

Assessing Risk of Exceedance Events with Buffered Probability of Exceedance and Superquantiles

Matthew Norton

Assistant Professor, Dept. of Operations Research, Naval Postgraduate School,
Monterey, CA

ABSTRACT: Many systems, or components of systems, often have associated with them a critical safety threshold. If events then occur that have magnitude larger than such a threshold, such as a weather event or physical stress on a component, critical failures become likely and/or overall system integrity can become critically compromised. Therefore, regulations and risk assessments are often formulated mathematically in terms of Probability of Exceedance (POE). This characteristic, however, can hide important information about the frequency, magnitude, and overall risk of exceedance events. This includes the magnitude of events that do exceed the threshold. Additionally, the frequency and magnitude of near-exceedance events, just below the threshold, are ignored entirely. We overview a new probabilistic characterization of exceedance risk call Buffered Probability of Exceedance (bPOE), also reviewing a closely related concept called the superquantile. We show that bPOE simultaneously assess both the frequency and magnitude of both exceedance events and near-exceedance events. After introducing bPOE and superquantiles, we show how it can be viewed as superior to POE as a measure of exceedance risk. We then present a simple parametric distribution fitting procedure that utilizes bPOE and the superquantile, two characteristics that we see are advantageous to consider when estimation of exceedance risk and tail density are the focus of the fitting procedure.

1. MEASURING RISK

When faced with a random outcome, represented here by a real valued random variable X , it is often critical in engineering applications to properly characterize its inherent uncertainty or risk. Frequently, in this context, engineers will have some idea about what might be considered a *large* outcome or realization of X . For example, one might say that realizations with magnitude larger than some threshold $z \in \mathbb{R}$ would be considered large. In this case, it would be desirable to characterize the uncertainty or riskiness of X relative to the threshold z .

An intuitive and common characterization of uncertainty that considers such a threshold is Probability of Exceedance (POE), $p_z(X) = P(X > z)$. Here, one characterizes the riskiness of X by how frequently its value exceeds z . However, even

though this is a simple and intuitive concept, POE has some limitations.

First, POE does not consider the magnitude of exceedance events or near-exceedance events. For example, suppose you calculate that $P(X > z) = .05$. This tells you how frequently X turns out to be large, but it does not tell you how large X might be when it does exceed z . This can be critical information, such as when X represents a financial loss or structural stressor. In this context, a single stressor, if large enough, could be catastrophic. Additionally, when X nearly exceeds z , how close to z does it come and how frequently does it come so close to z from below? This is also important information that is not considered by POE. Assume, again, that X represents an unknown financial loss and that $P(X > z) = .05$. Then, it is important to know if 95% of the losses equal $z - \$1$, i.e. are

near-exceedance events, or if 5% of the losses equal $z - \$100$ with the remaining 90% being very far below z .

Second, POE is often discontinuous w.r.t. the threshold. For example, assume that X is discretely distributed with finite potential outcomes and suppose you would like to know how $P(X > z)$ varies over $z \in [z_l, z_h]$. One will then discover that $P(X > z)$ is discontinuous w.r.t. the threshold and, furthermore, sensitivity analysis will be difficult as small changes in threshold can lead to large jumps in probability. Thus, the difference $P(X > z_l) - P(X > z_h)$ may be very large, even if $z_h - z_l$ is very small.

Third, when working with optimization of tail probabilities, one frequently works with constraints or objectives involving POE, or its associated quantile $q_\alpha(X) = \min\{z | P(X \leq z) \geq \alpha\}$, where $\alpha \in [0, 1]$ is a probability level. The quantile is a popular measure of tail probabilities in financial engineering, called within this field Value-at-Risk by its interpretation as a measure of tail risk. The quantile and POE, though, when included in optimization problems via constraints or objectives, can often be difficult to treat with continuous (linear or non-linear) optimization techniques.

This paper overviews a recently introduced concept called Buffered Probability of Exceedance (bPOE), a counterpart of POE that does not suffer from these difficulties. Like POE, we have that bPOE measures the frequency with which outcomes are large w.r.t. some threshold z . Unlike POE, though, it is continuous w.r.t. the threshold, considers the magnitude of events beyond or near to the threshold, and can be integrated into optimization frameworks efficiently, often with convex, sometime even linear programming.

1.1. bPOE and Superquantiles

As mentioned above, POE suffers from many difficulties. The quantile, being its inverse, suffers from identical issues. A significant advancement was made in Rockafellar and Uryasev (2000, 2002) in the development of a replacement for the quantile called the superquantile, also referred to as Conditional Value-at-Risk (CVaR) in the financial engineering literature. The superquantile is a measure of uncertainty similar to the quantile, but with su-

perior mathematical properties. Formally, the superquantile (CVaR) for a continuously distributed X is defined as,

$$\bar{q}_\alpha(X) = E[X | X > q_\alpha(X)] = \frac{1}{1 - \alpha} \int_\alpha^1 q_p(X) dp. \quad (1)$$

Similar to $q_\alpha(X)$, the superquantile can be used to assess the tail of the distribution. The superquantile, though, does not suffer from the difficulties outlined above like POE and the quantile. First, the superquantile accounts for the magnitude of events in the tail, which can be seen by noticing that it is simply a form of tail expectation. Therefore, in situations where a distribution may have a heavy tail, the superquantile accounts for magnitudes of low-probability large-loss tail events while the quantile does not account for this information. Second, it is continuous w.r.t. the parameter α . Third, it is far easier to handle in optimization contexts.

The notion of *buffered probability* was originally introduced by Rockafellar and Royset (2010) in the context of the design and optimization of structures as the Buffered Probability of Failure (bPOF). Working to extend this concept, bPOE was developed as the inverse of the superquantile by Mafusalov and Uryasev (2018) in the same way that POE is the inverse of the quantile. Specifically, for continuously distributed X , bPOE at threshold z is defined in the following way, where $\sup X$ denotes the essential supremum of random variable X and threshold $z \in [E[X], \sup X]$.

$$\bar{p}_z(X) = \{1 - \alpha | \bar{q}_\alpha(X) = z\}. \quad (2)$$

In words, bPOE calculates one minus the probability level at which the superquantile, the tail expectation, equals the threshold z . Roughly speaking, bPOE calculates the proportion of worst-case outcomes which average to z .¹ Figure 1 presents an illustration of bPOE for a Lognormal distributed

¹Empirical \bar{p}_z and \bar{q}_α can be easily calculated by sorting data and calculating tail averages. When q_α or z lie in-between sample ‘‘atoms’’, more care is needed to get exact results which include a weighted average of the left and right atom in the overall tail average. (see e.g. Mafusalov and Uryasev (2018) or Pavlikov and Uryasev (2014) for details.)




Figure 1: Shown is the Probability Density Function (PDF) of $X \sim \text{Lognormal}(\sigma = 1, \mu = 0)$. Given threshold $z \in \mathbb{R}$, POE equals $P(X > z)$ the cumulative density in red. For the same threshold z , bPOE equals $\bar{p}_z(X)$ the combined cumulative density in red and blue. By definition, the expectation of the worst-case $1 - \alpha = \bar{p}_z(X)$ outcomes equals $z = \bar{q}_\alpha(X)$. These worst-case outcomes are those that are larger than the quantile $q_\alpha(X)$.

random variable X . We note that there exist two slightly different variants of bPOE, called Upper and Lower bPOE which are identical for continuous random variables. For the interested reader, details regarding the difference between Upper and Lower bPOE can be found in Mafusalov and Uryasev (2018).

Similar to the superquantile, bPOE is a more robust measure of tail risk, as it considers not only the probability that events/losses will exceed the threshold z , but also the magnitude of these potential events. In addition, it also considers the magnitude of the largest events/losses that are less than, z . Thus, it is more robust to slight variations in z , particularly if there is a lot of probability mass right below or above z .

Also, much like the superquantile, bPOE can be represented as the unique minimum of a one-dimensional convex optimization problem with the formulas given as follows, where $[\cdot]^+ = \max\{\cdot, 0\}$.

$$\bar{p}_z(X) = \min_{a \geq 0} E[a(X - z) + 1]^+, \quad (3)$$

$$\bar{q}_\alpha(X) = \min_{\gamma} \gamma + \frac{E[X - \gamma]^+}{1 - \alpha}. \quad (4)$$

Note that these formulas are valid for general real valued random variables, not only continuously distributed random variables. Although we do not address it in this paper, these formulas allow for easy optimization of bPOE or the superquantile. We refer readers to Norton and Uryasev (2016); Mafusalov and Uryasev (2018); Rockafellar and Royset (2010) for optimization examples.

The bPOE concept is also closely related to the concept of a superdistribution function $\bar{F}(z)$, introduced by Rockafellar and Royset (2014). For the CDF, we have that POE equals $P(X > z) = 1 - F(z)$ and we have the inverse CDF given by $F^{-1}(\alpha) = q_\alpha(X)$. The superdistribution function $\bar{F}(z)$ is motivated by the inverse relation $\bar{F}^{-1}(\alpha) = \bar{q}_\alpha(X)$. Thus, bPOE equals $1 - \bar{F}(z)$. The superdistribution function of a random variable X can also be understood as the CDF of an auxiliary random variable $\bar{X} = \bar{q}_u(X)$ where $u \sim U(0, 1)$ is a uniformly distributed random variable. In this case, $\bar{F}_{\bar{X}}(z) = F_{\bar{X}}(z)$ where the subscript indicates that it is the distribution function associated with a particular random variable.

2. MEASURING RISK OF EXCEEDANCE EVENTS

Risk in engineering applications is often measured with POE. For example, the risk of monetary losses associated with seismic activity in an urban area is commonly characterized by the Complementary CDF (CCDF), or the probability that monetary losses will exceed various threshold levels. For example, Mahsuli (2012) utilize simulated seismic scenarios to estimate the probability that losses from seismic activity over time will exceed various threshold levels. Assessing the risk of a tropical storm is another example given by Davis and Uryasev (2016), where insurers and disaster relief organizations alike are concerned with the probability that tropical storm damage will exceed various thresholds when making landfall.

The accuracy and robustness of these measurements is of critical importance, particularly given the downstream effects they have upon decision

making and planning processes. For example, consideration of such statistics is obviously important for organizations like the Federal Emergency Management Agency (FEMA). Projecting budgetary needs given a seismic or tropical storm event involves consideration of the likelihood that damages will exceed particular limits. Under-estimation of such requirements can lead to slow and inadequate relief efforts, furthering the long-term effects of the event. Hurricane Maria of 2017 presents an example of this, with the 2017 Hurricane Season FEMA After-Action Report citing budgetary and resource shortfalls as a hinderance to a more effective response. Donald Trump himself said that the disaster “threw our budget a little out of whack.” (see Robles (2018))

Underestimation of the magnitude of such events can potentially be traced to two causes: 1) Event magnitudes which follow a heavy-tailed distribution. 2) Inadequate characterization of the risk associated with the frequency *and* magnitude of such tail-events. The heavy tailed nature of loss/damage distributions has been observed in a multitude of engineering applications. For example, in seismic risk assessment of structures, the probability of exceeding limiting values of structural demands, damage states, or losses is often modeled with heavy-tailed log-normal distributions or, when applied to modeling the tail, distributions from Extreme Value Theory (EVT) like the Weibull are used. See e.g., Qin et al. (2015). The focus, however, on measuring POE results in assessment of exceedance risk which can be misleading, particularly since it ignores the magnitude of events beyond the threshold and only counts their frequency.

The new characteristic called bPOE serves to *simultaneously* characterize the frequency and magnitude of events exceeding a specified threshold. This leads to risk-averse assessments of exceedance risk that can be modeled as a risk-averse variant of the CCDF.

To illustrate, we present an example from Mahsuli (2012) which simulates potential seismic losses over a 50 year time span, in billion \$CAD, for the Vancouver metropolitan area. Specifically, both the seismic activity emanating from the Cascade

Mountains crustal source over a 50 year timespan and the associated accumulated losses are simulated 100,000 times. As emphasized by Mahsuli (2012), a common method to assess seismic risk is to use the simulated data to estimate the probability that losses will exceed certain large-loss thresholds.

Figure 2 shows the 10% tail of the empirical POE curve of the simulated losses. The tail is indeed long, but the cumulative density which shrinks rapidly to zero could lead one to believe that the probability of exceedance for large thresholds z is insignificant. For example, consider the threshold $z = \$20$ billion. The probability that losses L exceed z is $P(L > z) = .038$.

Using this as an indication of the risk of seismic losses, however, is highly misleading. While the density in the tail beyond $z = \$20$ billion is small, the values within this tail are dramatically larger than \$20 billion. For example, $q_{.99}(L) = \$72.3$ billion, $q_{.995}(L) = \$109$ billion, and $q_1(L) = \$373$ billion. These events are so severe, their occurrence would lead to dramatic consequences for the Vancouver metropolitan area and its 50 year development if not properly accounted for and addressed ahead of time during risk assessment.

Utilizing bPOE, (2), as our assessment metric, however, we see a dramatic difference emerge in the perceived risk of exceedance events. Figure 2 displays the empirical bPOE curve² of the same simulated loss data, along with a curve representing the difference between bPOE and POE at different threshold levels. First, consider again the threshold $z = \$20$ billion. Comparing bPOE and POE, we have that $P(L > z) = .038$ while bPOE yields $\bar{p}_z(L) = .152$. bPOE is almost 4 times larger than POE, reflecting the risk of extremely large events in the tail exceeding z .

3. PARAMETRIC DISTRIBUTION FITTING

These concepts can not only be used to measure the risk of exceedance events, but can also be plugged into standard statistical estimation procedures as replacements for POE (or the CDF) and the quantile when heavy tails are a factor and exceedance risk is of central importance. In particular, we can

²solving (2) for many values of z .




Figure 2: Shown is the empirical Probability of Exceedance (POE) and Buffered Probability of Exceedance (bPOE) for high loss thresholds. Damage data from Mahsuli (2012).

consider parametric density estimation, which is often used to fit heavy tailed data to some parametric distribution allowing for further analysis of exceedance risk and its sensitivity the choice to threshold. For example, Qin et al. (2015) use a fusion technique to fit parametric models of the demand placed upon a structure by seismic activity. In this context, critical emphasis is placed upon accurate estimation of exceedance risk, specifically the probability that demands will exceed limiting values that represent danger zones for structural integrity. As seen in Qin et al. (2015), it is common to use heavy tailed distributions like the lognormal or weibull to parametrically fit a distribution to data. In particular, common tools include maximum likelihood (ML), method of moments (MM), and the Matching of Quantile's (MOQ) procedure which is much like MM but where parametric and empirical quantiles at various alpha levels $\alpha_1, \dots, \alpha_k$ are used instead of the typical parametric and empirical moments (see e.g., Sgouropoulos et al. (2015); Karian and Dudewicz (1999)). An important note to keep in mind is that MOQ is effectively equivalent to an identical procedure but where parametric POE (i.e. the CCDF) is fit to match empirical POE at various threshold levels z_1, \dots, z_k . This can be easily realized by noting the fact that $q_\alpha(X) = x \iff P(X > x) = 1 - \alpha$.

The ML and MM procedures, however, do not place direct emphasis on fitting the tail and can thus be poor choices when accurate and risk-averse estimation of exceedance probabilities is of the utmost important. Furthermore, while the MOQ procedure does put direct focus on making sure that parametric and empirical exceedance probabilities match, it is nevertheless considering POE as its primary risk assessment, which as already emphasized, ignores important aspects of the magnitude and frequency of exceedance events.

This, then, provides great opportunity to utilize bPOE and the related superquantile as criteria for distribution fitting. As we will show, as long as formulas for bPOE or the superquantile are available in closed-form³ for the considered parametric family, parametric procedures can be easily adapted to find parametric distributions with best matching bPOE or, equivalently, superquantiles. We illustrate this idea by using a simple variation of MOQ, which we call the Method of Superquantile's (MOS), where superquantile's at varying levels of α take the place of moments. Note that, just as with MOQ, the MOS procedure is equivalent to fitting a bPOE criteria, which can be understood similarly by noting that $\bar{q}_\alpha(X) = z \iff \bar{p}_z(X) = 1 - \alpha$.

Our numerical example utilizes a heavy tailed Weibull to fit the loss data from Section 2, since it is particularly well-suited for asymmetric heavy-tailed data. In addition, we also introduce a new method for visually assessing goodness-of-fit of this method. Instead of utilizing a Q-Q plot, comparing empirical quantiles against predicted quantiles, we use an SQ-SQ plot, which compares the empirical superquantiles with superquantiles predicted by our parametric distribution. Overall, we find that by utilizing superquantiles, or equivalently bPOE, for distribution fitting, we arrive at a parametric distribution that provides a more accurate and conservative estimate of tail behavior for further assessment of exceedance risk.

3.1. Method of Superquantile's

The MM is a well known tool for estimating the parameters of a distribution when moments are avail-

³Formulas for multiple distributions can be found in Norton et al. (2018).

able in parametric form and desired moments are either assumed to be known or are measured from empirical observations. It looks for the distribution $f_{\Theta}(x)$, parameterized by Θ , with moments equal to some known moments or, if moments are unknown, empirical moments. With n moments used, the problem reduces to solving a set of n equations w.r.t. the set of parameters Θ of the distribution family.

This method, though, can be generalized where moments are replaced by other distributional characteristics, such as the superquantile and quantile. We utilize superquantile's in this context. This method provides flexibility through the choices of different α , allowing the user to focus the fitting procedure on particular portions of the distribution. This flexibility is advantageous compared to other methods such as MM or ML since these fitting methods treat each portion of the distribution equally. When fitting the tail is important, for example, and there are many samples around the mean with few samples in the tail, it can be desirable to focus the fitting procedure on carefully fitting the tail samples. As will be shown, one can focus MOS by choice of α .

We formulate the following problem, where $\hat{q}_{\alpha}(X)$ denotes either a known superquantile or an empirical estimate from a sample and $\bar{q}_{\alpha}(X_{f_{\Theta}})$ denotes parameterized formulas for the superquantile when X has density function f_{Θ} with the set of parameters Θ :

Method of Superquantiles (MOS): Fix $\alpha_1, \dots, \alpha_k \in [0, 1]$ and choose a parametric distribution family f_{Θ} with parameters Θ . Solve for Θ such that,

$$\bar{q}_{\alpha_i}(X_{f_{\Theta}}) = \hat{q}_{\alpha_i}(X) \text{ for all } i = 1, \dots, k,$$

which is a system of k equations in $|\Theta|$ unknowns.

This problem, however, may not have a solution. For example, if $k = 2$ and the parametric family only has a single parameter (i.e. $|\Theta| = 1$). In this case, one can solve the following surrogate Least Squares minimization problem:

LS Method of Superquantiles (LS-MOS): Fix

$\alpha_1, \dots, \alpha_k \in [0, 1]$ and choose a parametric distribution family f_{Θ} with parameters Θ . Choose weights $c_1, \dots, c_k > 0$ and solve for,

$$\Theta \in \operatorname{argmin}_{\Theta} \sum_i c_i (\bar{q}_{\alpha_i}(X_{f_{\Theta}}) - \hat{q}_{\alpha_i}(X))^2 .$$

This procedure finds the distribution which has superquantile's that are *close* to the empirical superquantile's. The freedom to select α_i as well as c_i provides the user with much flexibility as to which portion of the distribution should match more exactly the empirical superquantile's.

Notice that we could have fit bPOE at various threshold levels z_1, \dots, z_k instead of superquantiles by solving,

$$\Theta \in \operatorname{argmin}_{\Theta} \sum_i c_i (\bar{p}_{z_i}(X_{f_{\Theta}}) - \hat{p}_{z_i}(X))^2 .$$

where $\hat{p}_z(X)$ denotes empirical bPOE and $\bar{p}_z(X_{f_{\Theta}})$ is a parameterized formula for bPOE for the distribution family f_{Θ} . Importantly, however, note that the approaches are effectively equivalent due to the fact that $\bar{q}_{\alpha}(X) = z \iff \bar{p}_z(X) = 1 - \alpha$. For our example, we use superquantiles since a closed-form equation for the superquantile of a weibull is known and a closed-form equation for bPOE is not.

3.1.1. Example: Fitting Seismic Loss Data

To illustrate the approach, we fit a Weibull distribution to the seismic loss data given before from Mahsuli (2012). Note that the superquantile of $X \sim Weibull(\lambda, k)$ is provided in closed form by Norton et al. (2018) as,

$$\bar{q}_{\alpha}(X) = \frac{\lambda}{1 - \alpha} \Gamma_U \left(1 + \frac{1}{k}, -\ln(1 - \alpha) \right), \quad (5)$$

where $\Gamma_U(a, b) = \int_b^{\infty} p^{a-1} e^{-p} dp$ is the upper incomplete gamma function.

We estimated the weibull parameters using MM, ML, and the LS-MOS. The MM was solved using the first two moments. The LS-MOS was solved twice. It was first solved with $\alpha_1 = .15, \alpha_2 = .75, c_1 = c_2 = 1$. Then, to put more emphasis on the tail observations and form a more conservative estimate of the tail and associated exceedances, it was also solved with $\alpha_1 = .5, \alpha_2 = .75, \alpha_3 = .95, c_1 =$




Figure 3: Shown is a plot of pairs $(\bar{q}_\alpha(X_E), \bar{q}_\alpha(X_{f_\Theta}))$ for every $\alpha \in \{.01, .02, \dots, .99\}$ and every $f_\Theta = ML, MM, LS1, E$ where f_Θ indicates the method used to fit parameters and $\bar{q}_\alpha(X_E)$ is the empirical superquantile.

$c_2 = c_3 = 1$. We denote these solutions as LS1, LS2 respectively. The ML solution is available in closed-form and we solved MM, LS1, and LS2 using Scipy's optimization library.⁴

Looking at Figure 3, we have a visual assessment of the goodness-of-fit of each distribution given by the SQ-SQ plot, where we have plotted the superquantiles of each distribution paired with the superquantiles of the empirical distribution. This is the same as a Q-Q, quantile-quantile, plot but with superquantiles. We see first that the distribution produced by ML, as predicted, is not a good fit and severely underestimates the superquantiles. The MM distribution is better, but we see that the distribution produced by LS1 and LS2 are clearly the best. Additionally, for large α levels (values in the upper-right corner), we see that LS2 provides a conservative estimate of the superquantile. This can be a desirable property of this method, particularly since random samples often under-sample high-magnitude/low-probability events in the tail.

Figure 4 shows the bPOE curve of the data, and the distributions produced by MM and LS2. LS1 and ML are omitted since ML is very poor and the




Figure 4: Shown is the tail portion of the bPOE curve for MM and LS2.




Figure 5: Shown is the POE curve, with x-axis on log-scale to highlight differences.

bPOE curve from LS1 is indistinguishable from the one given by LS2. We see from this that LS2 provides a very tight fit of the empirical bPOE curve with MM consistently underestimating bPOE. Figure 5 compares the tail of the POE curves, with the x-axis on a log-scale to help highlight the differences between curves. As can be seen on the POE plot, LS1 and LS2 again provide the best fit. The distribution produced by MM consistently overestimates POE below the $\alpha = .95$ quantile. And, as before, we see that ML provides a poor fit.

⁴Specifically, we used the *leastsq* function which implements MINPACK's *lmdif* routine. This routine requires function values and calculates the Jacobian by a forward-difference approximation.

4. CONCLUSION

Assessing the risk of exceedance events with probability of exceedance, while intuitive, can often hide

important information about the magnitude of exceedance events beyond the critical threshold and near-exceedance events below the threshold. Using, instead, bPOE to assess the risk of exceedance events provides a simultaneous assessment of the frequency and magnitude of exceedance events, as well as the frequency and magnitude of near-exceedance events. Ignoring these aspects of the tail can lead to false confidence in a low-risk assessment of large losses, damages, or failures in system risk analysis. We illustrated a use of bPOE for measuring the risk of exceedance events on seismic loss data, showing that bPOE provides a much more conservative assessment of exceedance risk compared to POE which is well justified by the potential impact of events that, while being low-probability, have magnitude that is catastrophically large.

With bPOE serving as a measure of tail risk, one would naturally like to incorporate it into distribution fitting procedures, particularly when the primary goal of the fitting is to accurately estimate the tail and subsequently assess the risk of exceedance events. In this direction, we showed how one can incorporate bPOE into a simple parametric distribution fitting procedure via the use of superquantiles. Using a closed-form expression for the superquantile of a weibull distribution, we were able to perform a method of moments style fitting procedure, where we utilized superquantiles instead of moments. We were also able to easily adapt the simple Q-Q plot goodness-of-fit test to create a SQ-SQ plot. This served as a natural way to assess and compare the fits of competing methods where the fitness criteria is based upon bPOE and superquantiles.

5. REFERENCES

- Davis, J. R. and Uryasev, S. (2016). "Analysis of tropical storm damage using buffered probability of exceedance." *Natural Hazards*, 83(1), 465–483.
- Karian, Z. A. and Dudewicz, E. J. (1999). "Fitting the generalized lambda distribution to data: a method based on percentiles." *Communications in Statistics-Simulation and Computation*, 28(3), 793–819.
- Mafusalov, A. and Uryasev, S. (2018). "Buffered probability of exceedance: Mathematical properties and optimization." *SIAM Journal on Optimization*, 28(2), 1077–1103.
- Mahsuli, M. (2012). "Probabilistic models, methods, and software for evaluating risk to civil infrastructure." Ph.D. thesis, University of British Columbia, University of British Columbia.
- Norton, M., Khokhlov, V., and Uryasev, S. (2018). "Calculating cvar and bpoe for common probability distributions with application to portfolio optimization and density estimation." *arXiv:1811.11301*.
- Norton, M. and Uryasev, S. (2016). "Maximization of auc and buffered auc in binary classification." *Mathematical Programming*, 1–38.
- Pavlikov, K. and Uryasev, S. (2014). "Cvar norm and applications in optimization." *Optimization Letters*, 8(7), 1999–2020.
- Qin, J., Mackie, K. R., and Stojadinovic, B. (2015). "Improved probability distribution models for seismic fragility assessment." *Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12)*, University of British Columbia.
- Robles, F. (2018). "Ema was sorely unprepared for puerto rico hurricane, report says.." *The New York Times*.
- Rockafellar, R. and Royset, J. (2010). "On buffered failure probability in design and optimization of structures." *Reliability Engineering & System Safety*, Vol. 95, 499-510.
- Rockafellar, R. and Uryasev, S. (2000). "Optimization of conditional value-at-risk." *The Journal of Risk*, Vol. 2, No. 3, 2000, 21-41.
- Rockafellar, R. T. and Royset, J. O. (2014). "Random variables, monotone relations, and convex analysis." *Mathematical Programming*, 148(1-2), 297–331.
- Rockafellar, R. T. and Uryasev, S. (2002). "Conditional value-at-risk for general loss distributions." *Journal of banking & finance*, 26(7), 1443–1471.
- Sgouropoulos, N., Yao, Q., and Yastremiz, C. (2015). "Matching a distribution by matching quantiles estimation." *Journal of the American Statistical Association*, 110(510), 742–759.