} = \frac{\widetilde{p}_{1t}^{(0)}(1 - \widetilde{\delta}^{(0)})}{\widetilde{p}_{1t}^{(0)}(1 - \widetilde{\delta}^{(0)}) + \widetilde{p}_{2t}^{(0)} \widetilde{\delta}^{(0)}} \text{ and } \widetilde{w}_{2t}^{(0)} = \frac{\widetilde{p}_{2t}^{(0)} \widetilde{\delta}^{(0)}}{\widetilde{p}_{1t}^{(0)}(1 - \widetilde{\delta}^{(0)}) + \widetilde{p}_{2t}^{(0)} \widetilde{\delta}^{(0)}},$$

$$\widetilde{w}_{1t}^{(0)} + \widetilde{w}_{2t}^{(0)} = 1, t = 1, \ldots, n$$

The mixture parameter at next iteration is thereafter found based on these estimated weights (expressed as a function of the estimated probabilities), namely

$$\widetilde{\delta}^{(1)} = n^{-1} \sum_{t=1}^{n} \{1 - \widetilde{w}_{1t}^{(0)}\} = n^{-1} \sum_{t=1}^{n} \widetilde{w}_{2t}^{(0)}$$

* Step d.1: M-Step at iteration $i = 1$: Estimate parameter vectors of a probability weighted time series

Consider the time series $\widetilde{r}_{1t}^{(1)} = \widetilde{w}_{1t}^{(1)} r_t^{(0)}$ and $\widetilde{r}_{2t}^{(1)} = \widetilde{w}_{2t}^{(1)} r_t^{(0)}$, $r_t^{(0)} = r_t, t = 1, \ldots, n$

$$\widetilde{\mu}_1^{(1)} = \sum_{t=1}^{n} \widetilde{r}_{1t}^{(1)} / \sum_{t=1}^{n} \widetilde{w}_{1t}^{(1)}, \widetilde{\sigma}_1^{(1)} = \sqrt{\sum_{t=1}^{n} (\widetilde{r}_{1t}^{(1)} - \widetilde{\mu}_1^{(1)})^2 / \sum_{t=1}^{n} \widetilde{w}_{1t}^{(1)}}, \ \widetilde{p}_{1t}^{(1)} = \varphi_{\widetilde{\theta}_1^{(1)}}(\widetilde{r}_{1t}^{(1)})$$

$$\widetilde{\mu}_2^{(1)} = \sum_{t=1}^{n} \widetilde{r}_{2t}^{(1)} / \sum_{t=1}^{n} \widetilde{w}_{2t}^{(1)}, \widetilde{\sigma}_2^{(1)} = \sqrt{2} \sum_{t=1}^{n} |\widetilde{r}_{2t}^{(1)} - \widetilde{\mu}_2^{(1)}| / \sum_{t=1}^{n} \widetilde{w}_{2t}^{(1)}, \widetilde{p}_{2t}^{(1)} = h_{\widetilde{\theta}_2^{(1)}}(\widetilde{r}_{2t}^{(1)})$$

* Step c.i: E-Step at iteration $i = 1, \ldots, m-1$: Calculate weights $\widetilde{w}_{1t}^{(i)}$ and $\widetilde{w}_{2t}^{(i)}$, and find the mixture parameter estimate $\widetilde{\delta}^{(i+1)}$

* Step d.i: M-Step at iteration $i = 2, \ldots, m$: Find the estimates $\widetilde{\theta}_1^{(i)}$ and $\widetilde{\theta}_2^{(i)}$

* Step e: Iterate steps c and d until convergence and after m iterations finally provide the estimates

$$\widehat{\delta} = \widetilde{\delta}^{(m)}, \widehat{\theta}_1 = \widetilde{\theta}_1^{(m)} \text{ and } \widehat{\theta}_2 = \widetilde{\theta}_2^{(m)}$$

The third example utilizes TERES for risk assessment of pre-selected monthly and weekly rebalanced portfolios with time changing component weights.

4.1. Equity risk exposure

Consider the selected equity indices between 2007 and 2016 and observe the return time series and the corresponding estimated probabilities, i.e. weights, that the returns follow a Laplace distribution, see figure 6, and the resulting TERES estimated density in figure 7. It is evident that during the financial crisis (2008) and the European debt crisis (2012) in risk modeling, more importance is given to the Laplace distribution. Looking at the two sub-periods – 2007 to 2011 and 2012 to 2016 in table 3 – reveals the same finding from the perspective of the estimated mixture parameter δ (above 0.50 in any case).

From the considered time series, DAX and S&P 500 returns, at times the FTSE returns are related with relatively higher estimated mixture parameter readings. This finding is also reflected by the density plot in figure 7 related to the DAX series risk assessment. Interestingly, more probability mass is located in the left tail. This may resemble the discussed standard power exponential distribution with the

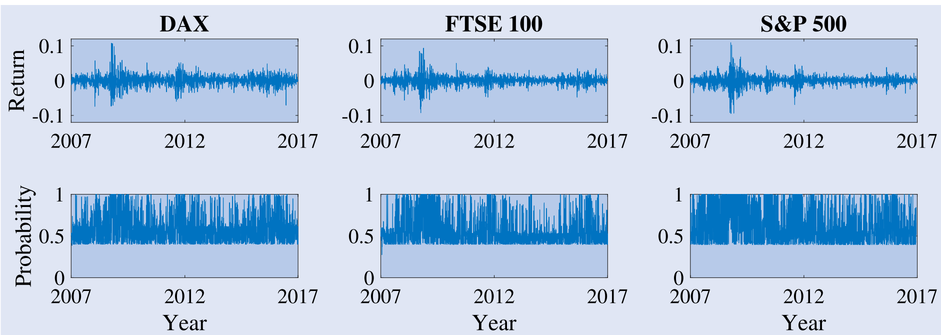

Figure 6. Time series of daily DAX, FTSE 100 and S&P 500 returns and corresponding estimated probabilities, i.e. weights $\widehat{w}_{2t} = \widetilde{w}_{2t}^{(m)}$, that the returns follow a Laplace distribution from 2 January 2007 to 31 December 2016.

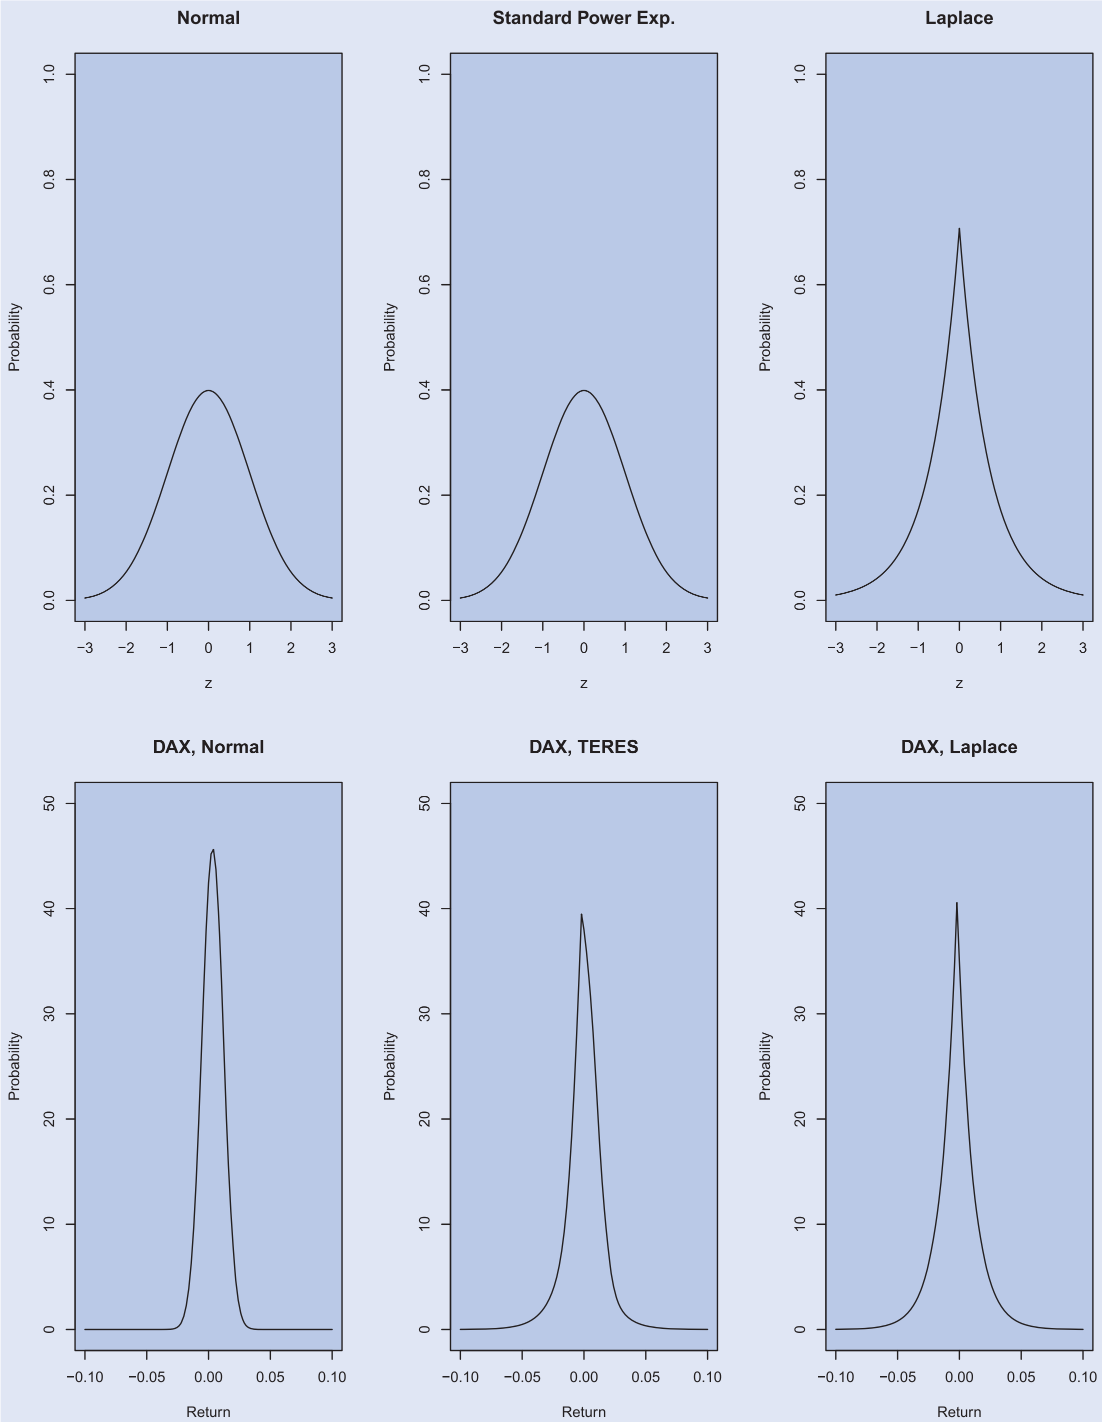

Figure 7. Theoretical densities (upper panel): normal, standard power exponential with power parameter 1.38 and Laplace. Estimated densities in risk management of daily DAX time series between 2007 and 2016 (lower panel): normal, TERES with $\widehat{\delta} = 0.60$ and Laplace.

power parameter equal to 1.38, which was estimated for the DAX series. Other index series that are modeled by TERES exhibit similar characteristics.

Our expectile-based normal–Laplace mixture model offers an economically sound risk quantification while 'moving' from the normal to the Laplacian case. Most prominently, it gives us an opportunity to evaluate the effect that 'shortening' and 'lengthening' of the distribution tail exercises on the displayed risk measures. Consider correspondingly a simulation study, with four cases: (a) TERES estimated expected shortfall for DAX time series for a selected mixture parameter $\delta \in \{1.00, 0.90, \ldots, 0.00\}$, (b) TERES calculated expected shortfall for a mixture between the standard normal and standard Laplace (zero mean, unit variance) distributions given $\delta \in \{1.00, 0.90, \ldots, 0.00\}$, (c) expected shortfall based on the standard power exponential distribution with power parameter $c \in \{1.00, 1.10, \ldots, 2.00\}$ and (d) TERES related power exponential distribution with matching rule-of-thumb power parameter $\widetilde{c}_{\delta}$. Based on the estimated expected shortfall using TERES in case (b), we report on the power parameter values that would lead to economically equal expected shortfall estimates. All simulation cases, based on risk level

Table 3. Estimated parameters of the normal–Laplace mixture for the selected index returns from 2007 to 2016.

	2007–2016			2007–2011			2012–2016		
	DAX 30	FTSE 100	S&P 500	DAX 30	FTSE 100	S&P 500	DAX 30	FTSE 100	S&P 500
$\widehat{\mu}_1$	0.00337	0.00167	0.00228	0.00237	0.00235	0.00337	0.00444	0.00149	0.00380
$\widehat{\sigma}_1$	0.00872	0.00702	0.00490	0.00970	0.00868	0.00436	0.00870	0.00633	0.00564
$\widehat{\mu}_2$	−0.00191	−0.00119	−0.00102	−0.00202	−0.00218	−0.00125	−0.00290	−0.00042	−0.00091
$\widehat{\sigma}_2$	0.01730	0.01531	0.01543	0.02073	0.01899	0.01877	0.01389	0.01035	0.00862
$\widehat{\delta}$	0.59735	0.56652	0.63536	0.55751	0.54028	0.75016	0.53635	0.67878	0.70871

Table 4. Estimated value at risk $\hat{q}_\tau$ and expected shortfall $\hat{s}_{e_{w_\tau}}$ of the normal-Laplace mixture for the selected index returns from 2007 to 2016. For convinience the measures are reported as positive numbers.

	2007–2016			2007–2011			2012–2016		
	DAX 30	FTSE 100	S&P 500	DAX 30	FTSE 100	S&P 500	DAX 30	FTSE 100	S&P 500
$\hat{q}_{0.01}$	0.044	0.035	0.041	0.052	0.049	0.051	0.033	0.025	0.023
$\hat{q}_{0.05}$	0.024	0.020	0.021	0.026	0.025	0.027	0.020	0.015	0.014
$\hat{s}_{e_{w_{0.01}}}$	0.054	0.057	0.044	0.063	0.052	0.055	0.053	0.038	0.027
$\hat{s}_{e_{w_{0.05}}}$	0.036	0.031	0.033	0.044	0.039	0.043	0.028	0.022	0.018

$\tau = 0.01$, are evaluated and summarised in table 5 and graphically illustrated in figure 8, and can be applied in TERES-based risk management practice.

The simulation experiments display several attractive properties of the proposed TERES framework. First, a larger value of the mixture parameter δ indeed 'lengthens' the distribution tail. Both the expected shortfall and expectile figures become, as expected, quite distant as mixture parameter δ indicates the presence of heavy-tails. Second, the mass probability indicator, $\frac{w_\tau}{\tau}$ clearly contributes to the tail behavior of the resulting risk measures. Third, while employing our TERES framework in a simulation study, we can easily suggest the economically sound parameter constellations of a benchmark approach. Here we show how the standard power exponential distribution modeling can be adjusted to attain the same risk level as proposed by TERES while 'moving' in an economically interpretable way from the normal to the heavy-tailed Laplacian case.

Most notably, related to the TERES equity risk estimation in practice, a higher change and thereby a large sensitivity was reported for the value at risk figures relative to expected shortfall estimates especially at the 1% across all indices; see table 4. This supports the current findings that value at risk indeed leads to increasing margin requirements in practice – here due to the change in the risk level – as compared to expected shortfall estimates. Expected shortfall estimates are more robust in the context of model (distribution) misspecification, which may be amplified in a quantile-based approach. In our framework, the potential risk is controllable because of the presented stability property. Therefore, by using TERES, decision makers may be able to identify a potential shift in the tail structure.

4.2. High-frequency finance

While analyzing the high-frequency data of 16 selected NASDAQ companies, one observes that the estimated mixture parameter changes most notably with the trading day selected; see figure 9. One observes a relatively large estimated mixture parameter values during the upward movement of the S&P

Table 5. Expected shortfall $s_{e_{w_\tau}}$ and the expectile e_{w_τ} in a simulation study. For convenience all measures are reported as positive numbers: (a) TERES-based normal–Laplace mixture for the selected DAX index returns series from 2007 to 2016 for specified δ, (b) TERES based standard normal–standard Laplace mixture for given δ, (c) Standard power exponential distribution with power parameter c and (d) TERES measure related rule-of-thumb power parameter $\widetilde{c}_\delta$.

(a) TERES, DAX			(b) TERES, Standard			(c) Power exp.			(d) TERES, Power		
δ	$s_{e_{w_\tau}}$	e_{w_τ}	δ	$s_{e_{w_\tau}}$	e_{w_τ}	c	$s_{e_{w_\tau}}$	e_{w_τ}	$\widetilde{c}_\delta$	$s_{e_{w_\tau}}$	e_{w_τ}
1.00	0.0623	0.0497	1.00	3.50	2.76	1.00	3.50	2.76	1.00	3.50	2.76
0.90	0.0610	0.0484	0.90	3.43	2.71	1.10	3.35	2.70	1.07	3.42	2.71
0.80	0.0594	0.0470	0.80	3.36	2.66	1.20	3.24	2.64	1.16	3.34	2.65
0.70	0.0576	0.0454	0.70	3.29	2.61	1.30	3.13	2.59	1.24	3.28	2.60
0.60	0.0558	0.0435	0.60	3.23	2.56	1.40	3.05	2.54	1.33	3.22	2.55
0.50	0.0538	0.0412	0.50	3.12	2.52	1.50	3.00	2.49	1.42	3.13	2.51
0.40	0.0510	0.0385	0.40	3.07	2.47	1.60	2.93	2.45	1.53	3.07	2.46
0.30	0.0474	0.0350	0.30	2.97	2.43	1.70	2.82	2.42	1.65	2.96	2.42
0.20	0.0424	0.0301	0.20	2.89	2.39	1.80	2.80	2.38	1.77	2.88	2.38
0.10	0.0343	0.0230	0.10	2.76	2.36	1.90	2.74	2.35	1.89	2.77	2.35
0.00	0.0200	0.0169	0.00	2.70	2.32	2.00	2.70	2.32	2.00	2.70	2.32

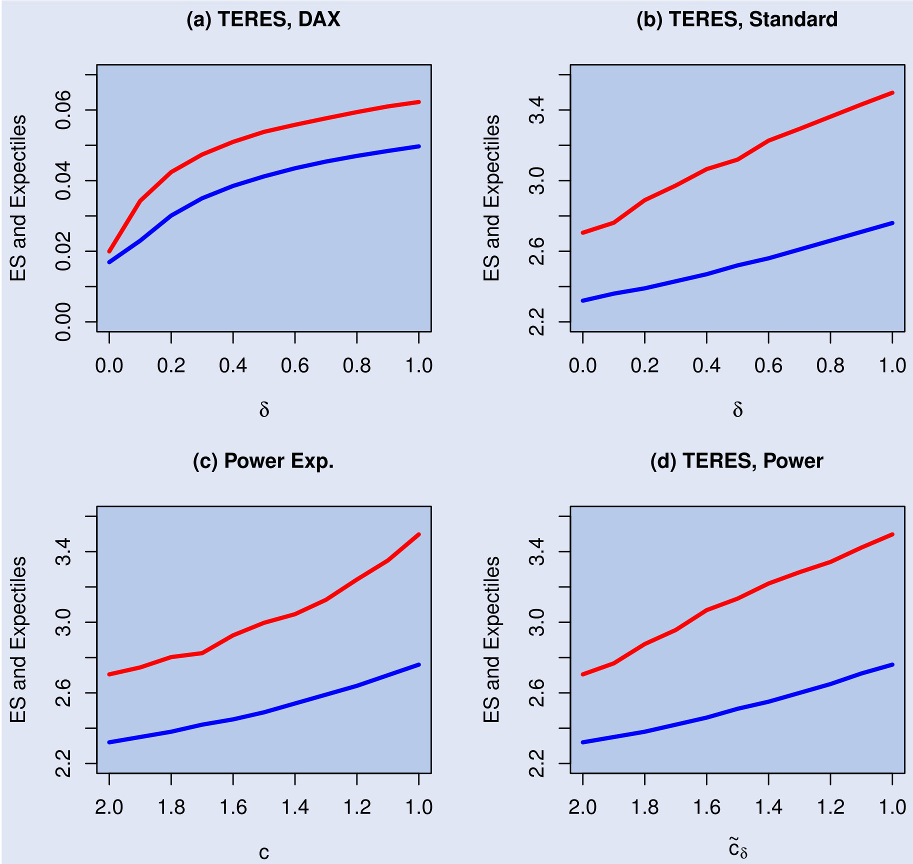

Figure 8. Expected shortfall $s_{e_{w_\tau}}$ (red) and the expectile e_{w_τ} (blue) in a simulation study. For convenience all measures are reported as positive numbers: (a) TERES-based normal–Laplace mixture for the selected DAX index returns series from 2007 to 2016, (b) TERES-based standard normal–standard Laplace mixture, (c) standard power exponential distribution and (d) TERES-based power exponential selection.

500 index on 30 June 2016, which was induced largely by the demand factor; see Mihoci (2017).

The prices of the analyzed stocks reacted positively on 30 June 2016, and the risk exposure, as measured by the expected shortfall using our TERES framework and the value at risk, see figure 10, indicates a correspondingly smaller risk. During relatively 'turbulent' trading days, a higher risk is implied by the model.

From the sector perspective, we see that there are several different risk-scheme transformations between the companies. Related to the mixture parameter, the consumer services sector exhibits relatively narrow changes between the analyzed days, whereas the largest deviations are present in the technological sector. While looking at the risk measures, again the consumer services market shows relatively unchanged (risk) readings (except AMZN), followed by the health care and the technological sector.

Related to tail sensitivity, a striking shift of the mixture parameter and risk quantities in the technological sector draws our attention. Taking FB and CSCO time series as examples, one observes a visible soar on the mixture parameter on 30 June 2016. An increasing expectile-based expected shortfall reflects a rising tail risk through a detected heavier tail structure. Quantile-based measures, however, remain constant or even move in the reverse direction. Less tail sensitivity by the quantile-based measures can also be evident in other sectors; see the risk assessment for AMZN. As evident,

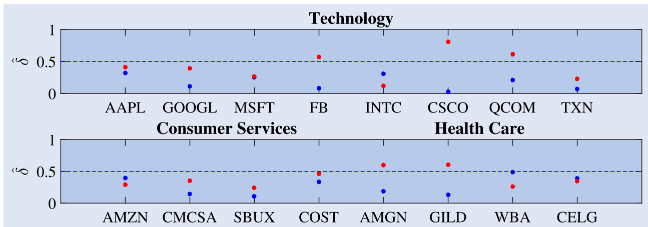

Figure 9. Estimated mixture parameter $\widehat{\tilde{\delta}}$ for selected NASDAQ companies on 27 (blue) and 30 (red) June 2016.

TERES-based risk management successfully captures the risk structure dynamics.

Referring to the intra-day variation of the time series, one expects a higher risk during the first day and relatively lower risk afterwards; see figure 2. An employment of TERES-based expected shortfall in the case of AMZN reconciles our expectation with the evidence. Not surprisingly, the quantile-based measures seemingly remain unchanged and are more vulnerable to the presence of outliers, especially during the market opening. This relatively noisy time period makes TERES a preferable choice in quantitative practice.

4.3. *Portfolio allocation*

The tail optimized TEDAS return series that was introduced in section 2 consist of 41 and 73 observations, respectively.

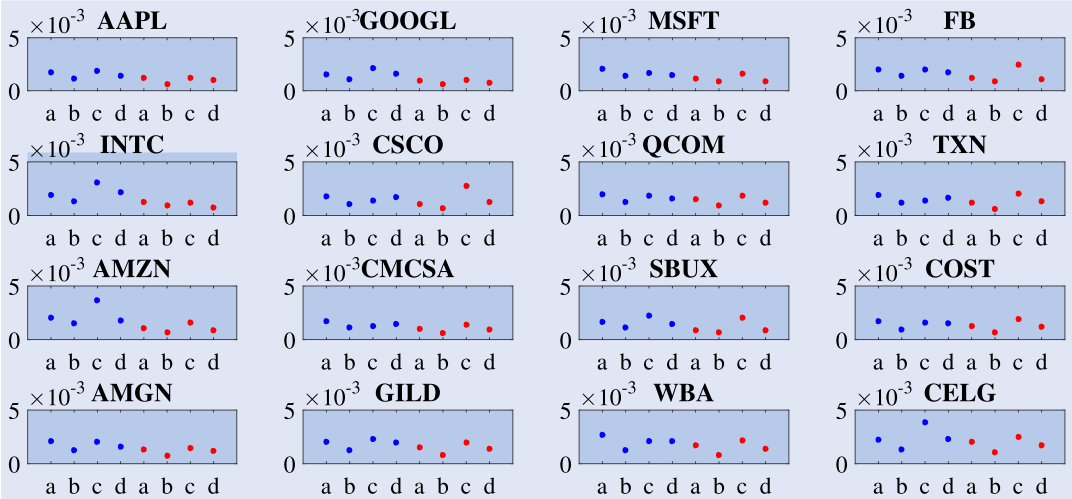

Figure 10. Estimated value at risk $\widehat{q}_\tau$ and expected shortfall $\widehat{s}_{e_{w_\tau}}$ of the normal–Laplace mixture for selected NASDAQ companies on 27 (blue) and 30 (red) June 2016. The four displayed cases include: (a) $\widehat{q}_{0.01}$, (b) $\widehat{q}_{0.05}$, (c) $\widehat{s}_{e_{w_{0.01}}}$ and (d) $\widehat{s}_{e_{w_{0.05}}}$. For convenience, the measures are reported as positive numbers.

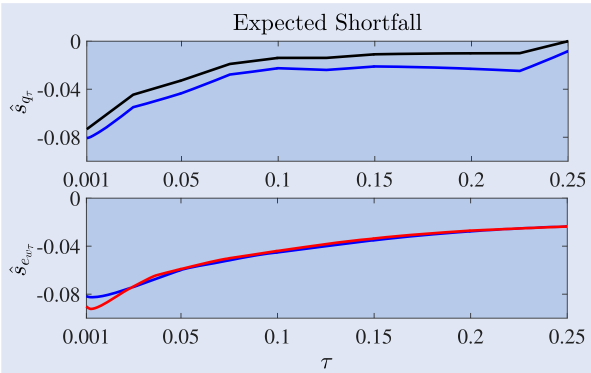

Figure 11. Estimated expected shorfall for the TEDAS (DAX) portfolio across different risk levels. Top: Empirical quantiles (black) and the resulting normal scenario expected shortfall (blue). Bottom: TERES-based expected shortfall with normal scenario ($\delta = 0$, blue) and high-tail risk scenario ($\delta = 1$, red).

This sparse availability of data makes it difficult to estimate the tail risk of the selected portfolio. By using TERES, we are able to overcome these challenges for small sample selections; see the risk estimates for the TEDAS portfolio based on DAX constituents and 41 observations, figure 11. TERES estimates are smoother, they induce monotonicity in the expected shortfall, and they deliver more suitable results; as shown in figure 11.

The normal and 'tail-heavy' Laplace contamination scenarios show a considerably higher degree of smoothness with respect to variation in the risk level τ. A comparison of the quantile and TERES-based expected shortfall realizations under a normal scenario reveals that our approach better captures the risk. Note that the empirical quantile employs linear interpolation, whereas the evidently exponential tail structure that is present here is more appropriately captured by our framework.

With inherited properties from expectiles, TERES naturally forms a smooth curve and ensures monotonicity. Its behavior is relatively stable and smooth, even with the sparse availability of data. As previously, a low variation over the scenarios is observed at the worldwide mutual funds market (the results available upon request), which becomes less pronounced as one considers risk levels τ closer to the distributional center. Due to these appealing properties, we recommend TERES for risk assessment practice.

5. Conclusion

The proposed TERES framework utilizes expectiles to estimate expected shortfall. Motivated by quantitative finance practice, we employ a normal–Laplace mixture in modeling expected shortfall. The key advantages of the approach include a relatively pronounced smoothness of the results, more robust evidence compared to quantile-based figures and advances in tail risk assessment of data using small samples.

Our empirical section presented successful TERES applications over several asset classes. First, during market distress periods, in normal–Laplace mixture modeling a relatively higher probability is given to the Laplace distribution. Interestingly, expected shortfall exhibits relatively robust results compared to value at risk when changing the underlying risk level. Second, while modeling high-frequency data, a smaller risk occurs during a trading day with an increasing price trend. Therefore, we additionally account for tail risk sensitivity for selected blue chip companies. Finally, TERES is considered advantageous in the case of monthly and weekly portfolio allocation. This makes TERES a desirable choice for risk management of selected asset classes on a monthly, weekly, daily and intra-day basis. Future research regarding TERES can consider the time-varying nature of the underlying distribution while capturing tail events and associated risk measures in time.

Acknowledgements

The authors are thankful for the research assistance provided by Philipp Gschöpf and to the reviewers for their valuable comments and suggestions. Humbold-Universität zu Berlin, Czech Science Foundation (19-28231X) and the Taiwan YuShan Scholarship, are gratefully acknowledged.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

Financial support from the Deutsche Forschungsgemeinschaft via Collaborative Research Center (CRC) 649 "Economic Risk" and International Research Training Group (IRTG) [SFB 649, IRTG 1792] "High Dimensional Non Stationary Time Series".

References

Acerbi, C., Coherent measures of risk in everyday market practice. *Quant. Finance*, 2007, **7**, 359–364.

Acerbi, C., Nordio, C. and Sirtori, C., Expected shortfall as a tool for financial risk management. arXiv preprint cond-mat/0102304, 2001.

Acerbi, C. and Tasche, D., On the coherence of expected shortfall. *J. Bank. Financ.*, 2002, **26**, 1487–1503.

Artzner, P., Delbaen, F., Eber, J.M. and Heath, D., Coherent measures of risk. http://www.math.ethz.ch/delbaen/ftp/preprints/CoherentMF.pdf, 1998.

Basel Committee on Banking Supervision, Fundamental review of the trading book: A revised market risk framework. Consultative Document, 2013.

Bassett, G.W., Koenker, R. and Kordas, G., Pessimistic portfolio allocation and Choquet expected utility. *J. Financ. Econom.*, 2004, **2**, 477–492.

Bollerslev, T., Generalized autoregressive conditional heteroskedasticity. *J. Econom.*, 1986, **31**, 307–327.

Breckling, J. and Chambers, R., M-quantiles. *Biometrica*, 1988, **75**, 761–771.

Cizek, P., Härdle, W.K. and Weron, R., *Statistical Tools for Finance and Insurance*, 2nd ed., 2011 (Springer Verlag: Heidelberg).

Cont, R., Deguest, R. and Scandolo, G., Robustness and sensitivity analysis of risk measurement procedures. *Quant. Finance*, 2010, **10**, 593–606.

Delbaen, F., Coherent risk measures on general probability spaces. *Adv. Finance Stoch.*, 2002, 1–37.

Dempster, A., Laird, N. and Rubin, D., Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Seri. B (Methodological)*, 1977, **39**, 1–38.

Engle, R.F., Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 1982, **50**, 987–1008.

Fama, E.F., The behavior of stock-market prices. *J. Business*, 1965, **38**, 34–105.

Franke, J., Härdle, W.K. and Hafner, C.M., *Statistics of Financial Markets: An Introduction*, 5th ed., 2019 (Springer Verlag: Berlin, Heidelberg).

Härdle, W.K., Lee, D.K.C., Nasekin, S., Ni, X. and Petukhina, A., Tail event driven ASset allocation: Evidence from equity and mutual funds' markets. SFB 649 DP 2015-045, 2015.

Hartley, H. and Rao, J., Maximum-likelihood estimation for the mixed analysis of variance model. *Biometrika*, 1967, **54**, 93–108.

Huber, P., Robust estimation of location parameter. *Annal. Math. Stat.*, 1964, **35**, 73–101.

Hull, J.C. and White, A.D., Value at risk when daily changes in market variables are not normally distributed. *J. Derivat.*, 1998, **5**, 9–19.

Jones, M., Expectiles and M-quantiles are quantiles. *Statist. Probab. Lett.*, 1994, **20**, 149–153.

Koenker, R., When are expectiles percentiles?. *Econ. Theor.*, 1993, **9**, 526–527.

Koenker, R., *Quantile Regression*, Econometric Society Monographs (Book 38), 2005 (Cambridge University Press).

Kusuoka, S., On law invariant coherent risk measures. *Adv. Math. Econom.*, 2001, 83–95.

Mandelbrot, B., The variation of certain speculative prices. *J. Business*, 1963, **36**, 394–419.

McNeil, A.J. and Frey, R., Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach. *J. Empir. Finance*, 2000, **7**, 271–300.

McNeil, A.J., Frey, R. and Embrechts, P., *Quantitative Risk Management, Concepts, Techniques and Tools*, revised ed., Princeton Series in Finance, 2015 (Princeton University Press: Princeton, NJ).

Mihoci, A., Modelling limit order book volume covariance structures. In *Advances in Statistical Methodologies and Their Applications to Real Problems*, edited by T. Hokimoto, 2017 (InTech: Rijeka).

Nolan, J.P., *Stable Distributions - Models for Heavy Tailed Data*, 2015 (Birkhauser: Boston). In progress, Chapter 1 online at academic2.american.edu/ ∼ jpnolan.

Otto-Sobotka, F., Salvati, N., Ranalli, M.G. and Kneib, T., Adaptive semiparametric M-quantile regression. *Econom. Stat.*, 2019, **11**, 116–129.

Overbeck, L., *Spectral Capital Allocation*, 2004 (RISK Books: London).

Rigby, R.A. and Stasinopoulos, D.M., Smooth centile curves for skew and kurtotic data modelled using the Box–Cox power exponential distribution. *Stat. Med.*, 2004, **23**, 3053–3076.

Schulze Waltrup, L., Sobotka, F., Kneib, T. and Kauermann, G., Expectile and quantile regression - David and Goliath?. *Stat. Modelling.*, 2015, **15**, 433–456.

Stahl, G., Zheng, J., Kiesel, R. and Rühlicke, R., Conceptualizing robustness in risk management. Available at SSRN 2065723, 2012.

Taylor, J., Estimating value at risk and expected shortfall using expectiles. *J. Financ. Econom.*, 2008, **6**, 231–252.

Wang, R., Regulatory arbitrage of risk measures. *Quant. Finance*, 2016, **16**, 337–347.

Ziegel, J., Coherence and elicitability. *Math. Finance*, 2016, **26**, 901–918.