

Innovations in Insurance, Risk- and Asset Management

Kathrin Glau
Daniël Linders
Aleksey Min
Matthias Scherer
Lorenz Schneider
Rudi Zagst

editors



KPMG Center of Excellence
in Risk Management



World Scientific

Chapter 6

Data and Uncertainty in Extreme Risks: A Nonlinear Expectations Approach

Samuel N. Cohen

*Mathematical Institute, University of Oxford
Oxford, UK*

**samuel.cohen@maths.ox.ac.uk
<http://people.maths.ox.ac.uk/cohens/>*

Estimation of tail quantities, such as expected shortfall or Value at Risk, is a difficult problem. We show how the theory of nonlinear expectations, in particular the Data-robust expectation introduced in [1], can assist in the quantification of statistical uncertainty for these problems. However, when we are in a heavy-tailed context (in particular when our data are described by a Pareto distribution, as is common in much of extreme value theory), the theory of [1] is insufficient, and requires an additional regularization step which we introduce. By asking whether this regularization is possible, we obtain a qualitative requirement for reliable estimation of tail quantities and risk measures, in a Pareto setting.

Keywords: nonlinear expectation, extreme value theory, risk measure, uncertainty quantification.

1. Introduction

Statistics is, in part, the art of combining data and probabilistic models to predict the unknown. This necessarily results in uncertainty, not only as described by the estimated probabilistic model,¹ but also in the correctness of the model itself. In many situations, it is important to describe this uncertainty, and to incorporate it in our decision making.

One area where these issues are particularly prominent is when considering extreme outcomes. By their very nature, this involves looking at unlikely events, often beyond the range of available data. Justifying this extrapolation from observations is a key purpose of ‘extreme value theory’. In particular, the Fisher–Tippett–Gnedenko and Pickands–Balkema–de Haan theorems tell us that, assuming they exist, the possible distributions of

¹Since the work of Knight [2], uncertainty that is described by a probabilistic model is sometimes called ‘risk’, while lack of knowledge of the model is called ‘uncertainty’. Open Access chapter published by World Scientific Publishing Company and distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives (CC BY-NC 4.0) License.

renormalized extremes are in a small class, allowing asymptotic approximation methods to be used. Nevertheless, one must still estimate parameters of these distributions, and intuition suggests that there remains a significant level of uncertainty when working with these models.

The question of how to incorporate uncertainty in decision making has a long history. Bayesian approaches to statistics do this by placing a distribution over the unknown quantities, conditioning on observations, and then integrating over the possible unknowns. This effectively forces all uncertainty to be treated ‘linearly’, and one cannot describe an agent’s aversion to uncertainty separately to their aversion to risk (i.e. to the randomness which is included within the model). Conversely, classical methods in statistics allow for uncertainty in parameters (for example, through the use of confidence sets), but these are frequently used with relatively limited axiomatic decision-theoretic support.

One approach, building from an axiomatic approach to valuations of outcomes, is ‘nonlinear expectations’. These are mathematical generalizations of the classical expected value, allowing for risk and uncertainty to be represented in a rigorous and precise manner. Nonlinear expectations can be seen as generalizations of worst-case valuation, considered in a statistical context by Wald [3], but with an increased focus on making valuations of *multiple* random outcomes, rather than on minimizing a particular fixed risk. In finance, nonlinear expectations have played a particular role in the theory of convex risk measures, see for example Föllmer and Schied [4]. In a recent paper [1], a method of connecting these nonlinear expectations with statistical estimation was introduced.

In this paper, we shall look at some of the implications that this way of thinking has on extreme value theory. In particular, we will consider how quantifying our uncertainty leads to skepticism in extreme extrapolation, and how different approaches to measuring the riskiness of a decision require different quantities of data.

This paper proceeds as follows. Following an initial summary of the theory of nonlinear expectations and their connection with estimation, we will consider how to combine these expectations with measures of risk, corresponding in some cases to extrapolation into the tail. We will then look at methods of regularizing these nonlinear expectations, which are often needed when considering unbounded outcomes. Finally, we will apply this theory to a classic estimation problem with heavy tails, and draw conclusions for the estimation of a variety of measures of risk.

Remark 1.1. In their well known textbook on extreme value theory in finance and insurance, Embrechts, Klüppelberg and Mikosch [5] consider the challenges of extrapolation. In particular, they discuss their preference for looking at only moderately high quantiles of distributions, rather than estimating deeper into the tail. They state,

The reason why we are *very reluctant* to produce plots for high quantiles like 0.9999 or more, is that we feel that such estimates are to be treated *with extreme care*. [...] The statistical reliability of these estimates becomes, as we have seen, very difficult to judge in general. Though we can work out approximate confidence intervals for these estimators, such constructions strongly rely on mathematical assumptions which are unverifiable in practice.

— Embrechts, Klüppelberg and Mikosch, [5], p. 364
(emphasis in original)

In this paper, we give a partial attempt at a formal analysis supporting this reticence. We shall see that the theory of nonlinear expectations gives certain bounds (which depend on the amount of data at hand) beyond which the uncertainties from estimation dominate the computation of high quantiles (and similar quantities), hence statistical inference is problematic.

2. DR-expectations

Consider the following problem. Suppose we have a family $\{X\} \cup \{X_i\}_{i \in \mathbb{N}}$ of real-valued random variables on the canonical space $(\Omega, \mathcal{F}) = (\mathbb{R}^{\mathbb{N}}, \mathcal{B}(\mathbb{R}^{\mathbb{N}}))$, where $\mathcal{B}(\mathbb{R}^{\mathbb{N}})$ denotes the Borel cylinder σ -algebra. We observe $\mathbf{x}_N = \{X_n\}_{n=1}^N$, and seek to draw conclusions about the likely values of $\phi(X)$, where ϕ is some (known, Borel measurable) real function. Our aim, therefore, is to estimate the distribution of X , accounting for the uncertainty in our estimation.

Beginning classically, we first propose a family \mathcal{Q} of measures, assumed equivalent on $\sigma(\mathbf{x}_N)$ for each $N < \infty$, corresponding to possible joint distributions of $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$. For each of these distributions, we obtain the log-likelihood

$$\ell(\mathbf{x}_N; Q) = \log \left(\frac{dQ|_{\sigma(\mathbf{x}_N)}}{dQ_{\text{ref}}|_{\sigma(\mathbf{x}_N)}} \right) (\mathbf{x}_N)$$

(relative to some reference measure Q_{ref} on Ω). We can then define the divergence, or negative log-likelihood ratio,

$$\alpha_{Q|\mathbf{x}_N}(Q) = -\ell(\mathbf{x}_N; Q) + \sup_{Q \in \mathcal{Q}} \ell(\mathbf{x}_N; Q).$$

If we knew the ‘true’ measure Q , a natural estimate for $\phi(X)$ given our observations \mathbf{x}_N would be the conditional expectation $E_Q[\phi(X)|\mathbf{x}_N]$. Formally, as this is only defined for almost every observation \mathbf{x}_N , and we have a possibly uncountable family of measures \mathcal{Q} , this may cause technical difficulties. As we are working with countably generated Borel spaces,² simultaneously for every $Q \in \mathcal{Q}$, we can define the regular conditional Q -probability distribution of X given \mathbf{x}_N , that is, a probability kernel $\mathbf{Q} : \mathbb{R}^N \times \mathcal{B}(\mathbb{R}) \rightarrow [0, 1]$ such that

$$\mathbf{Q}(\mathbf{x}_N(\omega), A) = Q(X \in A | \mathbf{x}_N)(\omega),$$

and so obtain a ‘regular’ conditional expectation

$$E_Q[\phi(X)|\mathbf{x}_N] := \int_{\mathbb{R}} \phi(x) \mathbf{Q}(\mathbf{x}_N, dx).$$

We use this regular conditional expectation in what follows.

Remark 2.1. We shall focus our attention on the simpler case where $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$ are iid under each $Q \in \mathcal{Q}$. In this case, each measure in \mathcal{Q} can be described by the density f_Q of X under Q , and we obtain the simplifications $\ell(\mathbf{x}_N; Q) = \sum_{i=1}^N \log f_Q(X_i)$ and $E[\phi(X)|\mathbf{x}_N] = E[\phi(X)]$.

The key idea of [1] is to define the following operator.

Definition 2.1. For fixed constants $k > 0$, $\gamma \in [1, \infty]$, define the nonlinear expectation

$$\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma}(\phi(X)) = \sup_{Q \in \mathcal{Q}} \left\{ E[\phi(X)|\mathbf{x}_N] - \left(\frac{1}{k} \alpha_{Q|\mathbf{x}_N}(Q) \right)^\gamma \right\},$$

where we write $|x|^\infty = 0$ for $|x| \leq 1$, and ∞ otherwise. We call this the DR-expectation (with parameters k, γ). ‘DR’ is a deliberately ambiguous acronym, and refers to either ‘divergence robust’ or ‘data-driven robust’.

We can see that $\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma}$ is built from considering the expected values under a range of models Q , penalized by how well each model fits the observed data \mathbf{x}_N . This allows the DR-expectation to encode the uncertainty inherent in statistical estimation.

²See, for example, Cohen and Elliott [6], Section 2.6 or Shiryaev [7], Section II.7 for further details on the mathematics of this construction.

This nonlinear expectation has many elegant properties. From the perspective of the theory of nonlinear expectations, it is monotone, convex, constant equivariant and assigns all constants their values. One can obtain a risk measure, in the sense of Föllmer and Schied [4], Frittelli and Rosazza-Gianin [8] and Artzner, Delbaen, Eber and Heath [9], by specifying $\rho(\phi(X)) = \mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma}(-\phi(X))$, or by treating $\phi(X)$ as ‘losses’. The nonlinear expectation is coherent (in particular, it is positively homogeneous) if $\gamma = \infty$.

At the same time, unlike most risk measures in the literature, the DR-expectation is directly built from the observed data to represent statistical uncertainty. In order to perform classical statistical estimation, an agent needs to specify a class of models \mathcal{Q} to be considered (and in a Bayesian analysis, a prior distribution over models); the only remaining inputs to the problem are the observations \mathbf{x}_N . In the DR-expectation framework, one additionally needs the uncertainty aversion parameter k and the curvature parameter γ . Rather than separating statistical inference from the valuation of outcomes, the DR-expectation combines the two, as the likelihood is part of the definition of the expectation. Observations \mathbf{x}_N do not simply affect our expectation through conditioning (as in a Bayesian analysis), they also determine which models $Q \in \mathcal{Q}$ are ‘good’, in the sense of fitting our past observations well.

Remark 2.2. A key advantage of this approach is that we define the DR-expectation for every function ϕ . This allows the comparison of different ϕ to be consistently carried out, and the impact of statistical uncertainty on different ϕ may vary. While we have assumed above that X is real-valued, there is no issue if X is taken to be \mathbb{R}^d -valued (or more generally, Borel measurable and valued in some separable Banach space), which allows a great deal of flexibility in modeling.

In the setting where $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$ are iid, and under some regularity assumptions on \mathcal{Q} detailed in [1], one can show that the DR-expectation is a consistent estimator of the expected value of $\phi(X)$, that is,

$$\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma}(\phi(X)) \rightarrow_P E_P[\phi(X)] \quad \text{as } N \rightarrow \infty, \text{ for every } P \in \mathcal{Q},$$

for any value of $k > 0$ and $\gamma \in [1, \infty]$. It is also clear that the case $\gamma = \infty$ is closely related to the Neyman–Pearson lemma from hypothesis testing, as it considers models where the negative log-likelihood-ratio $\alpha_{\mathcal{Q}|\mathbf{x}_N}(Q)$ is sufficiently small.

In order to give precise asymptotic statements, the following definition will prove useful. Here P^* is the outer measure associated with P , to deal with any potential lack of measurability.

Definition 2.2. Consider sequences $f = \{f_n\}_{n \in \mathbb{N}}$ and $g = \{g_n\}_{n \in \mathbb{N}}$ of functions $\Omega \rightarrow \mathbb{R}$.

- (i) We write $f = O_P(g)$ whenever f_n/g_n is stochastically bounded, that is, $P^*(|f_n/g_n| > M) \rightarrow 0$ as $M \rightarrow \infty$ for each n ,
- (ii) We write $f = o_P(g)$ whenever $\lim_{n \rightarrow \infty} P^*(|f_n/g_n| > \epsilon) = 0$ for all $\epsilon > 0$.

Note that this depends on the choice of measure P .

Remark 2.3. In [1], a range of large-sample asymptotic results were obtained, for iid observations and X , under varying assumptions on the family \mathcal{Q} . Suppose that \mathcal{Q} is a subset of a ‘nice’ exponential family parametrized by $\theta \in \Theta$ (in particular, Assumption 3.1, as detailed below, holds) and the maximum likelihood estimator based on a sample of size N , denoted $\hat{\theta}_N$, is well defined. For bounded ϕ , the following large sample approximations were obtained:

$$\begin{aligned}\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,1}(\phi(X)) &= E_{\hat{\theta}_N}[\phi(X)] + \frac{k}{2N}V(\phi, \hat{\theta}_N) + O_P(N^{-3/2}), \\ \mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\infty}(\phi(X)) &= E_{\hat{\theta}_N}[\phi(X)] + \sqrt{\frac{2k}{N}}V(\phi, \hat{\theta}_N) + O_P(N^{-3/4}),\end{aligned}$$

for every $P \in \mathcal{Q}$, where $V(\phi, \hat{\theta}_N)$ is the ‘local variance’ at the MLE, defined by

$$V(\phi, \hat{\theta}) := \left(\frac{\partial E_{\theta}[\phi(X)]}{\partial \theta} \Big|_{\hat{\theta}} \right)^{\top} (\mathfrak{I}_{\hat{\theta}}^{-1}) \left(\frac{\partial E_{\theta}[\phi(X)]}{\partial \theta} \Big|_{\hat{\theta}} \right),$$

for $\mathfrak{I}_{\hat{\theta}_N}$ the observed information matrix (the Hessian of the negative log-likelihood) at $\hat{\theta}_N$.

Remark 2.4. The results in [1] focus on the case where $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$ are iid under every model in \mathcal{Q} . This is clearly restrictive from a modeling perspective. However, it is clear from the definition that the DR-expectation depends on the set \mathcal{Q} only through the conditional expectation and the log-likelihood, evaluated at the observed data \mathbf{x}_N . Therefore, provided our models are ‘close’ to an iid setting, in the sense that our expectations and likelihoods are similar to those from an iid model (possibly after some transformation of the space), the asymptotic behavior of the DR-expectation should not be significantly affected.

Remark 2.5. The DR-expectation is particularly designed to highlight statistical uncertainty, rather than an individual's risk aversion. In particular, the above results (particularly when combined with the classic analysis of White [10]) show that as the sample size $N \rightarrow \infty$, in the iid case,

$$\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma}(\phi(X)) \approx -\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma}(-\phi(X)) \approx E_P[\phi(X)],$$

where P is the measure in \mathcal{Q} which is closest to the empirical distribution of the data. In this way, an individual's preferences can be seen not to play any role in the DR-expectation when samples are large (that is, when statistical uncertainty is low).

2.1. Data-robust risk measures

To focus attention on extreme events, or to include an individual's risk aversion, we may wish to extend our focus from the DR-expectation. Classically, when we ignore statistical uncertainty (so the DR-expectation would be replaced by the expectation with respect to the known distribution), a standard approach would be to take X to denote losses and ϕ to be a utility function. In our setting, following this approach would lead to a version of expected utility theory where the expectation is replaced with the DR-expectation.

Further interesting cases can be obtained by directly combining the DR-expectation with a risk assessment depending directly on the law of the random outcome. As we shall see, many common risk assessments are of this type, for example the value at risk and expected shortfall, as well as classical statistics such as the expectation and variance. Writing \mathcal{M}_1 for the space of probability measures on \mathbb{R} , we consider a map $\mathcal{R} : \mathcal{M}_1 \rightarrow \mathbb{R}$ which represents the 'riskiness' of a gamble with specified distribution. We can then combine the risk aversion of \mathcal{R} (which treats the law of X as fixed) with the uncertainty aversion of the DR-expectation (which considers our uncertainty in this law). We formalize this construction in the following definition.

Definition 2.3. Write \mathcal{M}_1 for the space of probability measures on \mathbb{R} and $\mathcal{L}_Q(\xi|\mathbf{x})$ for the (regular conditional) law of a random variable ξ under the measure Q (given observations \mathbf{x}). Let $\mathcal{R} : \mathcal{M}_1 \rightarrow \mathbb{R}$ be a map with $\mathcal{R}(\mathcal{L}_Q(\xi|\mathbf{x}))$ representing the risk of our position ξ assuming $Q \in \mathcal{Q}$ is a 'true' model. We combine $\mathcal{E}_{Q|\mathbf{x}}^{k,\gamma}$ and \mathcal{R} through the definition

$$(\mathcal{E}_{Q|\mathbf{x}}^{k,\gamma} \circ \mathcal{R})(\xi) := \sup_{Q \in \mathcal{Q}} \left\{ \mathcal{R}(\mathcal{L}_Q(\xi|\mathbf{x})) - \left(\frac{1}{k} \alpha_{Q|\mathbf{x}}(Q) \right)^\gamma \right\}$$

Remark 2.6. For the sake of simplicity, we will identify this construction by adding a prefix ‘DR-’. For example, combining the DR-expectation and the expected shortfall, as in the example below, we obtain the ‘DR-expected shortfall’, similarly the ‘DR-Value at Risk’.

Example 2.1. Consider the risk assessment given by the (upper) expected shortfall at level ϵ (also called the conditional value at risk, tail value at risk or average value at risk, by various authors). For a model P such that ξ has a continuous distribution, this can be written

$$\mathcal{R}(\mathcal{L}_P(\xi)) = E_P[\xi | \xi \geq F_P^{-1}(1 - \epsilon)] = \text{ES}_\epsilon(\xi)$$

where F_P is the cdf of X under the measure P . This is a coherent risk measure (or nonlinear expectation), and can also be written (see Föllmer and Schied [4] or further discussion in McNeil, Frey and Embrechts [11])

$$\text{ES}_\epsilon(\xi) = \sup_{P' \sim P} \left\{ E_{P'}[\xi] - \tilde{\alpha}(P'; P) \right\}$$

where

$$\tilde{\alpha}(P'; P) = \begin{cases} 0 & \text{if } \|dP'/dP\|_\infty < \epsilon^{-1}, \\ \infty & \text{otherwise.} \end{cases}$$

In the special case where \mathcal{R} corresponds to a convex expectation (based on a reference measure), this construction can simplify in a natural way.

Proposition 2.1. *Suppose that, for each measure $P \in \mathcal{Q}$, our risk assessment \mathcal{R} has representation*

$$\mathcal{R}(\mathcal{L}_P(\xi|\mathbf{x})) = \sup_{P' \in \tilde{\mathcal{Q}}} \{E_{P'}[\xi|\mathbf{x}] - \tilde{\alpha}(P'; P)\},$$

where $\tilde{\mathcal{Q}}$ is a collection of probability measures on Ω with $\mathcal{Q} \subseteq \tilde{\mathcal{Q}}$ and $\tilde{\alpha} : \tilde{\mathcal{Q}} \times \tilde{\mathcal{Q}} \rightarrow \mathbb{R}$ is an arbitrary penalty function. The composition of $\mathcal{E}_{\mathcal{Q}|\mathbf{x}}^{k,\gamma}$ and \mathcal{R} is a nonlinear expectation, and has representation

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}}^{k,\gamma} \circ \mathcal{R})(\xi) = \sup_{P' \in \tilde{\mathcal{Q}}} \{E_{P'}[\xi|\mathbf{x}] - \alpha^*(P')\}$$

where α^* is the inf-sum (cf. the inf-convolution in Barrieu and El Karoui [12])

$$\alpha^*(P') = \inf_{Q \in \mathcal{Q}} \left\{ \left(\frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}}(Q) \right)^\gamma + \tilde{\alpha}(P'; Q) \right\}.$$

Proof. To obtain the representation, simply expand

$$\begin{aligned} (\mathcal{E}_{\mathcal{Q}|\mathbf{x}}^{k,\gamma} \circ \mathcal{R})(\xi) &= \sup_{Q \in \mathcal{Q}} \left\{ \mathcal{R}(\mathcal{L}_Q(\xi|\mathbf{x})) - \left(\frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}}(Q) \right)^\gamma \right\} \\ &= \sup_{Q \in \mathcal{Q}, P' \in \tilde{\mathcal{Q}}} \left\{ E_{P'}[\xi|\mathbf{x}] - \left(\frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}}(Q) \right)^\gamma - \tilde{\alpha}(P'; Q) \right\} \\ &= \sup_{P' \in \tilde{\mathcal{Q}}} \left\{ E_{P'}[\xi|\mathbf{x}] - \alpha^*(P') \right\} \end{aligned}$$

This defines a nonlinear expectation by duality, as in Föllmer and Schied [4]. \square

3. Regularization from data

The above approach, based on simply penalizing with the divergence, often fails to give bounded values when considering unbounded random variables. In some sense, this is because of the difficulty in using data to determine the probabilities of extreme outcomes. To motivate our discussion, we consider the following simple example.

Example 3.1. Consider the case where $\xi = \phi(X) = \beta X$, and $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$ are iid $N(\mu, \sigma^2)$ distributed random variables, with $\mu \in \mathbb{R}$ and $\sigma^2 > 0$ unknown. Writing $\bar{X} = N^{-1} \sum_{i=1}^N X_i$ and $\hat{\sigma}^2 = N^{-1} \sum_i (X_i - \bar{X})^2$ for the MLE estimates of μ and σ^2 , it is straightforward to calculate (see [1] for details)

$$\begin{aligned} \mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,1}(\beta X) &= \sup_{\mu, \sigma^2} \left\{ \beta \mu - \frac{N}{2k} \left(\log(\sigma^2/\hat{\sigma}^2) + \frac{\frac{1}{N} \sum_{n=1}^N (X_n - \mu)^2}{\sigma^2} - 1 \right) \right\} \\ &= \sup_{\mu, \sigma^2} \left\{ \beta \mu - \frac{N}{2k\sigma^2} (\bar{X} - \mu)^2 - \frac{N}{2k} \left(\log(\sigma^2/\hat{\sigma}^2) + \frac{\hat{\sigma}^2}{\sigma^2} - 1 \right) \right\} \\ &= \beta \bar{X} + \sup_{\sigma^2} \left\{ \frac{\beta^2 k}{2N} \sigma^2 - \frac{N}{2k} \left(\log(\sigma^2/\hat{\sigma}^2) - 1 + \frac{\hat{\sigma}^2}{\sigma^2} \right) \right\}. \end{aligned} \tag{1}$$

This is problematic, as for any finite N , the growth of σ^2 will dominate the term in braces as $\sigma^2 \rightarrow \infty$, which implies $\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,1}(\beta X) = \infty$.

However, the function to be optimized is increasing near zero, and we can find a local minimum value of the derivative at $\sigma^2 = 2\hat{\sigma}^2$, where

$$\left. \frac{\partial}{\partial \sigma^2} \left(\frac{\beta^2 k}{2N} \sigma^2 - \frac{N}{2k} \left(\log(\sigma^2/\hat{\sigma}^2) - 1 + \frac{\hat{\sigma}^2}{\sigma^2} \right) \right) \right|_{\sigma^2=2\hat{\sigma}^2} = \frac{\beta^2 k}{2N} - \frac{N}{2k} \frac{1}{4\hat{\sigma}^2}.$$

Provided this quantity is negative, that is, $\hat{\sigma}^2 > (\beta k/2N)^2$, we know that the derivative has changed sign at least once on the interval $\sigma^2 \in (0, 2\hat{\sigma}^2)$, so we can be sure that the function will have a local maximum, and (we shall see that) this occurs near $\sigma^2 \approx \hat{\sigma}^2$.

This example leads us to consider further ways of restricting our attention, to ensure that it is a *local* maximum which is chosen in the DR-expectation, as this corresponds to a value close to the MLE, and avoids the explosion of the expectation caused by considering the implausible model $\sigma^2 \rightarrow \infty$.

The following are some natural approaches to preventing this explosion:

- We could place a prior probability on σ^2 , with density f , which collapses sufficiently quickly when $\sigma^2 \rightarrow \infty$, and then incorporate this prior into the penalty. This would result in a penalty given by the negative log-posterior density,

$$\alpha_{Q|\mathbf{x}_N}(Q) = -\ell(\mathbf{x}_N; Q) - \log f(\sigma_Q^2) + \text{normalization term}$$

where the normalizing term is constant with respect to Q , and ensures $\inf_{Q \in \mathcal{Q}} \alpha_{Q|\mathbf{x}_N}(Q) = 0$. In order to ensure $\mathcal{E}_{Q|\mathbf{x}_N}^{k,1}(\beta X) < \infty$ for large N , we require $f(\sigma^2)e^{a\sigma^2} \rightarrow 0$ as $\sigma^2 \rightarrow \infty$, for some $a > 0$, which is not the case for the classical (inverse Gamma) conjugate prior. This also seems a rather *ad hoc* fix, and raises the concern that the conclusions drawn, no matter how much data is available, depends principally on the choice of prior distribution $f(\sigma^2)$.

- We could *a priori* truncate σ^2 away from ∞ (i.e. set a maximum value $\bar{\sigma}^2$ which will be considered). This is equivalent to a uniform prior $f(\cdot) \propto I_{[0, \bar{\sigma}^2]}$ for σ^2 , and will result in the local maximum being chosen whenever N is sufficiently large. However this has the drawback that it requires us to posit a maximum value of σ^2 independently of the data, and again our conclusions will depend on the choice of this upper bound, for all finite values of N .
- We could restrict to the case $\gamma = \infty$, and thus only consider a likelihood interval of values for σ^2 . This has the implication that our DR-expectation is then positively homogeneous, and we ‘flatten out’ the interaction between the choice of model and the variable being considered. (In particular, the penalty term only takes the values 0 and ∞ — models are either excluded or are considered as reasonable as the best-fitting model, with no middle ground.)

The final of these options has the particular, advantage that the conclusions are purely based on the data, rather than our prior assumptions. At the same time, the flattening of our expectation may be undesirable, as it makes it difficult to see how different models will be chosen when evaluating different random variables. For this reason, we propose an alternative regularization technique, combining the $\gamma < \infty$ and $\gamma = \infty$ cases. To achieve this, we define the following transformation of the divergence.

Definition 3.1. The δ -truncation of the divergence, based on a sample of size N , is given by

$$\alpha_{\mathcal{Q}|\mathbf{x}_N}^{(\delta)}(Q) = \begin{cases} \alpha_{\mathcal{Q}|\mathbf{x}_N}(Q) & \text{if } \alpha_{\mathcal{Q}|\mathbf{x}_N}(Q) \leq N^\delta, \\ \infty & \text{otherwise.} \end{cases}$$

The corresponding DR-risk-assessment is defined by

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma,\delta} \circ \mathcal{R})(\xi) = \sup_{Q \in \mathcal{Q}} \left\{ \mathcal{R}(\mathcal{L}_Q(\xi)|\mathbf{x}_N) - \left(\frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}_N}^{(\delta)}(Q) \right)^\gamma \right\}. \quad (2)$$

Remark 3.1. For bounded random variables ξ , taking $\mathcal{R}(\mathcal{L}_Q(\xi)|\mathbf{x}) = E_Q[\xi|\mathbf{x}_N]$, the optimization step in calculating the DR-expectation implies that only those measures in $\{Q : \alpha_{\mathcal{Q}|\mathbf{x}_N}(Q) < \|\xi\|_\infty\}$ need be considered. Therefore, provided $N^\delta > \|\xi\|_\infty$, we have $\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma,\delta}(\xi) = \mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma}(\xi)$.

Remark 3.2. Suppose \mathcal{Q} is sufficiently regular and N sufficiently large that Wilks' theorem holds.³ Then the set of measures $\{Q : \alpha_{\mathcal{Q}|\mathbf{x}_N}(Q) < N^\delta\}$ corresponds to a confidence set with confidence level $F_{\chi_d^2}^{-1}(2N^\delta)$, for d the dimension of (a parametrization of) \mathcal{Q} , where $F_{\chi_d^2}^{-1}$ is the inverse cdf of the χ_d^2 distribution. In this case, we observe that the δ truncation corresponds to considering only those measures in a confidence set around the MLE (and then penalizing further using their divergences based on the data). However, the confidence level of this set will grow quickly with N (when $d = 1$, already for $N^\delta \approx 2$ we are considering a 95% confidence set).

We now consider the behavior of a sequence $\{\mathcal{R}_N\}_{N \in \mathbb{N}}$ of risk assessments, which may vary with our sample size N .

Definition 3.2. We say that $(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma,\delta} \circ \mathcal{R}_N)(\xi)$ is regular whenever the supremum in its definition is finite and attained at a point where $\alpha_{\mathcal{Q}|\mathbf{x}_N}(Q) < N^\delta$.

³This classical theorem states that $\alpha_{\mathcal{Q}|\mathbf{x}_N}(P)$ has an asymptotic χ_d^2 distribution under P , where d is the dimension of a parameter space representing \mathcal{Q} . Conditions under which this holds (which can be interpreted as requirements on the family \mathcal{Q}) can be found in Lehmann and Casella [13].

This definition gives us a criterion, which can be evaluated for a particular set of observations, to determine whether the estimation of our truncated DR-risk assessment is reliable. If our DR-risk assessment is not regular, then our estimate is based on parameters on the boundary of the admissible region $\{Q : \alpha_{Q|\mathbf{x}_N}(Q) < N^\delta\}$. This suggests our estimate depends critically on the choice of truncation parameter δ . On the other hand, when our DR-risk assessment is regular, we know that we have chosen a local maximum within the admissible region, which suggests that the variation in the truncation parameter δ has no marginal impact on our estimates.

The fact that Definition 3.2 depends on the particular observations \mathbf{x}_N is important, as it provides us with a qualitative criterion to assess reliability of estimates which can be calculated from our observations. At the same time, it is interesting to know whether there exists a choice of truncation parameter δ which will typically result in a regular DR-risk assessment.

Definition 3.3. We say $(\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma} \circ \mathcal{R}_N)(\xi)$ is (Q) -regularizable if, for some $\delta > 0$, for all $P \in \mathcal{Q}$, with P -probability tending to 1, as $N \rightarrow \infty$, we have that $(\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma,\delta} \circ \mathcal{R}_N)(\xi)$ is regular.

Remark 3.3. Typically, the divergence will grow like $O(N)$ outside of any neighborhood of the MLE. Therefore, regularizing with $\delta < 1$ is sufficient to guarantee that only points in a neighborhood of the MLE are considered in the optimization. The question is then whether, for risk assessments \mathcal{R}_N , we can choose δ sufficiently large that the local maximum is strictly in the interior of the restricted neighborhood. If this can be done, then the problem is regularizable. This leads to a delicate interplay between the class of risk assessments, the sample size, and the estimation problem itself.

Example 3.2. In the normal example above, $\mathcal{E}_{Q|\mathbf{x}_N}^{k,1,\delta}(\beta X)$ is regular whenever $\beta = o(N^{(\delta+1)/2})$ and $\delta < 1$ (for large N , with P -probability approaching 1 for every $P \in \mathcal{Q}$). Consequently, $\mathcal{E}_{Q|\mathbf{x}_N}^{k,1}(\beta X)$ is regularizable for any fixed β . This is suggested by the following sketch argument (and will be rigorously proven later).

As $\delta < 1$, we are restricting to a neighborhood of the MLE, so we can be confident that we will select a local extremum. We need only to verify that this occurs at a point with penalty $\alpha_{Q|\mathbf{x}_N}(\mu_*, \sigma_*^2) < N^\delta$. For large N , we approximate $\alpha_{Q|\mathbf{x}_N}$ by a quadratic in (μ, σ^2) . This gives us the approximate optimization problem

$$\mathcal{E}_{Q|\mathbf{x}_N}^{k,1,\delta}(\beta X) \approx \sup_{\mu, \sigma^2} \left\{ \beta \mu - \left\{ \frac{N}{2k\sigma^2} (\bar{X} - \mu)^2 + \left(\frac{N}{4k} \right) \left(\frac{\sigma^2}{\hat{\sigma}^2} - 1 \right)^2 \right\} \right\}.$$

The extremum of this approximation is attained by

$$\mu_* = \bar{X} + \frac{k}{N}\beta\sigma^2, \quad \sigma_*^2 = \hat{\sigma}^2 \left(1 + \frac{\beta^2 k^2}{N^2} \hat{\sigma}^2\right).$$

We now substitute back into the quadratic approximation, to obtain

$$\begin{aligned} \alpha_{\mathcal{Q}|\mathbf{x}_N}(\mu_*, \sigma_*^2) &\approx \frac{k\beta^2}{2N} \hat{\sigma}^2 \left(1 + \frac{\beta^2 k^2}{N^2} \hat{\sigma}^2\right) + \left(\frac{N}{4k}\right) \left(\frac{\beta^2 k^2}{N^2} \hat{\sigma}^2\right)^2 \\ &= O_P(N^{-1}\beta^2 + N^{-3}\beta^4). \end{aligned}$$

Provided $\beta = o(N^{(\delta+1)/2})$, this guarantees $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\mu_*, \sigma_*^2) = o_P(N^\delta)$, so our DR-expectation is regular.

For fixed β (or more generally $\beta = o(N^{1-\eta})$ for some $\eta > 0$), we can find a δ satisfying these conditions. Therefore, with P -probability approaching one for every $P \in \mathcal{Q}$, as $N \rightarrow \infty$, the DR-expectation is regular. Therefore, the DR-expectation is regularizable.

To make the argument above rigorous and more generally applicable, we proceed under the following assumption on our observations.

Assumption 3.1. *Assume that*

- (i) *The family \mathcal{Q} is parametrized by an open set $\Theta \subseteq \mathbb{R}^d$.*
- (ii) *Under each $Q \in \mathcal{Q}$, the observations $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$ are iid with density of the form*

$$f(x; Q) = h(x) \exp \left\{ \langle \theta, T(x) \rangle - A(\theta) \right\}.$$

(in other words, \mathcal{Q} is an exponential family with its natural parametrization), where T and A are fixed functions (called respectively the sufficient statistic and the log-partition function).

- (iii) *The variable of interest is $\xi = \phi(X)$, for some Borel function ϕ .*
- (iv) *The Hessian $\mathfrak{I}_\theta = \partial^2 A(\theta)$ (commonly known as the information matrix) is strictly positive definite at every point of Θ .*
- (v) *The \mathcal{Q} -MLE exists and is consistent, with probability tending to 1 as $N \rightarrow \infty$ (that is, for every $P \in \mathcal{Q}$, a maximizer \hat{Q}_N of the likelihood in \mathcal{Q} exists with P -probability approaching 1 and $\theta_{\hat{Q}_N} \rightarrow_P \theta_P$).*

The above assumption leads to the following estimate, which gives us a good understanding of the asymptotic behavior of $\alpha_{\mathcal{Q}|\mathbf{x}_N}$. For simplicity of notation, we shall write $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\theta)$ for $\alpha_{\mathcal{Q}|\mathbf{x}_N}(Q_\theta)$, where Q_θ is the measure parametrized by $\theta \in \Theta$.

Lemma 3.1. *Suppose Assumption 3.1 holds. Then for every $P \in \mathcal{Q}$, there exists a constant $C > 0$ depending on P but independent of N , such that we have the uniform bound*

$$P\left(\alpha_{\mathcal{Q}|\mathbf{x}_N}(\theta) \geq \frac{N}{C} \|\theta - \hat{\theta}_N\| (1 \wedge \|\theta - \hat{\theta}_N\|) \text{ for all } \theta\right) \rightarrow 1.$$

Proof. This follows from the proof of Lemma 3 in [1]. \square

The following abstract approximation result will be useful.

Theorem 3.1. *Consider the maximization of the general function*

$$f_N(\psi) := g_N(\psi) - \frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi)$$

Suppose that $\alpha_{\mathcal{Q}|\mathbf{x}_N}$ arises from a setting where the result of Lemma 3.1 can be applied and, for some $\epsilon \in]0, 1/2]$, for every $P \in \mathcal{Q}$, for every N the function g_N is C^3 and satisfies, for B a ball of constant radius around $\psi = 0$,

- (i) $g_N(0) = 0$,
- (ii) $\|g'_N(\psi)\| = o_P(N^{1-\epsilon})$ uniformly on B ,
- (iii) $\|g''_N(0)\| = o_P(N)$,
- (iv) $\|g'''_N(\psi)\| = O_P(N)$ uniformly on B .

Then defining

$$\delta = 1 - 2\epsilon$$

there exists a point ψ_N^ such that, for each $P \in \mathcal{Q}$, with P -probability approaching 1 as $N \rightarrow \infty$, we know that ψ_N^* maximizes the value of f_N on the set $\{\psi : \alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) \leq N^\delta\}$ and, furthermore, $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi_N^*) = o_P(N^\delta)$.*

Proof. See Appendix. \square

Example 3.3. In the normal setting considered in Example 3.2, we are seeking to maximize

$$\beta\mu - \frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}}(\mu, \sigma^2) = \beta\bar{X} + \{\beta\psi_\mu - \frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}}(\psi)\}$$

where $\psi = (\psi_\mu, \psi_{\sigma^2}) = (\mu - \bar{X}, \sigma^2 - \hat{\sigma}^2)$. Provided $\beta = o(N^{1-\eta})$ for some $\eta > 0$, we can take $g(\psi) = \beta\psi_\mu$, which satisfies the conditions of Theorem 3.1. By application of the theorem, it follows that the DR-expectation is regularizable in this case, giving a rigorous proof of the condition obtained by quadratic approximation in Example 3.2.

In the above example, we have considered the case $\gamma = 1$. We can also obtain a sufficient condition for the case $\gamma > 1$ using an extension of this method, assuming that $\epsilon = 1/2$.

Lemma 3.2. *Suppose that, in the case $\gamma = 1$, our truncated optimization problem (2) has a maximizer Q_1^* where $\alpha_{Q|\mathbf{x}_N}(Q_1^*) < 1$, when taking a truncation parameter $\delta \geq 0$. Then the truncated problem in the case $\gamma \geq 1$ certainly has a maximizer Q_γ^* with $\alpha_{Q|\mathbf{x}_N}(Q_\gamma^*) < 1$, using the same truncation parameter.*

Proof. This is immediate from the fact that $|x|^\gamma < |x|$ for $|x| < 1$ and $|x|^\gamma > |x|$ for $|x| > 1$. \square

Example 3.4. In the normal setting, we see that if $\beta = o(N^{1/2})$ then Theorem 3.1 yields the regularization parameter $\delta = 1 - 2\epsilon = 0$. In other words, we truncate to a set $\alpha_{Q|\mathbf{x}_N} < o_P(1)$. By Lemma 3.2, this implies that the same regularization will be sufficient for all choices of $\gamma < \infty$.

This argument suggests that $\beta = o(N^{1/2})$ is a fundamental requirement for reliable statistical estimation of $E[\beta X]$ from N observations, under the assumption that X and our observations are iid normal with unknown mean and variance. Comparing with the case $\gamma = \infty$, where the DR-expectation is the upper bound of a confidence interval, we would need $\beta = o(N^{1/2})$ to ensure that

$$|\mathcal{E}_{Q|\mathbf{x}_N}^{k,\infty}(\beta X) - E_{\hat{Q}}[\beta X]| \approx \sqrt{2k} \frac{|\beta|}{\sqrt{N}} \hat{\sigma} \rightarrow 0,$$

where \hat{Q} is the MLE.

Remark 3.4. The example above shows that there is a close asymptotic relationship between regularity and the convergence of confidence intervals. However, the concept of regularity has the advantage that it can be applied for a given sample (of fixed size), rather than only being meaningful in an asymptotic sense. In particular, for a given sample, we can evaluate (for fixed $\delta \geq 0$) whether our truncated DR-expectation is regular, which allows us to qualitatively assess the statistical reliability of our estimates *based on the data at hand, and for the variables of interest*. In particular, if the function to be maximized when calculating $(\mathcal{E}_{Q|\mathbf{x}_N}^{k,\gamma} \circ \mathcal{R})(\xi)$ does not have a local maximum near the MLE, then we know that the DR-risk assessment is not regular for any δ . This is unlike a confidence interval, which will generally not give a qualitative assessment of reliability for any fixed sample.

If we find that our problem is not regular, then we know that the observed data are insufficient to guide our calculation of DR-expectations — our choice of regularization has an overriding impact on our conclusions.

4. Heavy tails

The example given in the previous section shows us that, in a Gaussian setting, provided β is not too large relative to N (in particular $\beta = o(N^{1-\eta})$ for some $\eta > 0$), the DR-expectation $\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,1}(\beta X)$ is regularizable. This suggests asymptotic bounds on the size of random variables we are willing to consider, if we wish to account for our statistical uncertainty. In what follows, we shall seek to show that a similar relationship exists in a heavy tailed setting, when we calculate expected shortfall and related quantities.

We shall focus our attention on the following special case.

Assumption 4.1. *The family of measures \mathcal{Q} describes models under which our observations $\{X\} \cup \{X_n\}_{n \in \mathbb{N}}$ are iid from a Pareto distribution (with known minimal value 1), that is, for a model parametrized by θ , we have the density*

$$f(x; \theta) = \theta x^{-(1+\theta)}; x > 1.$$

The possible values of θ will be assumed to satisfy one of two cases:

- (i) \mathcal{Q} corresponds to all $\theta > 0$ (which implies a valid probability distribution for X , but no integrability),
- (ii) \mathcal{Q} corresponds to all $\theta > 1$ (which implies X is integrable).

We should note that this is an exponential family of distributions, so the above approximation results (in particular Lemma 3.1) can be applied. Clearly Assumption 4.1(ii) is stronger than Assumption 4.1(i).

It is an easy exercise to show the following:

Proposition 4.1. *In the setting of Assumption 4.1, the MLE is given by*

$$\hat{\theta} = \frac{N}{\sum_i \log(x_i)}$$

and the divergence by

$$\alpha_{\mathcal{Q}|\mathbf{x}_N}(\theta) = N \left(-\log \left(\frac{\theta}{\hat{\theta}} \right) - 1 + \frac{\theta}{\hat{\theta}} \right).$$

Remark 4.1. This is a significantly simplified version of a standard extreme value estimation problem. In particular, consider a model for excesses

over a threshold u . By the Pickands–Balkema–de Haan theorem (see for example Embrechts, Klüppelberg and Mikosch [5]), we know that for large u , the probability distribution in such a setting is typically approximated by the generalized Pareto distribution

$$P(X - u \leq y | X > u) \approx \begin{cases} 1 - \left(1 + \frac{ky}{\sigma}\right)^{-1/k} & \text{if } k \neq 0, \\ 1 - e^{-y/\sigma} & \text{if } k = 0. \end{cases}$$

If we assume that this is the correct distribution family (so we are in the limiting regime), the scale variable σ is known (for simplicity) and $k > 0$ (so our distribution is unbounded and not exponential), then estimating k is equivalent to estimating θ in the above (classical) Pareto setting (after reparametrization).

Of course, in practice, the estimation of σ and the choice of cut-off value u are significant issues in their own right. What we shall see is that, even without these additional concerns, the estimation error associated with k (equivalently θ) is sufficient to restrict which risk assessments can be reliably estimated.

In what follows, we shall give relationships, for a variety of common risk assessment methods, between the growth of a parameter β (which will describe the ‘extremity’ of our risk assessment) and the sample size N , such that the DR-risk assessment is regularizable. This gives a method of assessing the reliability of our statistical estimates. We should point out that, if this were to be used in the context of fitting the excesses over a threshold, our sample size N would refer exclusively to the number of excesses we observe, and any probabilities would be conditional on being over the threshold point.

4.1. Expected shortfall

The first example we shall consider is the DR-expected shortfall, as defined in Example 2.1. As expected shortfall is a convex expectation, we know that the DR-expected shortfall is also a convex expectation. Given the expected shortfall is only defined for integrable random variables, we shall work under Assumption 4.1(ii).

Lemma 4.1. *Without truncation, under Assumption 4.1(ii), the DR-expected shortfall is infinite for all N , and all probability levels β , that is*

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{ES}_\beta)(X) = \infty$$

for every set of observations \mathbf{x}_N and all choices of $k, \gamma < \infty$.

Proof. Under Assumption 4.1, for a given $\theta > 1$, we can calculate the Expected shortfall with tail probability β ,

$$\text{ES}_\beta = E[X|X > F_\theta^{-1}(1 - \beta)] = \frac{\theta}{\theta - 1} \beta^{-1/\theta}.$$

As $\theta \downarrow 1$, we observe that $\text{ES}_\beta \rightarrow \infty$, but $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\theta) \not\rightarrow \infty$. Therefore, for every $\gamma < \infty$, it is easy to see that $(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{ES}_\beta)(X) = \infty$. \square

Remark 4.2. In the case $\gamma = \infty$, by monotonicity we will have

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{ES}_\beta)(X) = \frac{\theta^*}{\theta^* - 1} \beta^{-1/\theta^*},$$

where θ^* is the smallest value⁴ such that $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\theta) \leq k$. This condition is the same for all N – in particular, using only a single observation, we give a finite estimate of the expected shortfall at any level $\beta \approx 0$ (i.e. arbitrarily far into the tail), which is intuitively unreasonable. This is the ‘confidence interval’ bound discussed in Embrechts, Klüppelberg and Mikosch [5].

This problem can be treated through regularization techniques, giving qualitatively different conclusions.

Theorem 4.1. *Under Assumption 4.1(ii), in the case $\gamma = 1$, the DR-expected shortfall with tail probability β is regularizable whenever $\beta^{-1} = o(N^{1-\eta})$, for some $\eta > 0$.*

Proof. As the MLE is consistent, we know that for large N , $\hat{\theta} > 1$ with P -probability approaching 1, for every $P \in \mathcal{Q}$. In what follows, we assume this is the case.

The quantity to be (locally) maximized when calculating (2) is

$$\frac{\theta}{\theta - 1} \beta^{-1/\theta} - \frac{N}{k} \left(\log(\hat{\theta}/\theta) - 1 + \theta/\hat{\theta} \right).$$

To apply Theorem 3.1, as we are interested in small values of β , we calculate the derivatives

$$\begin{aligned} \frac{d}{d\theta} \left(\frac{\theta}{\theta - 1} \beta^{-1/\theta} \right) &= \beta^{-1/\theta} \left(\frac{\log \beta}{\theta(\theta - 1)} - \frac{1}{(\theta - 1)^2} \right), \\ \frac{d^2}{d\theta^2} \left(\frac{\theta}{\theta - 1} \beta^{-1/\theta} \right) &= \beta^{-1/\theta} \left(\frac{2}{(\theta - 1)^3} - \frac{2 \log \beta}{\theta(\theta - 1)^2} + \frac{(\log \beta)^2}{\theta^3(\theta - 1)} \right), \\ \frac{d^3}{d\theta^3} \left(\frac{\theta}{\theta - 1} \beta^{-1/\theta} \right) &= O(\beta^{-1/\theta} (\log \beta)^3) \quad \text{uniformly near } \theta > 1. \end{aligned}$$

⁴This assumes $\theta^* > 1$, otherwise our risk assessment takes the value ∞ .

For N large, uniformly in a neighborhood of $\hat{\theta} > 1$, taking

$$g_N(\psi) = \frac{\hat{\theta} + \psi}{\hat{\theta} + \psi - 1} \beta^{-1/(\hat{\theta} + \psi)} - \frac{\hat{\theta}}{\hat{\theta} - 1} \beta^{-1/\hat{\theta}}$$

the requirements of Theorem 3.1 are guaranteed by the simple condition

$$\frac{1}{\beta} = o_P(N^{\hat{\theta} - \eta}) \quad \text{for some } \eta > 0, \text{ for all } P \in \mathcal{Q}.$$

As $\hat{\theta} > 1$ with P -probability approaching 1 for $P \in \mathcal{Q}$ (by consistency of the MLE), we obtain the deterministic bound desired. \square

Remark 4.3. This result places a qualitative bound on the extremity of the expected shortfall that can be reliably estimated, when we measure uncertainty through the DR-expected shortfall with $\gamma = 1$. Given N observations, for Pareto distributed models assuming only integrability, we cannot generally claim to estimate expected shortfalls for probabilities below c/N , where c is some constant. In particular, estimating extreme tails from small numbers of observations is shown to be unreliable.

Corollary 4.1. *Under Assumption 4.1(ii), in the case $\gamma < \infty$, the DR-expected shortfall with tail probability β is regularizable whenever $\beta^{-1} = o_P(N^{1/2 - \eta})$, in particular when $\beta^{-1} = o(N^{1/2 - \eta})$, for some $\eta > 0$.*

Proof. By assumption, we know that $\beta^{-1} = o_P(N^{\hat{\theta}/2 - \eta})$, we can apply Theorem 3.1 with $\epsilon = 1/2$. This gives us a regularization parameter $\delta = 0$, and from Lemma 3.2 we observe regularity for all $\gamma < \infty$. \square

Remark 4.4. In practice, this is still a very optimistic requirement for reliable estimation, as we have assumed that our simple Pareto model is correct (and no further parameters need to be estimated). In this sense, these results give ‘best-case’ bounds on how far into the tail we can look before losing reliability of expected shortfall estimation using a Pareto model.

4.2. Value at risk

The Value at Risk (with tail probability β) is not a convex risk measure, so it is not generally true that $\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k, \gamma} \circ \text{VaR}_\beta$ is a convex expectation. Nevertheless, we can calculate

$$\text{VaR}_\beta(X) = F_\theta^{-1}(1 - \beta) = \beta^{-1/\theta}.$$

Given this is well defined (and finite) for every distribution, we shall proceed under Assumption 4.1(i).

Lemma 4.2. *Without truncation, under Assumption 4.1(i), the DR-value at risk is infinite for all N , and all probability levels $\beta < 1$, that is*

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{VaR}_\beta)(X) = \infty$$

for every set of observations \mathbf{x}_N and all choices of $k, \gamma < \infty$.

Proof. As in the Expected Shortfall case, we observe that $\beta^{-1/\theta} \rightarrow \infty$ as $\theta \downarrow 0$, but $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\theta) \not\rightarrow \infty$ as $\theta \rightarrow 0$. The result follows. \square

Remark 4.5. If we assumed Assumption 4.1(ii), then for $\beta \rightarrow 0$ and finite N , we would obtain $(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{VaR}_\beta)(X) \approx 1/\beta$, independently of the observed values. Clearly this is not reliable statistically, as it is the assumption of integrability, rather than the observations, which is leading to finiteness of the estimate.

Theorem 4.2. *Under Assumption 4.1(i), for all $\gamma < \infty$, the DR-Value at Risk with tail probability β is regularizable whenever $\beta^{-1} = O(1)$ (as $N \rightarrow \infty$), for some $\eta > 0$.*

Proof. As in the Expected Shortfall case, we see that

$$\frac{d^n}{d\theta^n}(\beta^{-1/\theta}) = O(\beta^{-1/\theta}((\log \beta)^n + 1)).$$

Applying Theorem 3.1 leads to the proposed condition for regularizability (in the case $\gamma = 1$)

$$\frac{1}{\beta} = o_P(N^{\hat{\theta}-\eta}) \quad \text{for some } \eta > 0, \text{ for all } P \in \mathcal{Q}.$$

Again, using Lemma 3.2, for regularizability of the $\gamma > 1$ case, we obtain the condition $\beta^{-1} = o_P(N^{\hat{\theta}/2-\eta})$ for some $\eta > 0$. Consistency of the MLE (and the assumption $\theta > 0$ for all P), shows that this is guaranteed when $\beta^{-1} = O(1)$ as $N \rightarrow \infty$. \square

Remark 4.6. The requirement $\beta^{-1} = O(1)$ is quite restrictive, but comes from the fact we are assuming nothing beyond Assumption 4.1(i), i.e. that our distribution is well defined. Strengthening Assumption 4.1(i) to restrict to $\theta > \tilde{\theta}$ for some $\tilde{\theta} > 0$, we have regularizability whenever $\beta^{-1} = o(N^{\tilde{\theta}/2-\eta})$ for some $\eta > 0$.

4.3. Probability of loss

The probability of a loss exceeding a level β is given by (under Assumption 4.1(i))

$$\text{PL}_\beta(X) = P(X > \beta) = 1 - F_\theta(\beta) = \beta^{-\theta}.$$

This is not a convex expectation, however we can equivalently express it as $P(X > \beta) = E[I_{X>\beta}]$, and then consider the regularity of

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{PL}_\beta)(X) = \mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma}(I_{X>\beta}).$$

Theorem 4.3. *The DR-probability of a loss is regular, for all N , γ and β .*

Proof. As $I_{X>\beta}$ is bounded, this is well behaved without regularization, no matter what the choice of β . \square

Remark 4.7. The asymptotic behavior of the DR-probability of loss (and other bounded random variables) is described by [1], Section 3.2.

4.4. Integrated tail and Cramér–Lundberg failure probability

From an insurance perspective, it is sometimes of interest to look at the integrated tail, which under Assumption 4.1(ii) is given by

$$\text{IT}_\beta(X) := E[(X - \beta)^+] = \int_\beta^\infty (1 - F_\theta(x))dx = \int_\beta^\infty x^{-\theta}dx = \frac{\beta^{1-\theta}}{\theta - 1}.$$

Remark 4.8. For $\beta \geq 1$, this is a convex map, but not translation invariant, so the DR-integrated tail, $(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{IT}_\beta)(\cdot)$, is not generally a convex expectation.

As in the previous cases, without truncation, the DR-integrated tail poses some problems.

Lemma 4.3. *Without truncation, under Assumption 4.1(ii), the DR-integrated tail is infinite for all N , and all $\beta \geq 1$, that is*

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{IT}_\beta)(X) = \infty$$

for every set of observations \mathbf{x} and all choices of $k, \gamma < \infty$.

Proof. By direct calculation, considering $\theta \downarrow 1$,

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{IT}_\beta)(X) = \sup_\theta \left\{ \frac{\beta^{1-\theta}}{\theta - 1} - \frac{N}{k} \left(\log(\hat{\theta}/\theta) - 1 + \theta/\hat{\theta} \right) \right\} = \infty.$$

\square

The surprising result is that, while we need to truncate to avoid infinite values, the DR-integrated tail is always regularizable.

Theorem 4.4. *Under Assumption 4.1(ii), for all $\gamma < \infty$, the DR-integrated tail is regularizable for all choices of $\beta \geq 1$ (with any desired dependence on N).*

Proof. For all $P \in \mathcal{Q}$, as in the earlier settings, we calculate

$$\frac{d^n}{d\theta^n} \left(\frac{\beta^{1-\theta}}{\theta-1} \right) = O(\beta^{1-\theta}((\log \beta)^n + 1)).$$

However, for $\beta \geq 1$, $\theta > 1$, the right hand side is $o(1)$ with respect to β , in particular (by consistency of the MLE) it is $o_P(1)$ for every $P \in \mathcal{Q}$. Therefore, the integrated tail is regularizable with no restriction on the value of β , for every value of $\gamma < \infty$. \square

A closely related quantity is the related failure probability under a Cramér–Lundberg model, given by (under Assumption 4.1(ii))

$$\text{CL}_\beta(X) = \frac{\text{IT}_\beta}{E[X]} = \frac{\beta^{1-\theta}}{\theta}.$$

Theorem 4.5. *For all choices of $\beta \geq 1$, under Assumption 4.1(ii), the DR-Cramér–Lundberg failure probability is regularizable.*

Proof. We calculate

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{CL}_\beta)(X) = \sup_{\theta} \left\{ \frac{\beta^{1-\theta}}{\theta} - \frac{N}{k} \left(\log(\hat{\theta}/\theta) - 1 + \theta/\hat{\theta} \right) \right\}.$$

For $\beta \geq 1$, we know $\frac{\beta^{1-\theta}}{\theta}$ is (uniformly) bounded for all $\theta > 1$. Therefore, from the consistency of the MLE and Theorem 3.1, $(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ \text{CL}_\beta)(X)$ is regularizable. \square

4.5. Distortion risk

As a final example, we consider a distortion based nonlinear expectation. This is given by taking a convex increasing bijective map $\lambda : [0, 1] \rightarrow [0, 1]$, then calculating the expectation with the transformed cdf $F_\lambda(x) = \lambda(F(x))$. As λ is convex, it is differentiable almost everywhere, and we can calculate the transformed density under Assumption 4.1,

$$f_\lambda(x) = \lambda'(F(x))f(x) = \lambda'(1 - x^{-\theta})\theta x^{-(1+\theta)}.$$

so the distortion risk is given by

$$D_\lambda(X) = \int_1^\infty \lambda'(1 - x^{-\theta})\theta x^{-\theta} dx. \quad (3)$$

Lemma 4.4. *The distortion risk is finite whenever there exists $\zeta > 1/\theta$ such that*

$$\lambda'(y) = O((1 - y)^{-1+\zeta}) \quad \text{as } y \rightarrow 1.$$

Proof. In order for D_λ to be finite, as $\theta > 0$, for large x we require $\lambda'(1 - x^{-\theta})$ not to be too large, else we lose integrability in (3). As we know $x^{-(1+\zeta)}$ is integrable on $[1, \infty[$, the result follows by dominated convergence. \square

As λ is increasing and convex, for any $y < 1$, we know that $\{\lambda'(x)\}_{x < y}$ is bounded, so it is the behavior of λ' near 1 which is of interest. For this reason, we will focus our attention on the following example.

Definition 4.1. We say λ is the ‘minmaxvar’ transform (with parameter $\beta \in (0, 1]$) if $\lambda(x) = \lambda_\beta(x) := 1 - (1 - x)^\beta$.

This case is interesting from our perspective, as it describes the critical growth of λ near the boundary $x = 1$. Other classic examples, for example the Wang transform $\lambda(x) = \Phi(\Phi^{-1}(x) - \beta)$, where Φ is the normal cdf, are also of interest in some settings, but do not have this critical growth. This is closely related to the ‘minmaxvar’ risk measure considered by Cherny and Madan [14].

Lemma 4.5. *Without truncation, under Assumption 4.1(ii), the DR-minmaxvar risk is infinite for all N , and all $\beta \in (0, 1)$, that is*

$$(\mathcal{E}_{\mathcal{Q}|\mathbf{x}_N}^{k,\gamma} \circ D_{\lambda_\beta})(X) = \infty$$

for every set of observations \mathbf{x}_N and all choices of $k, \gamma < \infty$.

Proof. We know $\lambda'_\beta(x) = \beta(1 - x)^{\beta-1}$, so

$$D_{\lambda_\beta} = \int_1^\infty \lambda'_\beta(1 - x^{-\theta})\theta x^{-\theta} dx = \int_1^\infty \beta(x^{-\theta})^{\beta-1}\theta x^{-\theta} dx = \frac{\beta\theta}{\beta\theta - 1}$$

provided $\beta > 1/\theta$, and is otherwise infinite. As $\alpha_{\mathcal{Q}|\mathbf{x}}(1/\beta) < \infty$ for all $\beta \in (0, 1)$, our DR-minmaxvar risk will be infinite. \square

Theorem 4.6. *Under Assumption 4.1(ii), the DR-minmaxvar risk is not regularizable for any $\beta < 1$. However, under the stronger assumption that we restrict our models to those where $\theta > \tilde{\theta}$ for some $\tilde{\theta} > 1$, then the DR-minmaxvar risk is regularizable whenever*

$$\beta \geq \frac{1}{\tilde{\theta}} + \frac{1}{|O(N^\eta)|}.$$

where $\eta = 1/4$ in the case $\gamma = 1$, and $\eta < 1/4$ in the case $\gamma < \infty$.

Proof. For any fixed $\beta < 1$, Assumption 4.1(ii), is insufficient to guarantee that the MLE $\hat{\theta} > 1/\beta$ (as $\hat{\theta}$ may be arbitrarily close to 1). As this is the condition for finiteness of D_{λ_β} and the MLE satisfies $\alpha_{Q|\mathbf{x}_N}(\hat{\theta}) \equiv 0$, the non-regularizability follows.

Under our stronger assumption, to determine conditions on β such that the DR-minmaxvar risk is regularizable, we proceed as before. We can calculate

$$(\mathcal{E}_{Q|\mathbf{x}_N}^{k,1} \circ D_{\lambda_\beta})(X) = \sup_{\theta} \left\{ \frac{\beta\theta}{\beta\theta - 1} - \left(\frac{N}{k} \left(\log(\hat{\theta}/\theta) - 1 + \theta/\hat{\theta} \right) \right)^\gamma \right\}.$$

We have the derivatives

$$\begin{aligned} \frac{d}{d\theta} \left(\frac{\beta\theta}{\beta\theta - 1} \right) &= \frac{-\beta}{(\beta\theta - 1)^2}, \\ \frac{d^2}{d\theta^2} \left(\frac{\beta\theta}{\beta\theta - 1} \right) &= \frac{2\beta^2}{(\beta\theta - 1)^3}, \\ \frac{d^3}{d\theta^3} \left(\frac{\beta\theta}{\beta\theta - 1} \right) &= \frac{-6\beta^3}{(\beta\theta - 1)^4}. \end{aligned}$$

With $g_N(\psi) = \frac{\beta(\hat{\theta}+\psi)}{\beta(\hat{\theta}+\psi)-1} - \frac{\beta\hat{\theta}}{\beta\hat{\theta}-1}$, Theorem 3.1 is satisfied by assuming $(\hat{\theta}\beta - 1)^{-4} = O_P(N)$, or equivalently,

$$\beta \geq \frac{1}{\tilde{\theta}} + \frac{1}{|O_P(N^{1/4})|}.$$

Using the consistency of the MLE we know $\hat{\theta} > \tilde{\theta}$, and the result follows. If we make the further assumption that $\beta = \tilde{\theta}^{-1} + 1/|O(N^\eta)|$ for $\eta < 1/4$, we can take $\epsilon = 1/2$ in Theorem 3.1, to obtain regularizability for all $\gamma < \infty$. \square

Remark 4.9. Many further cases can also be considered, using this general approach. It would be interesting also to apply DR-regularization in a more general estimation problem, where we do not assume simply that we have a Pareto distribution, but must use a generalized extreme value or generalized Pareto model, with the associated estimation difficulties.

Appendix

To prove our key approximation result (Theorem 3.1) we begin with the following lemma.

Lemma A.1. *Consider a C^3 function $f : \mathbb{R}^N \rightarrow \mathbb{R}$, with negative definite Hessian at zero H . Suppose we wish to find a local maximum in a small ball B_δ of radius δ around 0. Let $c_\delta = \sup_{x \in B_\delta} \|f'''(x)\|$. If a local maximum exists in the interior of the ball, then it has position*

$$x^* = -H^{-1}f'(0) + c_\delta \|H^{-1}\| O(\delta^2).$$

Proof. By Taylor's theorem, we can write the expansion

$$f(x) = f(0) + x^\top f'(0) + \frac{1}{2} x^\top H x + R(x)$$

where R is some remainder term with $|R(x)| \leq (c_\delta/6)\|x\|^3$ and derivative $\|R'(x)\| \leq (c_\delta/2)\|x\|^2 = c_\delta O(\delta^2)$ on our ball. To find a local extremum x^* , we differentiate to obtain the vector equation

$$0 = f'(0) + Hx^* + R'(x^*)$$

which rearranges to give the approximation (which is true for every interior local extremum)

$$x^* = -H^{-1}f'(0) - H^{-1}R'(x^*).$$

For an extremum within the ball, we have the desired approximation

$$x^* = -H^{-1}f'(0) + c_\delta \|H^{-1}\| O(\delta^2).$$

□

We now combine Lemma A.1 and Lemma 3.1 to give a proof of Theorem 3.1, which we repeat here for the ease of the reader.

Theorem A.1. *Consider the maximization of the general function*

$$f_N(\psi) := g_N(\psi) - \frac{1}{k} \alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi)$$

Suppose that $\alpha_{\mathcal{Q}|\mathbf{x}_N}$ arises from a setting where the result of Lemma 3.1 can be applied and, for some $\epsilon \in]0, 1/2]$, for every $P \in \mathcal{Q}$, for every N the function g_N is C^3 and satisfies, for B a ball of constant radius around $\psi = 0$,

- (i) $g_N(0) = 0$,
- (ii) $\|g'_N(\psi)\| = o_P(N^{1-\epsilon})$ uniformly on B ,

- (iii) $\|g_N''(0)\| = o_P(N)$,
 (iv) $\|g_N'''(\psi)\| = O_P(N)$ uniformly on B .

Then there exists a point ψ_N^* such that, for each $P \in \mathcal{Q}$, with P -probability approaching 1 as $N \rightarrow \infty$, we know that ψ_N^* maximizes the value of f_N on the set $\{\psi : \alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) \leq N^{1-2\epsilon}\}$ and, furthermore, $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi_N^*) = o_P(N^{1-2\epsilon})$.

Proof. By Lemma 3.1, for all N sufficiently large, with arbitrarily high P -probability, for ψ in a neighborhood of zero of fixed radius, there exists a constant $C > 0$ (depending on P) such that

$$\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) > \frac{N}{C} \|\psi\|^2$$

and outside this neighborhood $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) \geq O(N)$. In all that follows, we restrict our attention to this constant radius ball.

Omitting subscript N for simplicity, from our assumptions on g we know

$$g(\psi) = o_P(N^{1-\epsilon}) \|\psi\|.$$

Consequently, except possibly on a ball of radius $O_P(N^{-\epsilon})$ around zero, we know that

$$\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) > \frac{N}{C} \|\psi\|^2 > o_P(N^{1-\epsilon}) \|\psi\| = |g(\psi)|.$$

It follows that, for all ψ outside a ball of radius $O_P(N^{-\epsilon})$, we know that $f(\psi) < 0$. As $f(0) = 0$, there must exist a local maximum within the ball of radius $O_P(N^{-\epsilon})$.

We know that,

$$\begin{aligned} f'(0) &= g'(0) = o_P(N^{1-\epsilon}) \\ \|f''(0)^{-1}\| &\leq \|(g''(0) - N/C)^{-1}\| = O_P(N^{-1}) \\ f'''(\psi) &= g'''(\psi) + \alpha_{\mathcal{Q}|\mathbf{x}_N}'''(\hat{\theta} + \psi) = O_P(N). \end{aligned}$$

Applying Lemma A.1, we know that any local maximum of f within a ball of radius $o_P(N^{-\epsilon/2})$ will be at a point satisfying

$$\psi^* = -f''(0)^{-1}f'(0) + O_P(N)\|f''(0)^{-1}\|o_P(N^{-\epsilon}) = o_P(N^{-\epsilon}).$$

Therefore, *all* local maxima within the ball of radius $o_P(N^{-\epsilon/2})$ will be within the ball of radius $o_P(N^{-\epsilon})$.

Taking a Taylor approximation of the C^3 function $\alpha_{\mathcal{Q}|\mathbf{x}_N}$ we see that, within the ball of radius $o_P(N^{-\epsilon})$,

$$\begin{aligned}\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) &= N\psi^\top (\mathfrak{J}_{\hat{\theta}} + O_P(\|\psi\|))\psi = \|\mathfrak{J}_{\hat{\theta}}\|o_P(N^{1-2\epsilon}) + o_P(N^{1-3\epsilon}) \\ &= o_P(N^{1-2\epsilon}).\end{aligned}$$

Conversely, outside the ball of radius $o_P(N^{-\epsilon/2})$, we know that $\alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) > N^{1-\epsilon}/C$ for some $C > 0$. Therefore, we can be certain that a point $\psi^* = o_P(N^{-\epsilon})$ will be the maximizer within the region $\{\psi : \alpha_{\mathcal{Q}|\mathbf{x}_N}(\hat{\theta} + \psi) < N^{1-2\epsilon}\}$, as desired. \square

Acknowledgments

Research supported by the Oxford–Man Institute for Quantitative Finance and the Oxford–Nie Financial Big Data Laboratory.

References

1. S. Cohen, Data-driven nonlinear expectations for statistical uncertainty in decisions, *Electronic Journal of Statistics* **11**, 1858–1889 (2017).
2. F. Knight, *Risk, Uncertainty and Profit* (Houghton Mifflin, 1921).
3. A. Wald, Statistical decision functions which minimize the maximum risk, *Annals of Mathematics* **46**, 265–280 (1945).
4. H. Föllmer and A. Schied, *Stochastic Finance: An introduction in discrete time* (de Gruyter, 2002).
5. P. Embrechts, C. Klüppelberg and T. Mikosch, *Modelling Extremal events for Insurance and Finance* (Springer, 1997).
6. S. N. Cohen and R. J. Elliott, *Stochastic Calculus and Applications*, second Edn. (Birkhäuser, 2015).
7. A. N. Shiryaev, *Probability*, 2nd Edn. (Springer, 2000).
8. M. Frittelli and E. Rosazza Gianin, Putting order in risk measures, *Journal of Banking & Finance* **26**(7), 1473–1486 (2002).
9. P. Artzner, F. Delbaen, J.-M. Eber and D. Heath, Coherent measures of risk, *Mathematical Finance* **9**(3), 203–228 (1999).
10. H. White, Maximum likelihood estimation of misspecified models, *Econometrica* **50**(1), 1–25 (1982).
11. A. J. McNeil, R. Frey and P. Embrechts, *Quantitative Risk Management* (Princeton, 2005).
12. P. Barrieu and N. El Karoui, Inf-convolution of risk measures and optimal risk transfer, *Finance and Stochastics* **9**(2), 269–298 (2005).

13. E. Lehmann and G. Casella, *Theory of Point Estimation* (Springer, 1998).
14. A. Cherny and D. Madan, New measures for performance evaluation, *The Review of Financial Studies* **22**(7), 2571–2606 (2009).