Operations Research Letters 39 (2011) 163–171

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Conditional value-at-risk in portfolio optimization: Coherent but fragile

Andrew E.B. Lim [a], J. George Shanthikumar [b], Gah-Yi Vahn [a,*]

[a] Department of Industrial of Engineering & Operations Research, University of California, Berkeley, CA 94720, United States

[b] Krannert School of Management, Purdue University, West Lafayette, IN 47907, United States

ARTICLE INFO

Article history:
Received 29 August 2010
Accepted 21 February 2011
Available online 17 March 2011

Keywords:
Portfolio optimization
Conditional value-at-risk
Expected shortfall
Coherent measures of risk
Mean-CVaR optimization
Mean-variance optimization

ABSTRACT

We evaluate conditional value-at-risk (CVaR) as a risk measure in data-driven portfolio optimization. We show that portfolios obtained by solving mean-CVaR and global minimum CVaR problems are unreliable due to estimation errors of CVaR and/or the mean, which are magnified by optimization. This problem is exacerbated when the tail of the return distribution is made heavier. We conclude that CVaR, a *coherent* risk measure, is fragile in portfolio optimization due to estimation errors.

1. Introduction

Conditional value-at-risk (CVaR, also known as Expected Shortfall, or ES) has gained considerable attention in the financial risk management literature as a viable risk measure. CVaR at level β refers to the conditional expectation of losses in the top $100(1-\beta)\%$, and is anticipated to be a superior risk measure to value-at-risk (VaR), which, at level β, refers to the threshold level for losses in the top $100(1-\beta)\%$. At a time when the use of VaR is partly blamed for the 2007–2008 financial crisis, CVaR is more appealing than VaR because it takes into account the contribution from the very rare but very large losses. Formally, CVaR is a "coherent" risk measure, in that it satisfies [26,1] the four coherence axioms of [2]: translation invariance, subadditivity, positive homogeneity and monotonicity (whereas VaR violates subadditivity [2,10], i.e. diversification can result in greater risk). See [12] and the references therein for an overview of discussion on risk measures.

There have been many studies on CVaR once its coherence was established. In statistics/econometrics, there have been studies about CVaR estimation. Another line of work has been in incorporating CVaR as a risk measure in portfolio optimization, led by [27], which developed numerical methods for computing optimal portfolios with CVaR as the objective. These papers demonstrate CVaR portfolio optimization from a purely data-driven approach, i.e. the investor optimizes the portfolio based on empirical estimates of mean and CVaR.

If the underlying return distribution is multivariate normal and the investor knows its parameters, then the portfolio that minimizes CVaR with an expected return R is equivalent to the portfolio that minimizes variance (or VaR) with the same expected return R [27,8]. As a consequence, the frontiers of mean-variance and mean-CVaR portfolios coincide if plotted on the same scale. The authors of [27,8] also consider the case where the investor does not know the model and computes optimal portfolios purely based on historical data. Using real market data, they show that the *empirical frontiers* (see Section 2 for a precise definition) of mean-variance and mean-CVaR portfolios are very similar for the (different) assets and time periods under consideration. This merely indicates that the data used were approximately multivariate normal, and deems the use of CVaR over variance in a 'normal' market unnecessary (of the four papers, only De Giorgi correctly identifies this point). Nevertheless, the proponents of CVaR argue that its usefulness will be evident when the return distributions deviate from normality; in particular when they have fat left tails, as is often the case in 'crisis' periods [23]. However, to date, we are not aware of any studies on the use of CVaR as a risk measure in such a market have been done to validate this claim.

One important point that has been omitted by Artzner et al. and subsequent discussions on risk measures ([12,3,13], to name a few) is the role of estimation errors in the computation of a risk measure, and their effect on decision-making, such as portfolio optimization. CVaR may be coherent, but a large number

* Corresponding author.
E-mail address: gyvahn@ieor.berkeley.edu (G.-Y. Vahn).

0167-6377/$ – see front matter
doi:10.1016/j.orl.2011.03.004

of observations are needed to estimate it accurately because it is a tail statistic. However, in financial risk management, historical data older than 5 years may be irrelevant because of non-stationarity in the underlying return distribution. Thus from a practical perspective, data-driven portfolio optimization that involves estimated statistics is subject to estimation errors that may be very significant. Such observations in the context of mean-variance optimization have been made by [21,25,5,6]. We believe that an understanding of the impact of estimation errors in CVaR (as well as other tail risk measures) is important given its increasing popularity. The only works we are currently aware of that investigate issues concerning CVaR estimation are [7,17], where CVaR estimation is shown to be less "robust" than VaR in the sense of robustness defined by [19,15].

We note that this work is different in that we are investigating the estimation errors associated with CVaR in the context of an optimization problem—and we will see that the optimization procedure introduces extra "bias" to the statistical estimation errors already present.

The goal of this paper is thus to provide an objective analysis of the use of CVaR as a risk measure in data-driven portfolio optimization. Specifically, we set out to answer the following questions:

- How do estimation errors affect data-driven portfolio optimization that minimize CVaR as an objective?
- Is CVaR a reliable risk measure, in terms of estimation errors, for portfolio optimization in a heavy-tailed market?

To address these questions, we look at the mean-CVaR frontiers associated with the solution of empirical mean-CVaR problems constructed from data generated under different market models. We will first show that these empirical mean-CVaR frontiers vary wildly when both mean and CVaR are empirically estimated (call this problem EMEC, for *empirical mean-empirical CVaR*). Of course, such a variation may be due to the well-known problem of estimation errors of the mean [24]. To isolate the effect of the mean, we also consider (i) global minimum CVaR portfolio optimization (GMC) and (ii) mean-CVaR problem where the true mean is known (TMEC for *true mean-empirical CVaR*). For comparative purposes, we also analyze *empirical mean-empirical variance* (EMEV), *global minimum variance* (GMV) and *true mean-empirical variance* (TMEV) problems.

To see the effect of the tail behavior of return distributions, we do the analysis mentioned above for three different market scenarios: relative return distribution is multivariate normal ($\mathcal{M}1$), mixture of multivariate normal and negative exponential tail ($\mathcal{M}2$), or mixture of multivariate normal and one-sided power tail ($\mathcal{M}3$). Such mixture distributions represent a normal market that undergoes a shock with a small probability, with increasing heaviness in the tail.

The details of the evaluation methodology can be found in Section 2, of the optimization problems in Section 3 and simulation results and discussion in Section 4.

2. Evaluation methodology

We consider single-period portfolio optimization with n risky assets. We denote the excess returns of the assets by the random vector $\mathbf{X} = [X_1, \ldots, X_n]'$. To see how estimation errors affect EMEC portfolio optimization, we employ the following procedure:

1. Choose β (usually, 95% or 99%) and a model $\mathcal{M}$ for the distribution of the underlying assets. For example, $\mathcal{M}$ could be a multivariate normal distribution with parameters $(\boldsymbol{\mu}, \mathbf{V})$.
2. Simulate asset returns $\mathbf{D} = [\mathbf{x}_1, \ldots, \mathbf{x}_q]$ for a time period of size q under $\mathcal{M}$. This is the historical data the investor observes.
3. Fix a portfolio return level R. Compute the optimal solution of the EMEC(R; $\mathbf{D}$, β) problem (see Section 3 for details on the optimization). For the same data set $\mathbf{D}$, vary the range of R to compute a family of optimal portfolios. Note this family is random since the input data $\mathbf{D}$ is random.
4. For each portfolio in the family of optimal portfolios computed in Step 3, compute its expected (excess) return and CVaR under the true model $\mathcal{M}$, and plot the resulting mean-CVaR values. This generates a curve, the *empirical frontier*, representing the mean and CVaR of the portfolios computed in Step 3 under the true model $\mathcal{M}$. We do this because we are interested in the true performance of the EMEC portfolios. Note the empirical frontier is also random.
5. Repeat Steps 3–4 for EMEV(R; $\mathbf{D}$, β) portfolio optimization.
6. Repeat Steps 2–5 many times (50 in our study), each time with fresh input data $\mathbf{D}$. We can now compare the distribution of EMEC and EMEV empirical frontiers.

We employ a similar procedure for TMEC/TMEV and GMC/GMV portfolio optimization; the only difference is in the optimization problem we solve in Step 3. To see the effect of the tail of return distributions, we consider market models $\mathcal{M}$ with increasingly heavier one-sided tail; the exact characterizations of these markets are in Section 4.2. Next, we provide details of the optimization.

3. Data-driven portfolio optimization

3.1. Data-driven mean-variance portfolio optimization

The optimal data-driven mean-variance portfolio $\boldsymbol{\pi}^*_{MV}$ is given by solving the quadratic program:

$$\min_{\boldsymbol{\pi}} \boldsymbol{\pi}' V^{(q)} \boldsymbol{\pi} \tag{1a}$$

$$\text{s.t.} \sum_{i=1}^{n} \pi_i = 1 \tag{1b}$$

$$\boldsymbol{\pi}'\mathbf{g} = R, \tag{1c}$$

where $V^{(q)}$ is the sample covariance matrix computed from the observed data. For the EMEV problem, $\mathbf{g} = q^{-1}\sum_{i=1}^{q} \mathbf{x}_i$, i.e. the sample mean, and for the TMEV problem, $\mathbf{g} = E(\mathbf{X})$. For the GMV problem, we omit the constraint (1c).

3.2. Data-driven mean-CVaR portfolio optimization

The optimal data-driven mean-CVaR portfolio $\boldsymbol{\pi}^*_{\text{CVaR}}$ is given by solving:

$$\min_{\boldsymbol{\pi},\alpha} \alpha + \frac{1}{q(1-\beta)} \sum_{i=1}^{q} (-\boldsymbol{\pi}'\mathbf{x}_i - \alpha)^+ \tag{2a}$$

$$\text{s.t.} \sum_{i=1}^{n} \pi_i = 1 \tag{2b}$$

$$\boldsymbol{\pi}'\mathbf{g} = R, \tag{2c}$$

where $\mathbf{D} = [\mathbf{x}_1, \ldots, \mathbf{x}_q]$ are vectors of observed asset returns. We know from [27] that (2) can be transformed into a linear program. Again, for the EMEC problem, $\mathbf{g} = q^{-1}\sum_{i=1}^{q} \mathbf{x}_i$, and for the TMEC problem, $\mathbf{g} = E(\mathbf{X})$, and for the GMC problem, we omit the constraint (2c).

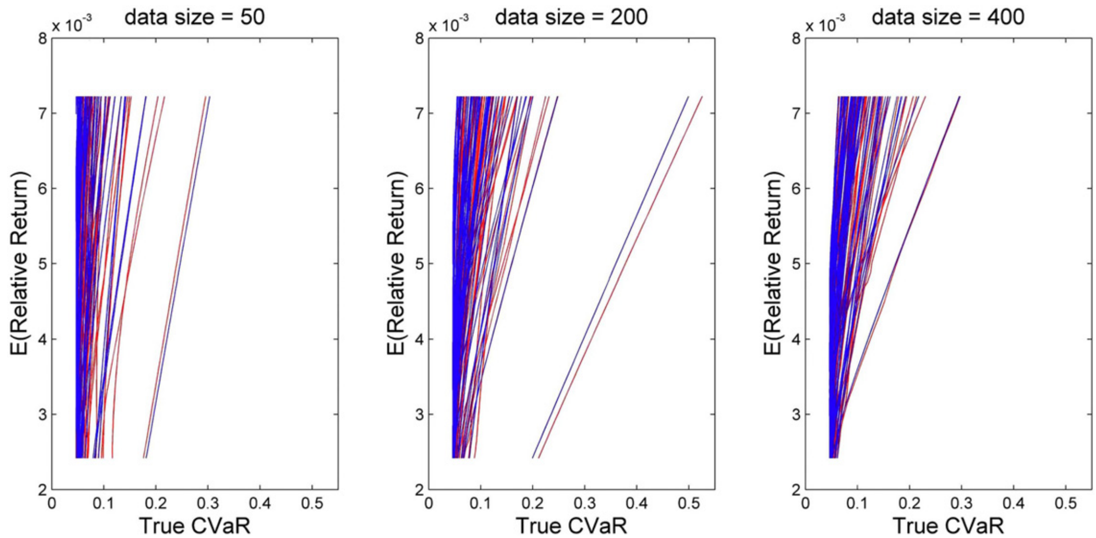

Fig. 1. The frontiers of empirical mean-empirical CVaR (red) and empirical mean-empirical variance (blue) portfolios under model. All scales are bps/mth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4. Results and discussion

4.1. *Returns $\sim$ multivariate normal*

4.1.1. *Model description*

We first consider the case where excess returns of $n = 5$ assets have a multivariate normal distribution:

$$\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{V}) \qquad (\mathcal{M}1)$$

where

$$\boldsymbol{\mu} = \begin{bmatrix}26.11, & 25.21, & 28.90, & 28.68, & 24.18\end{bmatrix} \times 10^{-4},$$

$$\boldsymbol{V} = \begin{bmatrix} 3.715, & 3.730, & 4.420, & 3.606, & 3.673 \\ 3.730, & 3.908, & 4.943, & 3.732, & 3.916 \\ 4.420, & 4.943, & 8.885, & 4.378, & 5.010 \\ 3.606, & 3.732, & 4.378, & 3.930, & 3.789 \\ 3.673, & 3.916, & 5.010, & 3.799, & 4.027 \end{bmatrix} \times 10^{-4}.$$

The vector $\boldsymbol{\mu}$ and the matrix $\boldsymbol{V}$ are the sample mean and covariance matrix of 299 monthly excess returns of 5 stock indices (NYA, GSPC, IXIC, DJI, OEX) from the period spanning August 3, 1984 to June 1, 2009. The histogram for the 10,000 sample returns of X_1 is shown in Fig. 2(a).

4.1.2. *Empirical mean-empirical CVaR (EMEC) problem*

As described in Section 2, we generate data using model $(\mathcal{M}1)$ and solve for EMEC and EMEV portfolios for a number of target expected returns R. We then simulate the returns for these portfolios under model $(\mathcal{M}1)$ to compute the true expected return and true CVaR, and generate the empirical frontiers in Fig. 1. Note we could have equally chosen true variance as the common risk scale. We also emphasize that the frontiers are not observed by the investor herself; they show the variability in the performance of empirical portfolio optimization under the true model.

The EMEC and EMEV frontiers both vary wildly; for example, with $q = 50$ ($\approx$4 years), the range of expected excess return of a portfolio with CVaR $=$ 1000 bps/mth per dollar invested (we will omit 'per dollar invested' hereafter) is greater than the range 25–72 bps/mth, which translates to a relative excess return ratio greater than $(1.0072^{12} - 1)/(1.0025^{12} - 1) \approx 300\%$ per year. The performance of EMEC and EMEV portfolios are very similar in that the positions and the spread of the frontiers are very similar. This is not very surprising, since, as previously mentioned, theoretical mean-variance and mean-CVaR problems yield equivalent solutions if the underlying asset returns have a multivariate normal distribution. Thus in a normal market, the EMEC problem is subject to large estimation errors of the mean, as is the EMEV problem. The same shortcomings apply to markets with heavier tails, as estimating the mean becomes more difficult as the underlying distribution becomes more irregular.

However, we cannot attribute the variation of the empirical frontiers solely to errors of the mean. What if we remove the expected return constraint, or assume the investor knows the true mean of asset returns? Removing the expected return constraint is a natural extension of EMEC and EMEV problems [20,9], and allows us to compare errors of CVaR and variance without errors of the mean. On the other hand, assuming the investor knows the true mean is an idealization where the investor has a good estimate of the mean obtained from alternatives to the sample average, e.g. based on CAPM [18], forecasts of exceptional returns [14], factor models [11], or incorporating investor knowledge via Black–Litterman [4]. In this regard, the TMEC model will allow us to evaluate whether the hard work put into estimating the mean can be destroyed by the errors associated with estimating CVaR. Hence, for the rest of the paper, we consider (i) GMC/GMV and (ii) TMEC/TMEV problems.

4.1.3. *Global minimum CVaR (GMC) problem*

We plot expected return vs. true CVaR of portfolios that solve the GMC problem (red) for $(\mathcal{M}1)$ in Fig. 4(a). For comparative purposes, GMV portfolios are also plotted (blue). In Table 1, we give the ranges for CVaR and expected return values corresponding to the solutions of GMC/GMV problems from our 50 simulations. We observe that GMV portfolios outperform GMC portfolios in that most blue stars are found on the left of the red stars (i.e. more accurate) and also are less spread out (i.e. more precise). The GMC portfolios are not precise at all; e.g. for our 50 simulations, when $q = 50$, the range of expected return is 9.15–43.2 bps/mth (481% per year relative difference) and for CVaR is 465–723 bps/mth.

What are the origins of the variation in the GMC empirical frontiers? In order to distinguish between the effects of inherent estimation errors and the optimization procedure, we plot in Fig. 3(a) *Perceived CVaR* and *Random Empirical CVaR* against *True CVaR*. By *Perceived CVaR* we mean the optimal objective from solving the GMC problem, i.e. the CVaR value perceived by the investor. By *Random Empirical CVaR* we mean the empirical

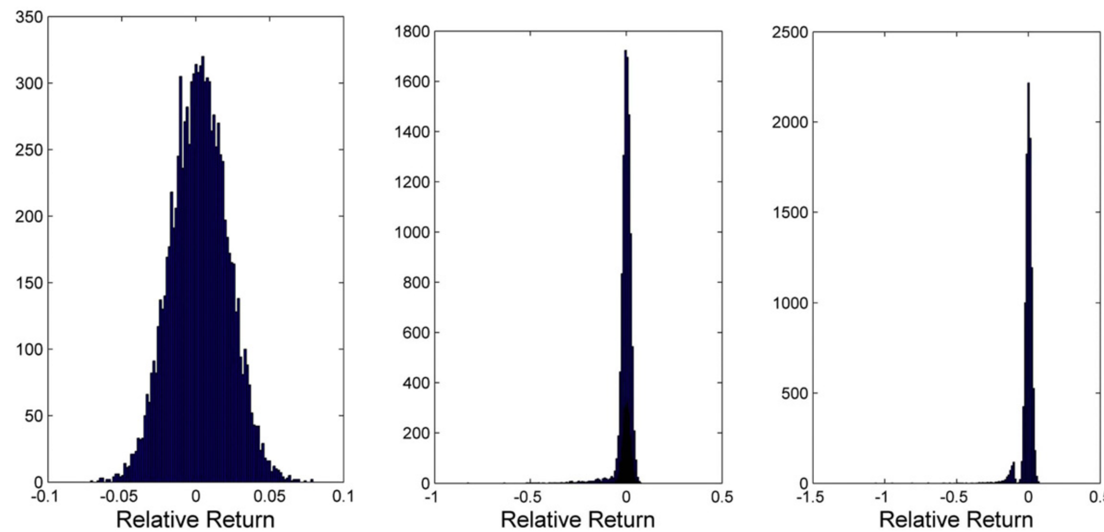

(a) Normal. (b) Laplace tail ($\lambda = 10$). (c) Power tail ($\lambda = 3.5$).

Fig. 2. Histogram for 10,000 samples of X_1 (bps/mth) under model (a) ($\mathcal{M}1$), (b) ($\mathcal{M}2$) and (c) ($\mathcal{M}3$).

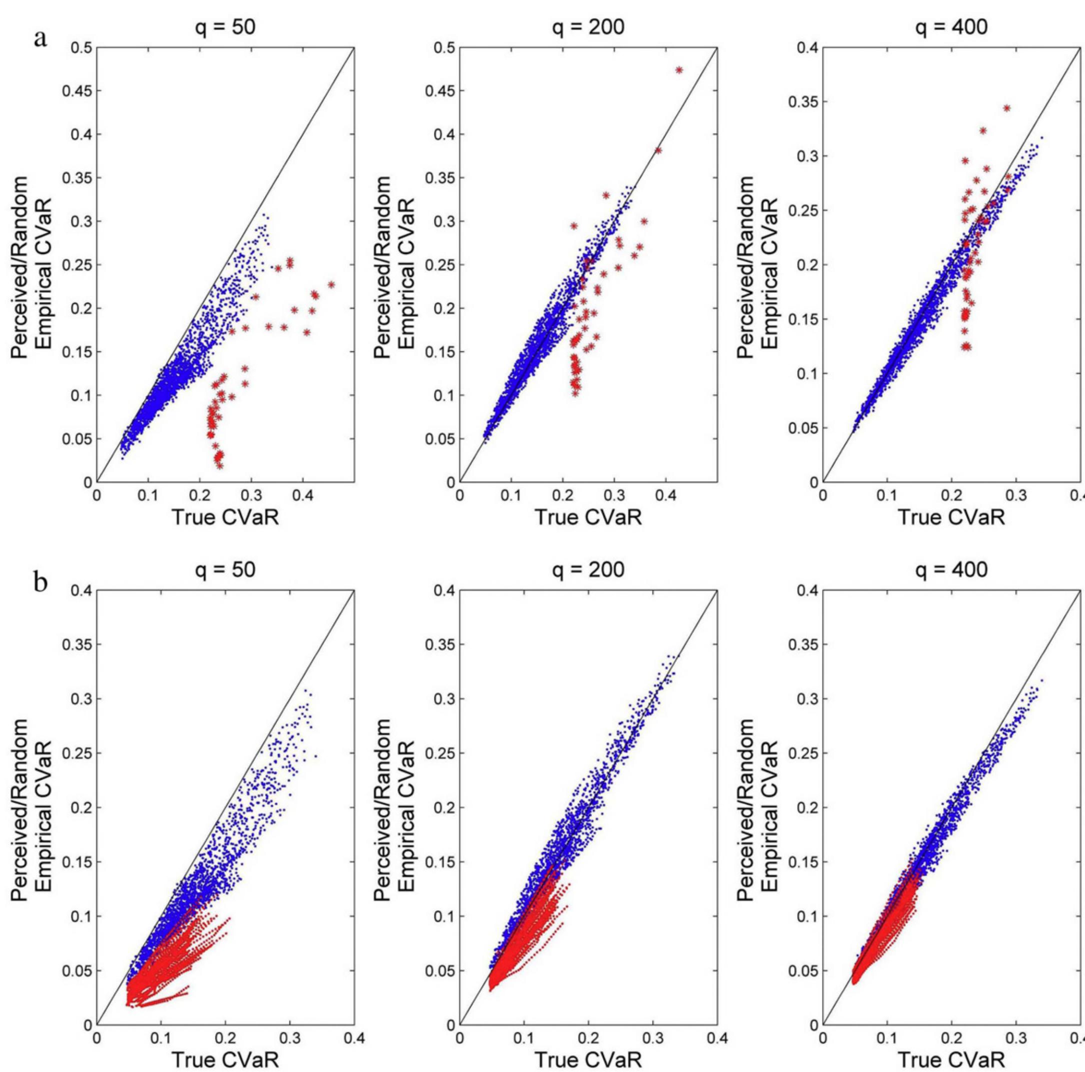

Fig. 3. Random Empirical CVaR vs. True CVaR (blue) and Perceived CVaR vs. True CVaR (red) for (a) global minimum CVaR problem under model ($\mathcal{M}1$) and (b) true mean-empirical CVaR problem. All scales are bps/mth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

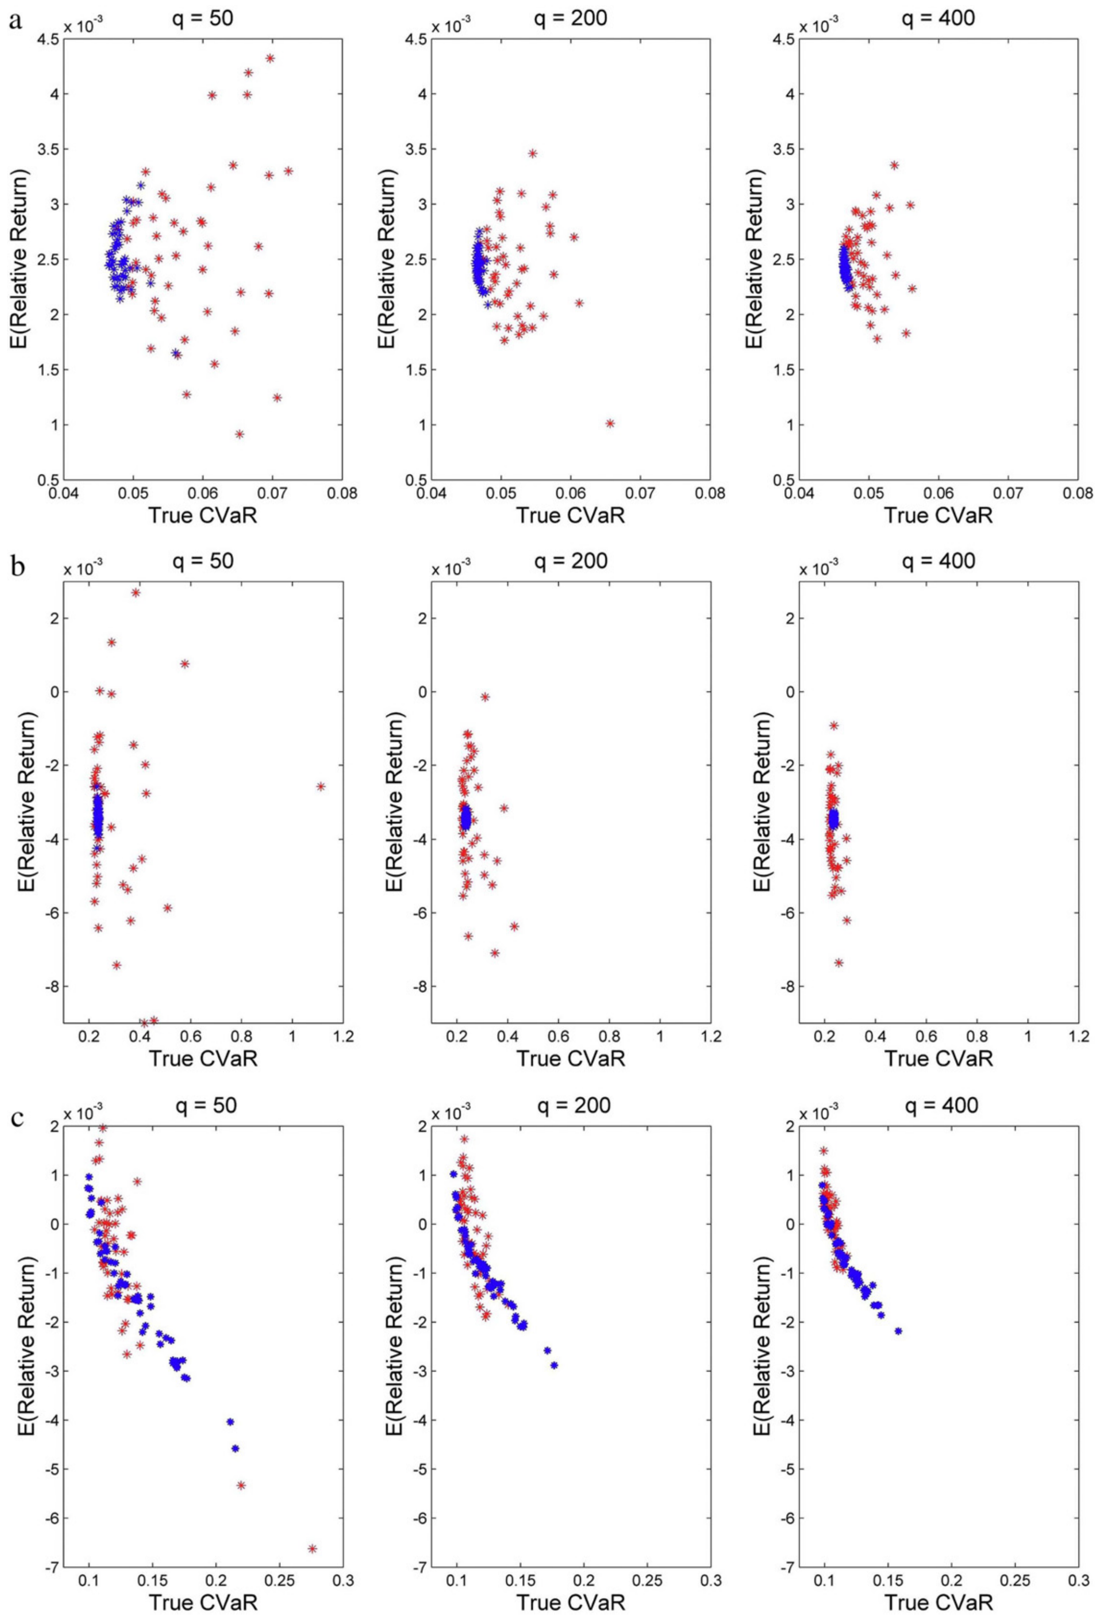

Fig. 4. The empirical frontiers of global minimum variance (blue) and global minimum CVaR (red) portfolios under models (a) ($\mathcal{M}1$) (b) ($\mathcal{M}2$) and (c) ($\mathcal{M}3$). All scales are bps/mth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

CVaR, defined by (2a), of portfolios that are randomly drawn (note the drawing procedure was not uniform) from the set of portfolios that satisfy the constraint (2b). *True CVaR* refers to the true CVaR value of a portfolio. Thus the scatter plot of Random Empirical CVaR against True CVaR (blue) gives an indication of the inherent estimation errors *before* optimization, and the scatter plot of Perceived CVaR against True CVaR (red) gives an indication of how the optimization affects the already present estimation

Table 1
Expected return and CVaR (true and perceived) ranges (bps/mth) for 50 GMC and GMV portfolios. In parenthesis is the size of the range.

	($\mathcal{M}1$)			($\mathcal{M}2$)			($\mathcal{M}3$)		
	GMV	GMC (true)	GMC (per)	GMV	GMC (true)	GMC (per)	GMV	GMC (true)	GMC (per)
CVaR									
$q = 50$	465 to 560 (96)	465 to 723 (257)	165 to 380 (257)	2298 to 2391 (93)	2204 to 11 120 (8916)	188 to 4706 (4518)	990 to 2151 (1161)	1043 to 2758 (1715)	255 to 687 (432)
$q = 200$	463 to 481 (18)	466 to 657 (191)	314 to 494 (180)	2332 to 2374 (42)	2202 to 4258 (2056)	1020 to 4738 (3718)	974 to 1764 (792)	1007 to 1403 (396)	480 to 837 (357)
$q = 400$	463 to 472 (9)	464 to 562 (98)	377 to 497 (121)	2339 to 2376 (37)	2194 to 2876 (682)	1240 to 3439 (2199)	980 to 1580 (600)	988 to 1177 (189)	555 to 946 (391)
Exp. return									
$q = 50$	16.5 to 31.7 (15.2)	9.15 to 43.2 (34.1)	−14.9 to 84.7 (99.6)	−42.4 to −25.8 (16.6)	−90.0 to 27.0 (117)	−318 to 426 (744)	−45.8 to 9.63 (55.4)	−66.3 to 19.6 (85.9)	−183 to 48.1 (232)
$q = 200$	20.9 to 27.5 (6.67)	10.1 to 34.6 (24.5)	−11.4 to 37.8 (49.2)	−36.9 to −31.6 (5.30)	−70.9 to −1.43 (69.5)	−34.4 to 98.5 (132.8)	−28.8 to 10.2 (39.0)	−18.9 to 17.3 (36.3)	−19.7 to −34.8 (54.6)
$q = 400$	22.4 to 25.9 (3.55)	17.8 to 33.5 (15.7)	8.31 to 41.5 (33.2)	−36.9 to −32.5 (4.40)	−73.6 to −9.21 (64.4)	−52.9 to 8.03 (61.0)	−21.8 to 7.94 (29.8)	−9.30 to 14.9 (24.2)	−20.4 to −5.65 (14.8)

Table 2
CVaR (true and perceived) ranges (bps/mth) for 50 mean-risk empirical frontiers at fixed expected return levels. In parenthesis is the size of the range.

Exp. return	EMEV/EMEC		TMEV/TMEC									
	($\mathcal{M}1$)		($\mathcal{M}1$)			($\mathcal{M}2$)			($\mathcal{M}3$)		
	TMEV	TMEC	TMEV	TMEC (true)	TMEC (per)	TMEV	TMEC (true)	TMEC (per)	TMEV	TMEC (true)	TMEC (per)
40 bps/mth											
$q = 50$	470–2210 (1740)	490–2150 (1660)	613–691 (78)	621–942 (321)	171–588 (417)	2721–2992 (271)	2736–10 610 (7874)	1053–4957 (3903)	1110–1242 (132)	1123–1778 (655)	469–916 (447)
$q = 200$	460–2970 (2510)	490–3120 (2630)	612–624 (12)	624–910 (286)	426–697 (271)	2720–2772 (52)	2739–4423 (1684)	1759–4998 (3239)	1109–1291 (182)	1113–1518 (405)	686–1159 (472)
$q = 400$	470–1200 (730)	470–1250 (780)	612–620 (8)	612–709 (97)	484–692 (208)	2720–2745 (25)	2728–3388 (660)	1904–3940 (2036)	1110–1179 (69)	1112–1299 (187)	846–1288 (443)
60 bps/mth											
$q = 50$	460–2720 (2260)	470–2650 (2180)	1039–1179 (140)	1056–1747 (691)	220–963 (742)	3139–3542 (403)	3161–10 570 (7409)	1282–5031 (3750)	1287–1530 (243)	1294–2072 (779)	565–1100 (536)
$q = 200$	480–4220 (3740)	520–4440 (3920)	1038–1063 (25)	1042–1294 (252)	692–1182 (490)	3140–3213 (73)	3155–4538 (1383)	2002–5126 (3124)	1286–1740 (455)	1288–1719 (431)	814–1362 (549)
$q = 400$	550–2280 (1730)	560–2300 (1740)	1038–1055 (17)	1045–1183 (138)	798–1186 (388)	3139–3175 (36)	3148–3555 (406)	2393–4122 (1729)	1286–1450 (164)	1287–1454 (167)	1008–1492 (484)

errors. We observe that empirical CVaR values of randomly drawn, unoptimized portfolios are slightly biased in the direction of underestimating the true CVaR. This is because the sample estimator from (2a) is biased in the direction of underestimation. However, this underestimation is aggravated by the optimization procedure, as can be seen by the position of the red dots—they are generally further below the black line (perfect estimation) than the blue dots. In Table 1, we highlight this phenomenon by listing the ranges of Perceived CVaR as well as the true CVaR for the GMC problem.

4.1.4. True mean-empirical CVaR (TMEC) problem

Fig. 5(a) shows TMEV empirical frontiers (blue) and TMEC($\beta = 0.99$) empirical frontiers (red) for different data sizes $q = 50, 200, 400$ (months). We also plot the theoretical mean-variance (equivalently, mean-CVaR) frontiers in green (left edge of the blue curves). In Table 2, we give the ranges for CVaR and expected return values corresponding to the solutions of TMEC/TMEV problems from our 50 simulations. In all three instances, the spreads of the TMEV empirical frontiers are significantly smaller than the TMEC empirical frontiers, and the TMEV empirical frontiers lie on the left-most side of the TMEC curves, closer to the theoretical frontier. Thus a portfolio manager who wishes to find the optimal TMEC portfolio should solve the TMEV problem to get a portfolio with higher accuracy and precision. For example, a manager with risk level CVaR $=$ 1000 bps/mth and $q = 50$ using TMEV optimization generates an average excess return of 56.6 bps/mth, but a manager with the same risk level using TMEC optimization generates 52.3 bps/mth—$\approx$ 8.5% per year higher excess return, on average. Furthermore, the TMEV manager is more reliable; e.g. for our 50 simulations and $q = 50$, the range of true CVaR is 613–691 bps/mth whereas for the TMEC manager it is 621–942 bps/mth for the same expected return of 40 bps/mth.

As with the GMC problem, the variation in the TMEC portfolios is due to inherent estimation errors of CVaR coupled with the effect of optimization. In Fig. 3(b) we plot the Perceived CVaR of TMEC portfolios in red, and empirical CVaR of random portfolios that satisfy (2b) in blue; notice the red dots are generally further below the black line (perfect estimation) than the blue dots. This is further verified by the CVaR (per) column in Table 2. This tells us that the TMEC investor can substantially underestimate the true CVaR value of her 'optimal' portfolio and be exposed to more risk than suggested by her perceived value.

Lastly, we comment on the difference between the performance of TMEC and TMEV empirical frontiers. Theoretically, the mean-CVaR and mean-variance empirical frontiers coincide; thus the discrepancy in the observed performance suggests that estimation errors of CVaR are more significant than of variance. This is not surprising since the mean vector and the covariance matrix are minimal sufficient statistics of a multivariate normal model (see, for example, [22]). Thus minimizing portfolio variance only contains errors of the covariance matrix, which are less significant than errors of the mean, whereas minimizing portfolio CVaR contains errors of both the mean vector and the covariance matrix. Of course, we could construct an artificial distribution that has

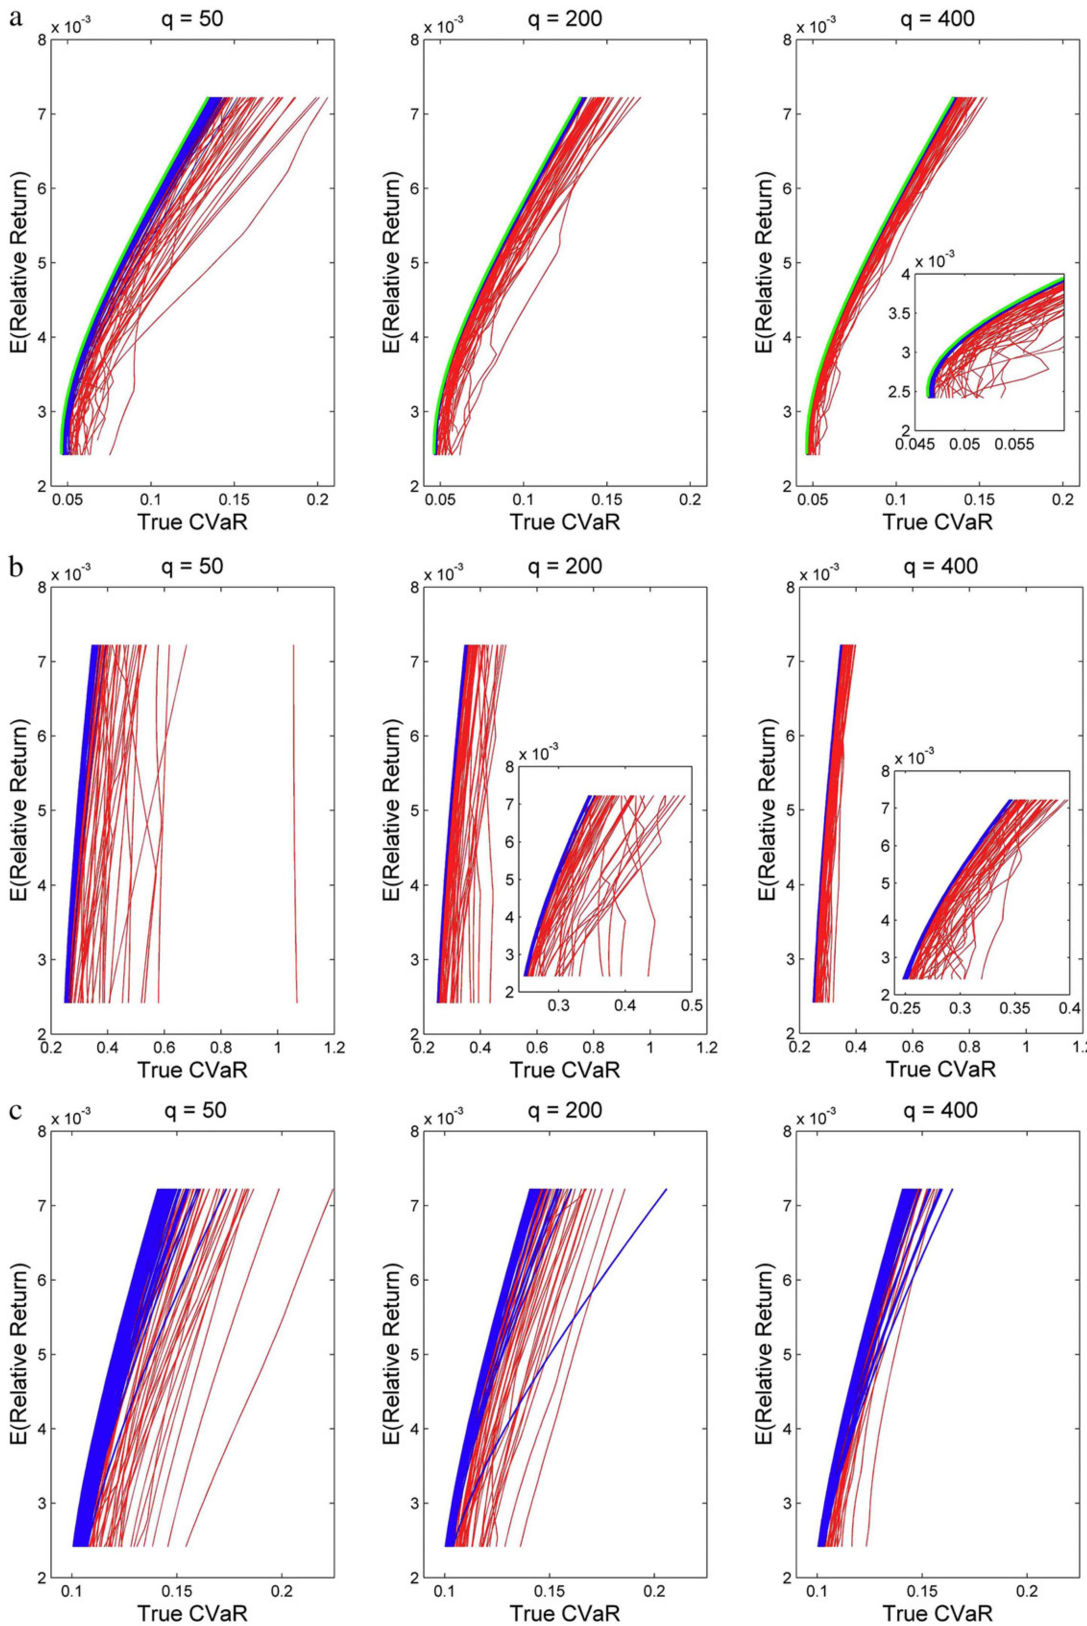

Fig. 5. The empirical frontiers of true mean-empirical variance (blue) and true mean-empirical CVaR (red) portfolios under models (a) ($\mathcal{M}1$) (b) ($\mathcal{M}2$) and (c) ($\mathcal{M}3$). In (a), we also plot the theoretical mean-variance (equivalently, mean-CVaR) frontier in green (found at the left-most edge of the blue curves). All scales are bps/mth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

CVaR (and/or the mean) being the minimal sufficient statistic rather than the mean and the covariance matrix and for which the solution to the TMEV problem is more reliable. However, as our purpose is to evaluate empirical CVaR optimization in the context of financial applications, such artificial distributions are not of interest in this study.

4.2. Returns $\sim$ multivariate normal + heavy loss tail

Recall that one of our objective is to evaluate CVaR portfolio optimization in markets with heavier tails. We now present analysis of GMC/GMV and TMEC/TMEV problems for two such markets.

4.2.1. Negative exponential tail

Let us consider returns being driven by a mixture of multivariate normal and negative exponential distributions, such that with a small probability, all assets suffer a perfectly correlated exponential-tail loss. Formally,

$$\mathbf{X} \sim (1 - I(\epsilon))N(\boldsymbol{\mu}, \mathbf{V}) + I(\epsilon)(Y\mathbf{e} + \mathbf{f}), \qquad (\mathcal{M}2)$$

where $I(\epsilon)$ is a Bernoulli random variable with parameter ϵ, $\mathbf{e}$ is a $n \times 1$ vector of ones, and $\mathbf{f} = [f_1, \dots, f_n]'$ is a $n \times 1$ vector of constants, and Y is a negative exponential random variable with density

$$P(Y = y) = \begin{cases} \lambda e^{\lambda y}, & \text{if } y \le 0 \\ 0 & \text{otherwise.} \end{cases}$$

In our simulations, we consider $\epsilon = 0.05$ (i.e. one shock every $\approx$ 1.7 years), $f_i = \mu_i - \sqrt{V_{ii}}$ and $\lambda = 10$. The histograms for 10,000 sample returns of X_1 is shown in (Fig. 2(b)).

4.2.2. One-sided power tail

Finally, we consider returns being driven by a mixture of multivariate normal and one-sided power distribution, such that with a small probability, all assets suffer a perfectly correlated power-tail loss. Formally,

$$\mathbf{X} \sim (1 - I(\epsilon))N(\boldsymbol{\mu}, \mathbf{V}) + I(\epsilon)Z(\gamma)\mathbf{f}, \qquad (\mathcal{M}3)$$

where $I(\epsilon)$ is a Bernoulli random variable with parameter ϵ, $\mathbf{f} = [f_1, \dots, f_n]'$ is a $n \times 1$ vector of constants, and $Z(\gamma)$ is a random variable defined for $\gamma \ge 1$ such that

$$P(Z(\gamma) = z) = \begin{cases} (\gamma - 1)(-z)^{-\gamma} & \text{if } z \le -1 \\ 0 & \text{otherwise.} \end{cases}$$

Note $\mathbf{X}$ under $(\mathcal{M}3)$ has finite variance for $\gamma > 3$. In our simulations, we consider $\epsilon = 0.05$, $f_i = \mu_i - 5\sqrt{V_{ii}}$, and $\gamma = 3.5$. The histogram for 10,000 sample returns of X_1 is shown in (Fig. 2(c)).

4.3. Discussion of results

4.3.1. Global minimum CVaR (GMC) problem

The expected return vs. true CVaR of GMC and GMV portfolios under models $(\mathcal{M}1)$, $(\mathcal{M}2)$ and $(\mathcal{M}3)$ are plotted in Fig. 4. In Table 1, we give the ranges for CVaR and expected return values corresponding to the solutions of GMC and GMV problems from our 50 simulations. In $(\mathcal{M}1)$ and $(\mathcal{M}2)$, the GMV portfolios perform better than GMC portfolios in that the blue stars are further to the left (i.e. more accurate) and have substantially smaller vertical and horizontal spreads (i.e. more precise) than the red stars. In $(\mathcal{M}3)$, the GMV portfolios are slightly more accurate and precise than GMC portfolios when $q = 50$, but the GMC portfolios converge faster with increasing data size. However, as financial data older than 5 years is rarely used in practice, $q = 50$ presents the most realistic scenario, and in this case, the GMV problem is clearly more reliable than the GMC problem across all models. In addition, when $q = 50$, the investor substantially underestimates the CVaR value—for our 50 simulations, True CVaR range is 2736–10610 bps/mth, whereas Perceived CVaR range is 1053–4957 bps/mth for the same expected return 40 bps/mth.

4.3.2. True mean-empirical CVaR (TMEC) problem

The empirical frontiers of TMEC and TMEV portfolios under models $(\mathcal{M}1)$, $(\mathcal{M}2)$ and $(\mathcal{M}3)$ are plotted in Fig. 5. In Table 2, we give the ranges for CVaR and expected return values corresponding to the solutions of TMEC/TMEV problems from our 50 simulations. For comparison, we also provide the ranges for EMEC/EMEV problems under $(\mathcal{M}1)$. Across all models, the TMEV empirical frontiers perform better than TMEV empirical frontiers in that the blue curves are further to the left and have smaller spreads than the red curves. There are differences between the models, however—the TMEV most outperform TMEC portfolios in $(\mathcal{M}2)$, whereas in $(\mathcal{M}3)$, the outperformance of TMEV empirical frontiers at $q = 50$ diminishes with increasing data. However, as previously mentioned, $q = 50$ presents the most realistic scenario, and in this case, the TMEV problem is clearly more reliable than the TMEC problem across all models. The investor remains too optimistic when $q = 50$—for our 50 simulations, True CVaR range is 1123–1778 bps/mth, whereas Perceived CVaR range is 461–916 bps/mth for the same expected return 40 bps/mth.

We can also see the effect of assuming the investor knows the true mean—from the $(\mathcal{M}1)$ columns in Table 2 and comparing Fig. 1 and (Fig. 5(a)), we see that the variation of both mean-variance and mean-CVaR empirical frontiers decrease substantially when the true mean is known. However, even if the investor estimates the mean exactly, estimation errors in CVaR can significantly affect the reliability of empirical frontiers, as is the case for $(\mathcal{M}2)$.

5. Concluding remarks

It is natural to ask whether these fragility problems could be alleviated if the investor had more knowledge about the asset returns and incorporated this into the model. This may be the case, but it would require alternative formulations of the mean-CVaR problem, since a general structure is not easily incorporated into Rockafellar–Uryasev's problem formulation. Furthermore, while incorporating structural information may improve the situation, substantial fragility problems may still remain. A simple illustration of this is the case where the investor correctly assumes that log-returns are i.i.d. Gaussian, but with unknown mean and covariance that need to be estimated from data. It is well known, however, that the Markowitz and mean-CVaR problems are equivalent for Gaussian models [27,8], and as the parametric Markowitz problem is susceptible to estimation errors if plug-in parameter estimates are used (see Fig. 1 and other studies of the Markowitz problem mentioned in the introduction), so is the mean-CVaR problem. Thus the impact of estimation errors can be substantial, even if investors are fortunate enough to have correct structural information (that the log-returns follow a Gaussian distribution).

The actual situation, of course, is much more difficult. It is well established that the tails of the return distributions are typically heavier than those of a Gaussian, and that the dependence structure between asset returns can be complex. This is important for CVaR because it is very sensitive to the tail of the return distribution. We know, however, that distinguishing between different tail behaviors is extremely difficult – for example, [16] show that even with 20 years of i.i.d. daily observations, one cannot distinguish between exponential and power-type tail – so the distributional information that is particularly relevant to CVaR is difficult to specify.

Several research directions are of interest: (i) investigating methods of reducing the impact of estimation errors in portfolio optimization problems when using risk measures with desirable theoretical properties such as CVaR, possibly by accounting for them explicitly in the model, (ii) developing risk measures that account for statistical errors and model uncertainty, and (iii) integrating these methods with structural information about the system being modeled.

Acknowledgments

This work is supported in part by an NSF *CAREER* Award CMMI-0348746 (Lim) and the NSF Awards CMMI-0500503 and CMMI-1031637 (Lim and Shanthikumar). Nevertheless, the opinions, findings, conclusions and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Andrew Lim acknowledges support from the NUS-Berkeley Risk Management Institute at the National University of Singapore and the Coleman Fung Chair in Financial Modelling. Gah-Yi Vahn acknowledges support from the Eleanor Sophia Wood Traveling Scholarship from The University of Sydney. The authors are grateful to Ankit Jain for comments and suggestions.

References

[1] C. Acerbi, D. Tasche, On the coherence of expected shortfall, Journal of Banking & Finance 26 (2001) 1487–1503.
[2] P. Artzner, F. Delbaen, J. Eber, D. Heath, Coherent measures of risk, Mathematical Finance 9 (3) (1999) 203–228.
[3] A. Ben-Tal, M. Teboulle, An old–new concept of convex risk measures: the optimized certainty equivalent, Mathematical Finance 17 (3) (2007) 449–476.
[4] F. Black, R. Litterman, Asset allocation: combining investor views with market equilibrium, The Journal of Fixed Income 1 (2) (1991) 7–18.
[5] M. Broadie, Computing efficient frontiers using estimated parameters, Annals of Operations Research 45 (1) (1993) 21–58.
[6] V. Chopra, W. Ziemba, The effect of errors in means, variances, and covariances on optimal portfolio choice, Journal of Portfolio Management 19 (2) (1993) 6–11.
[7] R. Cont, R. Deguest, G. Scandolo, Robustness and sensitivity analysis of risk measurement procedures, Quantitative Finance 10 (6) (2010) 593–606.
[8] E. De Giorgi, A note on portfolio selection under various risk measures, SSRN eLibrary.
[9] V. DeMiguel, L. Garlappi, F. Nogales, R. Uppal, A generalized approach to portfolio optimization: improving performance by constraining portfolio norms, Management Science 55 (5) (2009) 798.
[10] P. Embrechts, S. Resnick, G. Samorodnitsky, Extreme value theory as a risk management tool, North American Actuarial Journal 3 (1999) 30–41.
[11] E. Fama, K. French, Common risk factors in the returns on stocks and bonds, Journal of Financial Economics 33 (1) (1993) 3–56.
[12] H. Föllmer, A. Schied, Stochastic Finance: An Introduction in Discrete Time, De Gruyter, 2004.
[13] H. Föllmer, A. Schied, S. Weber, Robust preferences and robust portfolio choice, in: Bensoussan, Zhang (Eds.), Handbook of Numerical Analysis, in: Mathematical Modeling and Numerical Methods in Finance, vol. XV, 2009, pp. 29–89.
[14] R. Grinold, R. Kahn, Active Portfolio Management, McGraw-Hill, 2000.
[15] F. Hampel, E. Ronchetti, P. Rousseeuw, W. Stahel, Robust Statistics: The Approach Based on Influence Functions, John Wiley & Sons, New York, 1986.
[16] C. Heyde, S. Kou, On the controversy over tailweight of distributions, Operations Research Letters 32 (5) (2004) 399–408.
[17] C. Heyde, S. Kou, X. Peng, What is a good external risk measure: bridging the gaps between robustness, subadditivity, and insurance risk measures, Tech. Rep., Dept. of Industrial Engineering and Operations Research, Columbia University, 2007.
[18] C. Huang, R. Litzenberger, Foundations for Financial Economics, North Holland, 1988.
[19] P. Huber, E. Ronchetti, Robust Statistics, John Wiley & Sons Inc., 2009.
[20] R. Jagannathan, T. Ma, Risk reduction in large portfolios: why imposing the wrong constraints helps, The Journal of Finance 58 (4) (2003) 1651–1684.
[21] P. Jorion, International portfolio diversification with estimation risk, Journal of Business (1985) 259–278.
[22] E. Lehmann, G. Casella, Theory of Point Estimation, Springer Verlag, 1998.
[23] R. Litzenberger, D. Modest, Crisis and non-crisis risk in financial markets: a unified approach to risk management, SSRN eLibrary.
[24] R.C. Merton, On estimating the expected return on the market: an exploratory investigation, Journal of Financial Economics 8 (1980) 323–361.
[25] R. Michaud, The Markowitz optimization enigma: is optimized optimal? ICFA Continuing Education Series 1989 (4) (1989) 43–54.
[26] G. Pflug, Some remarks on the value-at-risk and the conditional value-at-risk, in: S. Uryasev (Ed.), Probabilistic Constrained Optimization: Methodology and Applications, Kluwer Academic Publishers, 2000, pp. 272–281.
[27] R. Rockafellar, S. Uryasev, Optimization of conditional value-at-risk, Journal of Risk 2 (2000) 21–41.