

# Mitigating optimistic bias in entropic risk estimation and optimization with an application to insurance

U. Sadana, E. Delage, A. Georghiou

G–2024–83

December 2024

---

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

**Citation suggérée** : U. Sadana, E. Delage, A. Georghiou (Décembre 2024). Mitigating optimistic bias in entropic risk estimation and optimization with an application to insurance, Rapport technique, Les Cahiers du GERAD G– 2024–83, GERAD, HEC Montréal, Canada.

**Suggested citation**: U. Sadana, E. Delage, A. Georghiou (December 2024). Mitigating optimistic bias in entropic risk estimation and optimization with an application to insurance, Technical report, Les Cahiers du GERAD G–2024–83, GERAD, HEC Montréal, Canada.

**Avant de citer ce rapport technique**, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2024-83>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

**Before citing this technical report**, please visit our website (<https://www.gerad.ca/en/papers/G-2024-83>) to update your reference data, if it has been published in a scientific journal.

---

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2024  
– Bibliothèque et Archives Canada, 2024

Legal deposit – Bibliothèque et Archives nationales du Québec, 2024  
– Library and Archives Canada, 2024

# Mitigating optimistic bias in entropic risk estimation and optimization with an application to insurance

Utsav Sadana <sup>a</sup>

Erick Delage <sup>b</sup>

Angelos Georghiou <sup>c</sup>

<sup>a</sup> CIRRELT, GERAD & Department of Computer Science and Operations Research, Université de Montréal, Québec, Canada, H3T 1N8

<sup>b</sup> GERAD & Department of Decision Sciences, HEC Montréal, Montréal (Québec), Canada, H3T 2A7

<sup>c</sup> Department of Business and Public Administration, University of Cyprus, CY-1678 Nicosia, Cyprus

utsav.sadana@umontreal.ca

erick.delage@hec.ca

georghiou.angelos@ucy.ac.cy

December 2024  
Les Cahiers du GERAD  
G–2024–83

Copyright © 2024 Sadana, Delage, Georghiou

---

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

**Abstract :** The entropic risk measure is widely used in high-stakes decision making to account for tail risks associated with an uncertain loss. With limited data, the empirical entropic risk estimator, i.e. replacing the expectation in the entropic risk measure with a sample average, underestimates the true risk. To mitigate the bias in the empirical entropic risk estimator, we propose a strongly asymptotically consistent bootstrapping procedure. The first step of the procedure involves fitting a distribution to the data, whereas the second step estimates the bias of the empirical entropic risk estimator using bootstrapping, and corrects for it. Two methods are proposed to fit a Gaussian Mixture Model to the data, a computationally intensive one that fits the distribution of empirical entropic risk, and a simpler one with a component that fits the tail of the empirical distribution. As an application of our approach, we study distributionally robust entropic risk minimization problems with type- $\infty$  Wasserstein ambiguity set and apply our bias correction to debias validation performance. Furthermore, we propose a distributionally robust optimization model for an insurance contract design problem that takes into account the correlations of losses across households. We show that choosing regularization parameters based on the cross validation methods can result in significantly higher out-of-sample risk for the insurer if the bias in validation performance is not corrected for. This improvement in performance can be explained from the observation that our methods suggest a higher (and more accurate) premium to homeowners.

**Keywords :** Entropic risk measure, bias correction, distributionally robust optimization, insurance pricing

---

**Acknowledgements:** The first author's research is supported by the NSERC, Canada, grant RGPIN-2024-05067. This research was also supported by compute resources provided by Calcul Quebec and the Digital Research Alliance of Canada. We thank Tianyu Wang for his helpful comments on an earlier version of the paper.

# 1 Introduction

The purpose of a risk measure is to assign a real number to a random variable, representing the preference of a risk-averse decision maker towards different risky alternatives. For instance, when faced with multiple options, a decision maker might prefer a guaranteed loss of zero over an uncertain option, even if the latter has a strictly negative expected loss. While this behavior can be explained using the mean-variance criterion (Markowitz 1952), which balances the expected loss and its fluctuations around the mean, the entropic risk measure offers greater flexibility by incorporating higher moments of the loss distribution.

The entropic risk measure is useful in high-stakes decision-making, where rare events and their associated extreme losses are a significant concern. A key advantage of using entropic risk in multi-stage decision-making is the time-consistency of the optimal policies. The entropic risk measure is widely used due to its interpretation as the certainty equivalent of the exponential utility function (Von Neumann and Morgenstern 1944), which represents the risk preferences of a decision-maker exhibiting constant absolute risk aversion (CARA – Arrow 1971, Pratt 1964). There has been significant growth in research on exponential utility functions, which appear in the literature under various names, including entropic risk minimization, tilted empirical risk minimization, constant absolute risk aversion, and as special cases of more general shortfall risk measures and optimized certainty equivalent risk measures (Ben-Tal and Teboulle 1986). Applications of these concepts are widespread, particularly in finance (Föllmer and Schied 2002, 2016, Smith and Chapman 2023), portfolio selection (Brandtner et al. 2018, Markowitz 2014, Chen et al. 2024b), revenue management (Lim and Shanthikumar 2007), economics (Svensson and Werner 1993), operations management (Choi and Ruszczyński 2011, Chen and Sim 2024), robotics (Nass et al. 2019), statistics (Li et al. 2023), reinforcement learning (Fei et al. 2021, Hau et al. 2023), risk-sensitive control (Howard and Matheson 1972, Bäuerle and Jaśkiewicz 2024), game theory (Saldi et al. 2020), and catastrophe insurance pricing (Bernard et al. 2020).

Since the seminal work by Föllmer and Schied (2002), which established the axiomatic foundations for convex risk measures, there has been growing interest in quantitative risk management using convex law-invariant risk measures, such as the entropic risk measure. Unlike coherent risk measures, like Conditional Value at Risk (CVaR – Artzner et al. 1999), convex law-invariant risk measures allow for non-linear variation in risk with the size of a position.

To formally define the entropic risk measure, let  $\ell(\boldsymbol{\xi})$  represent the uncertain loss associated with an uncertain parameter  $\boldsymbol{\xi} \in \Xi \subseteq \mathbb{R}^d$ . Then, the entropic risk associated with parameter  $\boldsymbol{\xi}$  is given by:

$$\rho_{\mathbb{P}}^{\alpha}(\ell(\boldsymbol{\xi})) := \begin{cases} \frac{1}{\alpha} \log(\mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\boldsymbol{\xi}))]) & \text{if } \alpha > 0, \\ \mathbb{E}_{\mathbb{P}}[\ell(\boldsymbol{\xi})] & \text{if } \alpha = 0, \end{cases} \quad (1)$$

where the loss  $\ell(\boldsymbol{\xi})$  is transformed by the increasing and convex exponential function, and  $\alpha$  is the risk aversion parameter. This formulation expresses the entropic risk as the certainty equivalent of the expected disutility  $\mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\boldsymbol{\xi}))]$ , reflecting the monetary value of the risk inherent in the uncertain outcome  $\ell(\boldsymbol{\xi})$ . By adjusting the risk-aversion parameter  $\alpha$ , also known as the Arrow-Pratt measure of risk aversion, the decision maker's sensitivity to extreme losses can be controlled. For the remainder of the paper, we simplify the notation by suppressing the dependence on  $\alpha$  where it is not essential, and denote entropic risk as  $\rho_{\mathbb{P}}$ .

In real-world applications, the distribution  $\mathbb{P}$  of the random variable  $\boldsymbol{\xi}$  is unknown, and decisions are often made using historical realizations of random variable  $\boldsymbol{\xi}$  that are assumed to be independent and identically distributed (i.i.d.) with distribution  $\mathbb{P}$ . Let the data set of  $N$  historical observations be denoted by  $\mathcal{D}_N = \{\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \dots, \boldsymbol{\xi}_N\}$ . A common approach to estimate the entropic risk is to replace the true distribution with the empirical distribution defined as  $\hat{\mathbb{P}}_N(\boldsymbol{\xi}) := \frac{1}{N} \sum_{i=1}^N \delta_{\boldsymbol{\xi}_i}(\boldsymbol{\xi})$ , where  $\delta_{\boldsymbol{\xi}}$  is a

Dirac distribution at the point  $\xi$ . The empirical entropic risk measure is then given by:

$$\rho_{\hat{\mathbb{P}}_N}(\ell(\xi)) := \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha \ell(\hat{\xi}_i)) \right). \tag{2}$$

Since the logarithm function is strongly concave, Jensen’s inequality implies that the empirical entropic risk strictly underestimates the true entropic risk:

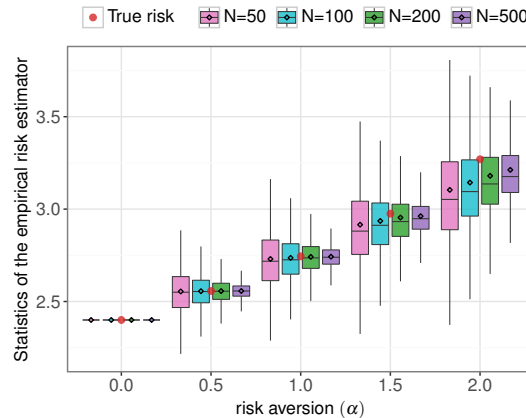
$$\mathbb{E}[\rho_{\hat{\mathbb{P}}_N}(\ell(\xi))] = \mathbb{E} \left[ \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha \ell(\hat{\xi}_i)) \right) \right] < \frac{1}{\alpha} \log \left( \mathbb{E} \left[ \frac{1}{N} \sum_{i=1}^N \exp(\alpha \ell(\hat{\xi}_i)) \right] \right) = \rho(\ell(\xi)), \tag{3}$$

unless  $\mathbb{E}[\ell(\xi)] = \ell(\xi)$  almost surely, and where the expectation is taken with respect to the randomness of the data  $\mathcal{D}_N$ .

A fundamental challenge in high-stakes decision-making lies in accurately estimating risk. Even with a large number of samples, the empirical entropic risk can significantly underestimate the true risk, especially for decision makers with heightened risk sensitivity. This challenge is demonstrated in the following example.

**Example 1.** In an insurance pricing problem, the insurer aims to determine the minimum premium  $\pi$  at which they can insure against the loss  $\ell(\xi) := \xi$ . Let risk aversion parameter of insurer be  $\alpha$ . Assuming full coverage for the losses  $\xi$ , loss of the insurer if they charge a premium  $\pi$  is a random variable given by  $\pi - \xi$ . Thus, minimum premium at which the insurer insures the risk should be such that the entropic risk of the insurer from insuring is at most equal to 0, i.e.,  $\frac{1}{\alpha} \log (\mathbb{E}_{\mathbb{P}}[\exp(\alpha(\pi - \xi))]) = 0$ . On rearranging the terms, one can show that the minimum premium equals the entropic risk associated with the loss  $\xi$ ,  $\pi = \frac{1}{\alpha} \log (\mathbb{E}_{\mathbb{P}}[\exp(\alpha\xi)])$ , also called the exponential premium (Gerber 1974). Suppose that the loss follows a Gamma distribution  $\Gamma(\kappa, \lambda)$  (see Fu and Moncher 2004, Bernard et al. 2020) with shape parameter  $\kappa$  and scale parameter  $\lambda$ . The moment-generating function of  $\Gamma$ -distributed random variable is known in closed form which allows us to analytically compute the optimal premium  $\pi = \frac{1}{\alpha} \log \left( (1 - \lambda\alpha)^{-\kappa} \right)$  if  $\lambda < 1/\alpha$ .

Suppose an insurer has access to  $N \in \{50, 100, 200, 500\}$  samples of the losses which are generated from a  $\Gamma(10, 0.24)$ -distribution. We use empirical distribution  $\hat{\mathbb{P}}_N$  over  $N$  samples to estimate the entropic risk for different risk aversion parameters,  $\alpha \in \{0, 0.5, 1, 1.5, 2\}$ . Figure 1 presents statistics of the distribution of the empirical risk estimator as a function of  $N$  and  $\alpha$ . We can see that more risk-averse insurers significantly underestimate the risk of the loss  $\xi$  and thus the premium to impose on the insuree even for the relatively large sample size  $N = 500$ . This phenomenon can also be



**Figure 1: Statistics of the empirical risk for different values of the risk aversion parameters  $\alpha \in \{0, 0.5, 1, 1.5, 2\}$  and training sample sizes  $N \in \{50, 100, 200, 500\}$  over 10000 repetitions. The true risk is given by  $\frac{1}{\alpha} \log \left( (1 - 0.24\alpha)^{-10} \right)$ .**

understood by analyzing the impact of an infinitesimal change in the distribution on the entropic risk, as quantified by the influence function (Hampel 1974). For this example, the influence function is given by  $\text{IF}(\hat{\xi}) = -\frac{1}{\alpha} + \frac{\exp(\alpha\hat{\xi})}{\alpha(1-0.24\alpha)^{-10}}$ . Consequently, missing tail scenarios can disproportionately affect the estimation of entropic risk for high values of  $\alpha$  due to the exponential term. Detailed calculations are provided in Appendix B.1, and Figure 14 offers a visual representation of the influence function across samples and their associated probabilities. The slow convergence of sample mean to the true mean for heavy tailed random variables is a well-known phenomenon (Catoni 2012, Lugosi and Mendelson 2019). Even when the underlying distribution of  $\ell(\boldsymbol{\xi})$  is light-tailed,  $\exp(\alpha\ell(\boldsymbol{\xi}))$  can be heavy-tailed (Nair et al. 2022). For instance,  $\exp(\alpha\ell(\boldsymbol{\xi}))$  is a lognormally distributed when  $\ell(\boldsymbol{\xi})$  is Gaussian random variable.

The contribution of this paper is to propose a scheme to produce (approximately) median-unbiased estimators of entropic risk to mitigate the bias of the empirical entropic risk estimator. We propose a bias correction term  $\delta(\mathcal{D}_N)$  that employs bootstrapping, using a distribution fitted to the data  $\mathcal{D}_N$ , to get

$$\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) \approx \text{median} \left[ \rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) + \delta(\mathcal{D}_N) \right], \quad (4)$$

where the median statistic is taken with respect to the randomness in  $\mathcal{D}_N$ . We establish mild conditions under which  $\delta(\mathcal{D}_N) \rightarrow 0$  almost surely as  $N \rightarrow \infty$ . In particular, basing the bootstrap on a constrained maximum likelihood estimation (MLE) of sampling distribution will yield an asymptotically consistent estimator. Unfortunately, our empirical experiments establish that a bootstrap correction based on MLE fails to adequately address the underestimation of entropic risk. Instead, we provide two procedures to fit “bias-aware” distributions that take into account the entropic risk estimation bias caused by tail events. The first one involves a distribution matching technique that tries to fit the entropic risk estimator’s distribution itself, and the second uses a simple mixture distribution with a component dedicated to fitting the tail of the empirical distribution.

Going beyond the estimation of entropic risk, we study the entropic risk minimization problem. Solving the sample average approximation (SAA) of the entropic risk minimization problem is known to produce a second source of bias, also known as the optimizer’s curse (Smith and Winkler 2006). Distributionally robust optimization (DRO) is widely used to address the optimistic bias of SAA policies because the decision-maker is protected against perturbations in the empirical distribution that lie in a distributional ambiguity set. Most of the literature on DRO with the Wasserstein ambiguity set assumes that the random variables involved in the expectation operator have light tails, a condition that is not satisfied for the entropic risk measure. It is well-known that worst-case loss in a DRO problem with type- $p$  Wasserstein ambiguity set ( $p < \infty$ ) is finite if and only if the loss function satisfies a growth condition (Gao and Kleywegt 2023). Thus, the worst-case entropic risk can be shown to be unbounded for type- $p$  Wasserstein ambiguity set with  $p \in [1, \infty)$ . So, we introduce a distributionally robust entropic risk minimization problem with type- $\infty$  Wasserstein ambiguity set and cast it into a finite-dimensional convex optimization problem for piecewise concave loss functions. To tune the radius of the ambiguity set, a typical approach is to use K-Fold cross validation (CV). We use our bias mitigation procedure to estimate the validation performance of the resulting decisions.

To demonstrate the effectiveness of our approach, we conduct a case study on an insurance pricing problem. The International Panel on Climate Change (IPCC) advocates using financial instruments like catastrophe insurance to mitigate risks associated with rare, high-impact events such as floods, earthquakes, and wildfires. These events have become more frequent due to climate change, making reliable risk estimation even more critical (Linnerooth-Bayer and Hochrainer-Stigler 2015). The correlated risks of floods, earthquakes, and wildfires can lead to significant payouts for insurers, limiting the availability of insurance for such events in commercial markets (Marcoux and H Wagner 2023). The true distribution of losses, which may be correlated across homeowners, is unknown to both the insurer and the homeowners. The insurer addresses this uncertainty by solving a distributionally robust insurance pricing problem, determining the coverage to offer and the premium to charge each homeowner.

Our results show that the insurer can achieve a significant improvement in out-of-sample entropic risk compared to the “traditional” K-Fold CV procedure, which selects the radius of the distributional ambiguity set solely by evaluating decisions on validation loss scenarios.

Our contributions can be described as follows:

1. On the theoretical side, we propose a strongly asymptotically consistent bootstrapping procedure to debias the empirical entropic risk estimator. Our main contribution lies in developing two bias-correction methods to mitigate the underestimation of the entropic risk in the finite sample case. In the first step, both methods fit a distribution to the samples to capture the bias in the samples. In the second step, bootstrapping is used to estimate the bias. Our methods could be of independent interest for debiasing more general risk measures.
2. We introduce a distributionally robust entropic risk minimization problem with a type- $\infty$  Wasserstein ambiguity set with bounded worst-case losses. We obtain its tractable robust counterpart for piecewise concave loss functions using Fenchel duality and provide conditions under which the optimal risk from the the DRO problem converges to the true optimal risk.
3. On the application side, our work contributes toward data-driven designing of insurance premium pricing and coverage policies. To the best of our knowledge, this is the first time that a distributionally robust version of the well-known risk-averse insurance pricing problem (Bernard et al. 2020) is introduced in the literature. Our model takes into account the different risk aversion attitudes of the insurer and homeowners as well as the systemic risk associated with catastrophe events.

The paper is organized as follows. Section 2 surveys the literature on three related topics, estimating risk measures, correcting optimistic bias associated with solving SAA problem, and catastrophe insurance pricing. Section 3 discusses the properties of entropic risk measure. Section 4 provides a bias correction procedure to mitigate the underestimation problem. In Section 5, we study the entropic risk minimization problem using the DRO framework. In Section 6, we introduce the distributionally robust insurance pricing problem and provide numerical results. Finally, conclusions are given in Section 7.

**Notations:**  $[m]$  denotes the set of integers  $\{1, 2, \dots, m\}$ .  $\|\cdot\|_*$  denotes the dual norm of  $\|\cdot\|$ .  $\delta_\xi$  is the Dirac distribution at the point  $\xi$ . Finally,  $\log(\cdot)$  refers to the natural logarithm.

## 2 Literature review

### 2.1 Risk estimation

Quantitative risk measurement often relies on precise estimation of risk measures commonly used in finance and actuarial science (McNeil et al. 2005). Calculating risk for multidimensional random variables can be challenging due to the need for complex integrals, often approximated using Monte Carlo simulation. When the underlying distribution isn’t directly accessible and only limited samples are available, Monte Carlo-based risk estimators tend to underestimate the actual risk. Kim and Hardy (2007) addressed this by using bootstrapping to correct bias in Value at Risk (VaR) and Conditional Tail Expectation (CTE) estimates, with Kim (2010) extending this approach to general distortion risk measures. Bootstrapping involves sampling with replacement from the empirical distribution, calculating the statistic, and averaging the outcomes over multiple iterations. However, bootstrap estimates still tend to underestimate risk in finite samples due to the lack of extreme tail scenarios (see Figure 4). In contrast, our bootstrapping procedure draws samples from a “bias-aware” fitted distribution that better accounts for tail scenarios.

Several approaches have been proposed in the literature for estimating tail risks. Lam and Mottet (2017) introduce a distributionally robust optimization based method to construct worst-case bounds

on tail risk, assuming the density function is convex beyond a certain threshold. Extreme Value Theory (EVT) is commonly used to estimate tail risk measures, such as CVaR, by fitting a Generalized Pareto Distribution to values exceeding a threshold. Troop et al. (2021) develop an asymptotically unbiased CVaR estimator by correcting the bias in the estimates obtained via maximum likelihood. In Kuiper et al. (2024), the authors derive DRO-based estimators for EVT statistics to account for model misspecification due to scenarios outside the asymptotic tails. While these methods focus on upper-tail risk, they are not directly applicable to the entropic risk measure. Instead of fitting an extreme value distribution, we utilize a parametric two-component Gaussian Mixture Model (GMM), which provides a closed-form expression for the entropic risk and captures both the mean and the tails of the data.

One related field of research is to derive concentration bounds on the risks estimates depending on whether the random variable is sub-gaussian, sub-exponential or heavy-tailed. For optimized certainty equivalent risk measures that are Hölder continuous, L.A. and Bhat (2022) link estimation error to the Wasserstein distance between the empirical and true distributions, for which concentration bounds are available. While the Central Limit Theorem (CLT) ensures asymptotic convergence of the sample average to the true mean, this guarantee doesn't always hold for finite samples unless the tails are Gaussian or sub-Gaussian (Catoni 2012, Bartl and Mendelson 2022). Robust statistics literature offers alternative estimators, like the median-of-means (MoM) estimator (Lugosi and Mendelson 2019), which ensure the estimator is close to the true mean with high confidence. However, these approaches differ from our focus, which is on constructing estimators with minimal bias.

## 2.2 Correcting optimistic bias

Our work on entropic risk minimization relates to correcting the optimistic bias of SAA policies to achieve true decision performance (Smith and Winkler 2006, Beirami et al. 2017). SAA is analogous to empirical risk minimization in machine learning, where the goal is to minimize empirical risk. Methods like DRO, hold-out, and K-fold CV are used to correct this bias. These approaches involve partitioning data into training, validation, and test sets, and then selecting the hyperparameter that results in smallest validation risk (Bousquet and Elisseeff 2000). Our hyper-parameter selection employs the debiased validation risk.

Several methods have been proposed to correct the bias of SAA in linear optimization problems under the assumption that the true data distribution is Gaussian. For example, Ito et al. (2018) introduce a perturbation technique that generates parameters around the true values under Gaussian error assumptions to achieve an asymptotically unbiased estimator of the true loss. Similarly, Gupta et al. (2024) derive estimators for the out-of-sample performance of in-sample optimal policies under a Gaussian distribution and offer extensions to approximate Gaussian cases, leveraging the structure of linear optimization problems. However, extending these methods to our nonlinear problem is challenging. Moreover, our objective is to find optimal policies with low out-of-sample risk rather than merely estimating this risk. Since CV risk could be biased estimate of the out-of-sample risk, we employ our bias correction procedure to correct it, enabling appropriate calibration of the regularization parameter. Importantly, our approach does not rely on Gaussian assumptions for the uncertain parameter or the structure of the objective function.

In Siegel and Wagner (2023), the authors analytically characterize the bias in SAA policies for a data-driven newsvendor problem, providing an asymptotically debiased profit estimator by leveraging the asymptotic properties of order statistics. Iyengar et al. (2023) introduce an Optimizer's Information Criterion (OIC) to correct bias in SAA policies, generalizing the approach by Siegel and Wagner (2023) to other loss functions. However, OIC requires access to the gradient, Hessian, and influence function of the decision rule, which can be challenging to obtain in general constrained optimization problems. Moreover, the form of optimal policy is known for risk-neutral newsvendor problems but not for entropic risk minimization problems.



### 2.3 Insurance pricing

The design of insurance contracts has been widely studied since the foundational work of Arrow (Arrow 1963, 1971). Under the assumption that premiums are proportional to the policy's actuarial value, it has been shown that an expected utility-maximizing policyholder will choose full coverage above a deductible. Various extensions of Arrow's model have been proposed to account for the risk aversion of both the insured and the insurer, using criteria such as mean-variance (Kaluszka 2004a,b), Value at Risk (VaR), and Tail VaR (Cai et al. 2008). Bernard and Tian (2010) incorporate regulatory constraints on the insurer's insolvency risk through VaR. Cheung et al. (2014) extend these models to multiple policyholders with fully dependent risks (comonotonicity), where the insurer utilizes convex law-invariant risk measures. Bernard et al. (2020) further explore different levels of dependence among policyholders, with both insurers and policyholders using exponential utility functions. However, these studies typically assume that the loss distribution is known. We extend the model proposed by Bernard et al. (2020) to account for ambiguity regarding the true loss distribution when only a limited number of samples are available.

## 3 Properties of entropic risk measure

Let  $(\Omega, \mathcal{F}, \mathbb{P})$  is a probability space and let  $\mathcal{L}^p := \mathcal{L}^p(\Omega, \mathcal{F}, \mathbb{P})$  denote the space of real-valued measurable functions,  $X : \Omega \rightarrow \mathbb{R}$  such that  $\mathbb{E}[|X|^p] < \infty$ , for some  $p \geq 1$ . The entropic risk measure is a convex, law invariant risk measure (Föllmer and Schied 2002), thus satisfying the following definition.

**Definition 1.** A functional  $\rho : \mathcal{L}^p \rightarrow \bar{\mathbb{R}}$ , where  $\bar{\mathbb{R}} := \mathbb{R} \cup \{\infty\}$ , is a convex law-invariant risk measure if

- (a)  $\rho(X - m) = \rho(X) - m$  for all  $X \in \mathcal{L}^p$  and  $m \in \mathbb{R}$  and  $\rho(0) = 0$ .
- (b)  $\rho(X) \leq \rho(X')$  if  $X \leq X'$  almost surely (a.s.) for all  $X \in \mathcal{L}^p$ .
- (c)  $\rho(\lambda X + (1 - \lambda)X') \leq \lambda\rho(X) + (1 - \lambda)\rho(X')$  for  $\lambda \in [0, 1]$  and for all  $X, X' \in \mathcal{L}^p$ .
- (d)  $\rho(X) = \rho(X')$  for all  $X, X' \in \mathcal{L}^p$  such that  $X = X'$  in distribution.

Condition (a), also known as cash-invariance property, states that  $m$  is the minimum amount that should be added to a risky position to make it acceptable to a regulator. Condition (b), ensures monotonicity, meaning lower losses are preferable. Condition (c), convexity, ensures that diversification reduces risk. Lastly, condition (d), law invariance, states that two random variables with the same distribution should have equal risk.

Letting  $\boldsymbol{\xi} : \Omega \rightarrow \mathbb{R}^d$  be a random vector and  $\ell(\boldsymbol{\xi}) \in \mathcal{L}^p$  the random loss that it produces, we will further impose the following assumption to ensure that the entropic risk of  $\ell(\boldsymbol{\xi})$  is finite, together with the mean and variance of  $\exp(\alpha\ell(\boldsymbol{\xi}))$ .

**Assumption 1.** The tails of  $\ell(\boldsymbol{\xi})$  are exponentially bounded:

$$\mathbb{P}(|\ell(\boldsymbol{\xi})| > a) \leq G \exp(-\alpha a C), \quad \forall a \geq 0,$$

for some  $G > 0$  and  $C > 2$ . Equivalently, the moment-generating function  $\mathbb{E}[\exp(t\ell(\boldsymbol{\xi}))] \in \mathbb{R}$  for all  $t \in [-\alpha C, \alpha C]$  for some  $C > 2$ , see Lemma 2 in Appendix A.8 for a proof of equivalence.

Assumption 1 further restricts the space of loss functions in Definition 1 in order to work with random variables that are “well-behaved” from the point of view of entropic risk estimation at a risk tolerance level of  $\alpha$ . Indeed, our assumptions will ensure that the empirical estimator is asymptotically consistent. We note that our assumption relates to  $\mathcal{L}_M$ , the set of random variables with finite-valued moment generating functions, through the following inclusion:  $\mathcal{L}^\infty \subseteq \mathcal{L}_M \subseteq \mathcal{L}_\alpha \subseteq \mathcal{L}^p$  with  $\mathcal{L}_\alpha$  as the set of random variables in  $\mathcal{L}_p$  that satisfy Assumption 1.

**Lemma 1.** Under Assumption 1,  $\mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))] \in \left[ \exp(-\frac{2G}{C}), \frac{G}{C-1} \right]$  and  $\text{Var}[\exp(\alpha\ell(\boldsymbol{\xi}))] \in \left[ 0, \frac{2G}{C-2} \right]$ .

## 4 Bias mitigation using bias-aware bootstrapping

In this section, we introduce our proposed estimators designed to address the underestimation problem associated with the empirical entropic risk estimator,  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))$ . The true bias is given by  $\delta = \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) - \mathbb{E}[\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi}))]$ , where the expectation is taken with respect to the randomness in  $\mathcal{D}_N$ . Since the true distribution  $\mathbb{P}$  is unknown, the exact bias cannot be determined. Several approaches rely on CLT to devise asymptotically unbiased estimators. However, with heavy-tailed losses (high risk aversion), a large number of samples is required before the estimator's error tails exhibit Gaussian behavior (see Figure 1). For instance, a typical approach in the literature to devise an unbiased estimator is to use bootstrapping which samples repeatedly from the empirical distribution. Such bootstrapping procedure has been shown to be weakly consistent (DasGupta 2008), however, it exhibits significant bias for small sample sizes. In this paper, we propose a modification to the classical bootstrap algorithm, namely, we first fit a distribution  $\mathbb{Q}_N$  using the  $N$  i.i.d. loss scenarios  $\mathcal{S} := \{\ell(\boldsymbol{\xi}_1), \ell(\boldsymbol{\xi}_2), \dots, \ell(\boldsymbol{\xi}_N)\}$  and then repeatedly sample from  $\mathbb{Q}_N$ , instead of resampling from the empirical distribution. We will demonstrate that fitting a distribution is a crucial step in reducing bias in finite samples, while still ensuring that the estimator is strongly asymptotic consistent. Nevertheless, merely fitting  $\mathbb{Q}_N$  using MLE does not fully resolve the underestimation issue, which is why we introduce bias-aware procedures to better fit the data, see Sections 4.1 and 4.2.

Let  $\zeta \sim \mathbb{Q}_N$  capture the loss  $\ell(\boldsymbol{\xi})$  associated with the uncertain parameter  $\boldsymbol{\xi}$ . Similar to Assumption 1, the following assumption ensures that the mean and variance of  $\exp(\alpha\zeta)$  are finite.

**Assumption 2.** Suppose that the tails of  $\zeta \sim \mathbb{Q}_N$  are almost surely uniformly exponentially bounded. Namely, with probability one with respect to the sample and the fitting procedure of  $\mathbb{Q}_N$ , there exists some  $G > 0$  and some  $C > 2$  such that  $\zeta \sim \mathbb{Q}_N$  satisfies Assumption 1 for all  $N \geq 1$ .

This assumption is not limiting since we have assumed that  $\ell(\boldsymbol{\xi})$  satisfies this assumption under  $\mathbb{P}$ . In practice, this assumption could be satisfied by properly defining the set of models used to estimate  $\mathbb{Q}_N$  from the loss scenarios  $\mathcal{S}$ .

Next we introduce the necessary notation to describe our estimation procedure. Let  $\rho_{\mathbb{Q}_N}(\zeta)$  denote the entropic risk for the distribution  $\mathbb{Q}_N$ . Let  $\hat{\mathbb{Q}}_{N,N}$  represent the empirical distribution of  $N$  values drawn i.i.d. from the estimated distribution  $\mathbb{Q}_N$ , and  $\rho_{\hat{\mathbb{Q}}_{N,N}}(\zeta)$  denote the corresponding empirical entropic risk. Notice that since  $\mathbb{Q}_N$  was estimated using  $N$  loss scenarios, it is itself random, thus  $\rho_{\mathbb{Q}_N}(\zeta)$  and  $\rho_{\hat{\mathbb{Q}}_{N,N}}(\zeta)$  are random variables as well. Our proposed estimator for the bias of the entropic risk is given by

$$\delta_N(\mathbb{Q}_N) := \text{median} \left( \rho_{\mathbb{Q}_N}(\zeta) - \rho_{\hat{\mathbb{Q}}_{N,N}}(\zeta) \mid \mathbb{Q}_N \right), \quad (5)$$

where the median is taken with respect to randomness of samples from  $\mathbb{Q}_N$ . Algorithm 1 provides an estimate for the bias. Given  $N$  loss scenarios  $\mathcal{S}$ , the algorithm first estimates  $\mathbb{Q}_N$ , and then repeatedly samples  $N$  i.i.d. scenarios from  $\mathbb{Q}_N$  to form the empirical distribution  $\hat{\mathbb{Q}}_{N,N}$ . For each of the  $M$  repetitions, it estimates the entropic risk, denoted by the sequence  $\rho_1, \rho_2, \dots, \rho_M$ . Finally, the bias is estimated through  $\hat{\delta}_N(\mathbb{Q}_N) = \text{median}\{\rho_{\mathbb{Q}_N}(\zeta) - \rho_n\}_{n=1}^M$ . Note that as  $M$  increases, the resampling (simulation) error in the bootstrap estimate decreases, and the bootstrap estimate  $\hat{\delta}_N(\mathbb{Q}_N)$  converges to the true estimate  $\delta_N(\mathbb{Q}_N)$ .

---

### Algorithm 1 Bootstrap bias correction.

---

```

1: function BOOTSTRAPBIASCORRECTION( $\mathcal{S}, M$ )
2:    $\mathbb{Q}_N \leftarrow$  Fit a distribution to the loss scenarios  $\mathcal{S}$ 
3:   for  $n \leftarrow 1$  to  $M$  do
4:      $\hat{\mathbb{Q}}_{N,N} \leftarrow$  Draw  $N$  i.i.d. samples from  $\mathbb{Q}_N$ 
5:      $\rho_n \leftarrow \rho_{\hat{\mathbb{Q}}_{N,N}}(\zeta)$ 
6:    $\hat{\delta}_N(\mathbb{Q}_N) \leftarrow \text{median}\{\rho_{\mathbb{Q}_N}(\zeta) - \rho_n\}_{n=1}^M$ 
7:   return  $\hat{\delta}_N(\mathbb{Q}_N)$ 

```

---

In the next theorem, we show that the bias-adjusted empirical risk given by  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) + \delta_N(\mathbb{Q}_N)$  is an asymptotically consistent estimator of the true risk  $\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$ , meaning that as the number of training samples  $N \rightarrow \infty$ , the bias-adjusted empirical risk almost surely converges to the true entropic risk.

**Theorem 1.** *Under Assumptions 1 and 2, the estimator  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) + \delta_N(\mathbb{Q}_N)$  is strongly asymptotically consistent.*

The proof involves two key steps: The first step is to establish that the empirical entropic risk converges to the true risk almost surely. This is achieved by using the strong law of large numbers and the continuous mapping theorem. The strong law of large numbers ensures that the average of exponentiated losses converges almost surely to its expected value, and the continuous mapping theorem extends this convergence to the logarithmic transformation involved in the entropic risk, leading to the almost sure convergence of the empirical risk to the true risk. The second step involves showing that the bias term,  $\delta_N(\mathbb{Q}_N)$ , converges to zero almost surely. This is accomplished by showing that the bias, calculated under any sequence of fitted distributions  $\{\bar{\mathbb{Q}}_N\}_{N=1}^{\infty}$  that satisfy Assumption 2 can be made arbitrarily small for sufficiently large  $N$ . We show that the almost sure convergence of bias to zero is equivalent to proving that the median of a random variable  $X_N$  representing the ratio of empirical and true risk under  $\bar{\mathbb{Q}}_N$  converges to 1 almost surely. To establish this, Chebyshev's inequality is used rather than the CLT, as it applies for any finite  $N \geq 1$ . By setting the upper tail probability to 25% in the Chebyshev's inequality, the median of  $X_N$  is bounded within an interval around 1 which becomes smaller as  $N$  increases. Consequently, the median of  $X_N$  converges to 1 almost surely. Thus, for any sequence of distributions  $\{\bar{\mathbb{Q}}_N\}_{N=1}^{\infty}$  which satisfy Assumption 2, the bias  $\delta_N(\mathbb{Q}_N)$  converges almost surely to zero.

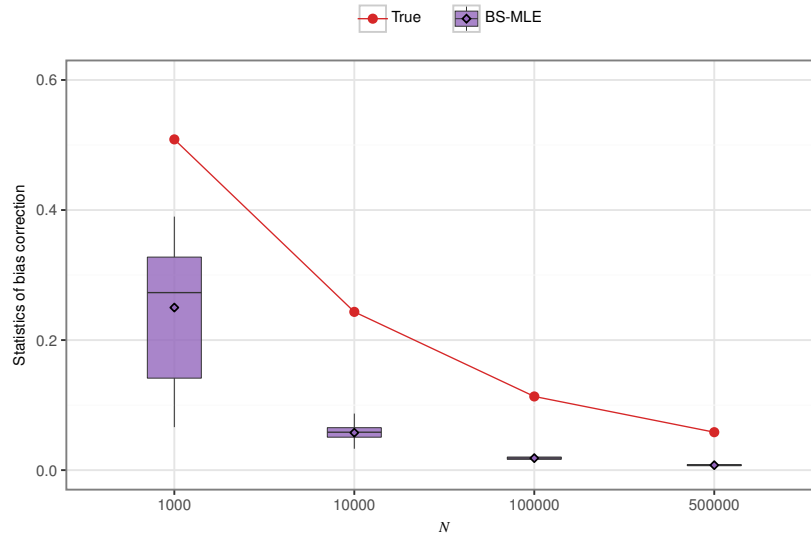
Our approach does not rely on asymptotic or parametric (Gaussian, for instance) assumptions made in the literature to correct the bias. Due to the heavy tailed loss, the asymptotically unbiased estimators proposed in the literature underestimate the risk in the finite sample case. In fact, the underestimation problem persists in our proposed method as well if we naively use our bootstrapping procedure in finite samples, e.g., using MLE to fit a distribution. Since our procedure is flexible in the choice of distributions to fit to the data and the fitting procedure, we will account for the bias due to the tail scenarios in the following subsections.

Among the options for choosing the distribution  $\mathbb{Q}_N$ , we utilize a Gaussian Mixture Model (GMM),  $\mathbb{Q}^{\boldsymbol{\theta}}$ , with parameters  $\boldsymbol{\theta} := (\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ , where  $\boldsymbol{\pi} \in \mathbb{R}^Y$  denotes the weights of the  $Y$  mixtures, and  $\boldsymbol{\mu} \in \mathbb{R}^Y$  and  $\boldsymbol{\sigma} \in \mathbb{R}^Y$  denote the means and standard deviations of the mixtures, respectively. There are two advantages for using GMM. First, GMMs are universal density approximators, meaning they can approximate any smooth density given sufficient data, and second, the moment-generating function of a random variable  $\zeta \sim \mathbb{Q}^{\boldsymbol{\theta}}$  exists for all  $\alpha$ , and thus the entropic risk  $\rho_{\mathbb{Q}^{\boldsymbol{\theta}}}(\zeta) = \frac{1}{\alpha} \log \left( \sum_{y=1}^Y \pi_y \exp(\alpha \mu_y + \frac{\alpha^2}{2} \sigma_y^2) \right)$  can be obtained in closed form. This eliminates the need to estimate the entropic risk through simulation in step 6 of Algorithm 1, which can be particularly beneficial when the risk aversion parameter  $\alpha$  is large, as this would otherwise require a large number of samples for accuracy.

The natural approach for fitting the parameters  $\boldsymbol{\theta}$  of a GMM is to use MLE, typically achieved via the Expectation-Maximization (EM) algorithm (Dempster et al. 1977). The following example demonstrates that with limited samples, if step 2 of Algorithm 1 utilizes the Expectation-Maximization algorithm to fit a GMM, the estimate will still underestimate the true risk.

**Example 2.** Consider the problem of estimating the entropic risk of a random variable that follows a Gaussian mixture model with two components  $\xi \sim \text{GMM}(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ ,  $\boldsymbol{\pi} = [0.7 \ 0.3]$ ,  $\boldsymbol{\mu} = [0.5 \ 1]$ , and  $\boldsymbol{\sigma} = [2 \ 1]$ . To obtain Figure 2, we draw  $N$  i.i.d. samples from  $\text{GMM}(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$  where  $N \in \{10^3, 10^4, 10^5, 5 \cdot 10^5\}$ . The true bias correction is obtained by first computing the true entropic risk and then subtracting the expected empirical entropic risk  $\bar{\rho} = \frac{1}{1000} \sum_{i=1}^{1000} \frac{1}{\alpha} \log \left( \mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha \xi_i)] \right)$  obtained

by bootstrapping with 1000 repetitions.<sup>1</sup> We fit a GMM  $\mathbb{Q}^\theta$  to the samples using the Expectation-Maximization algorithm and use Algorithm 1 to estimate the bias. The boxplots are plotted by resampling 100 times from GMM( $\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma}$ ). Fitting a GMM by MLE still underestimates the true bias for finite number of samples. Also, we can observe that as the number of training samples increase, the bias estimated by fitted GMM converges toward 0.



**Figure 2: Statistics of bias correction estimated from fitting a GMM by MLE followed by bootstrapping over 100 resampling from the underlying distribution.**

We next propose different distribution fitting strategies that can aid the bias mitigation in finite samples. The aim is to fit a distribution such that  $N$  i.i.d. samples drawn from the fitted distribution replicate the bias observed in  $N$  i.i.d. samples from the true distribution  $\mathbb{P}$ . We refer to this approach as “bias-aware” distribution matching. This concept is inspired by “decision-aware” learning methods in contextual optimization problems (Elmachtoub and Grigas 2021, Donti et al. 2017, Grigas et al. 2023, Sadana et al. 2025), where statistical accuracy is deliberately traded for improved decision outcomes.

#### 4.1 Entropic risk matching

In this section, we describe Algorithm 2 which is used to learn the parameters  $\boldsymbol{\theta}$  of the GMM,  $\mathbb{Q}^\theta$ , such that the distribution of entropic risk for  $n$  samples drawn from  $\mathbb{Q}^\theta$  matches the empirical distribution of entropic risk for  $n$  samples drawn from the empirical loss scenarios  $\mathcal{S}$ . To construct the latter distribution,  $\mathcal{S}$  is divided into  $B$  bins, with each bin containing  $n = N/B$  scenarios. The entropic risk is computed for each bin, forming the set  $\mathcal{R}_{\mathcal{D}_N}$ . The corresponding empirical distribution,  $\hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}$ , over the set  $\mathcal{R}_{\mathcal{D}_N}$ , captures the variability of entropic risk across the  $B$  bins. For the former distribution, with a fixed  $\boldsymbol{\theta}$ ,  $B' \times n$  i.i.d. samples are drawn from  $\mathbb{Q}^\theta$  and divided into  $B'$  bins. The entropic risk is then computed for the scenarios in each bin, yielding the set  $\mathcal{R}_\theta$ . The corresponding empirical distribution,  $\hat{\mathbb{P}}_{\mathcal{R}_\theta}$ , captures the variability of entropic risk across the  $B'$  bins.

Next, the algorithm compares the empirical distribution  $\hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}$  with the model-based distribution  $\hat{\mathbb{P}}_{\mathcal{R}_\theta}$  using the following Wasserstein distance:

$$\mathcal{W}^2(\hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_\theta}) = \left( \int_0^1 |F_{\mathcal{R}_{\mathcal{D}_N}}^{-1}(q) - F_{\mathcal{R}_\theta}^{-1}(q)|^2 dq \right)^{1/2},$$

<sup>1</sup>We repeat this procedure 100 times, compute the estimate of true entropic risk, and the 95% confidence interval for the true entropic risk is contained in the marker drawn on Figure 2 for the true entropic risk.

where  $F_{\mathcal{R}_{\mathcal{D}_N}}^{-1}$  and  $F_{\mathcal{R}_{\theta}}^{-1}$  are quantile functions associated with sets  $\mathcal{R}_{\mathcal{D}_N}$  and  $\mathcal{R}_{\theta}$ , respectively. This distance quantifies the discrepancy between the two distributions. The algorithm iteratively adjusts the GMM parameters to minimize this distance. It uses gradient descent to update the parameters  $\theta_t$  at each iteration as follows:

$$\theta_{t+1} = \theta_t - \gamma \nabla_{\theta_t} \mathcal{W}^2 \left( \hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_{\theta_t}} \right),$$

where  $\gamma$  is the step size. To enable computation of the gradients of the Wasserstein distance with respect to  $\theta$ , the algorithm employs differentiable sampling techniques that enable automatic differentiation through the sampling process (see Algorithm 6 in Appendix B.4). This is based on reparameterization approach (Kingma et al. 2015), which allows stochastic sampling operations to be expressed in a differentiable manner. The iterative process continues until the Wasserstein distance  $\mathcal{W}^2 \left( \hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_{\theta_t}} \right)$  falls below a predefined convergence threshold  $\epsilon$ , or until a maximum number of iterations  $T$  is reached. Further details of the algorithm can be found in Appendix B.3.

Even though computing the Wasserstein distance between distribution of losses has a worst-case complexity  $O(B' \log(B'))$  (Kolouri et al. 2019), there is a significant total cost associated with the gradient descent procedure described in Algorithm 2. In the next section, we provide a semi-analytic procedure to learn a two-component GMM that can account for the tail scenarios.

---

**Algorithm 2 Fit GMM by entropic risk matching.**


---

```

1: function BS-MATCH( $\mathcal{S}$ )
2:   Divide loss scenarios in  $\mathcal{S}$  into  $B$  bins, each of size  $n$ 
3:   Compute the entropic risk in  $B$  bins, forming the empirical distribution  $\hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}$ 
4:   Fit a GMM  $\mathbb{Q}^{\theta}$  to  $\mathcal{S}$  with  $Y$  components using EM algorithm
5:   Initialize the iteration counter  $t \leftarrow 0$  and  $\mathfrak{D} \leftarrow \infty$ 
6:   while  $d > \epsilon$  and  $t < T$  do
7:     Draw  $B' \times n$  i.i.d. samples from  $\mathbb{Q}^{\theta_t}$ , split into  $B'$  bins
8:     Compute entropic risk in each bin, forming  $\hat{\mathbb{P}}_{\mathcal{R}_{\theta_t}}$ 
9:     Update GMM parameters:  $\theta_{t+1} \leftarrow \theta_t - \gamma \nabla_{\theta_t} \mathcal{W}^2 \left( \hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_{\theta_t}} \right)$ 
10:    Project  $\theta_{t+1}$  onto the feasible region of a GMM
11:    Update distance:  $\mathfrak{D} \leftarrow \mathcal{W}^2 \left( \hat{\mathbb{P}}_{\mathcal{R}_{\mathcal{D}_N}}, \hat{\mathbb{P}}_{\mathcal{R}_{\theta_t}} \right)$ 
12:    Increment iteration counter:  $t \leftarrow t + 1$ 
13:  return  $\mathbb{Q}^{\theta_t}$ 

```

---

## 4.2 Matching the extremes

This section focuses on learning the parameters  $\theta$  of a GMM  $\mathbb{Q}^{\theta}$  by giving special attention to the tails of the loss distribution, aiming to improve the accuracy of the entropic risk estimation. Although the entropic risk depends on the entire distribution, its sensitivity to extreme values makes an accurate approximation of tail behavior essential. Figure 14 in Appendix B.1 illustrates the influence function, highlighting how tail events disproportionately affect the entropic risk. Our distribution fitting approach strikes a balance between capturing the tail behavior and retaining analytical tractability.

Our proposed procedure is motivated by the Fisher–Tippett–Gnedenko extreme value theorem (de Haan and Ferreira 2006) which states that given i.i.d. samples of  $\{\zeta_1, \zeta_2, \dots, \zeta_n\}$  with cumulative distribution function (cdf) given by  $F(\cdot)$ , the distribution of the (normalized) maxima  $M_n = \max\{\zeta_1, \zeta_2, \dots, \zeta_n\}$  converges to a non-degenerate distribution  $G$ :

$$\lim_{n \rightarrow \infty} \mathbb{P} \left( \frac{M_n - b_n}{a_n} \leq x \right) = \lim_{n \rightarrow \infty} F(a_n x + b_n)^n \rightarrow G(x),$$

where  $a_n$  and  $b_n$  are normalizing sequences of scale and location parameters, respectively, that ensure the limit exists and  $F(\cdot)^n$  is the cdf of  $M_n$ . The limit distribution  $G$  belongs to one of three extreme value distributions—Weibull, Fréchet or Gumbel distribution—depending on the tails of  $F(\cdot)$ .

Algorithm 3 constructs an equally-weighted two-component GMM to represent the loss distribution. The first component aims to estimate the distribution of maxima  $M_n$  of the loss scenarios. To this end, first we divide the loss scenarios  $\mathcal{S}$  into  $B$  bins, each of size  $n = N/B$ . We store the maximum within each bin in set  $\mathcal{M}$ , where  $F_{\mathcal{M}}$  denotes the cdf of scenarios in  $\mathcal{M}$ . This approach is typically referred to as the block maxima method (de Haan and Ferreira 2006). Motivated by the extreme value theory, the algorithm estimates the parameters of the first component by matching the cdf of maximum of  $n$  i.i.d. samples from  $\mathcal{N}(\mu^e, \sigma^e)$ , denoted by  $\Phi_{\mu^e, \sigma^e}^n$  to  $F_{\mathcal{M}}$ . The calculation of the parameters  $(\mu^e, \sigma^e)$  can be done in a semi-analytic way as discussed in Appendix B.3. For the second component of the GMM, we set the mean to  $2(\mu_{\mathcal{S}} - 0.5\mu^e)$ , where  $\mu_{\mathcal{S}}$  is the mean of the loss scenarios, and set the standard deviation to zero. This choice ensures that the mean of overall GMM is equal to the mean of the loss scenarios  $\mu_{\mathcal{S}}$ .

The reasoning behind the first component of the GMM is the following. Fitting an extreme value distribution directly to the maxima  $\mathcal{M}$  is computationally intensive, often inaccurate with finite samples, and typically only informative for upper tail risk, even with large datasets. Instead, fitting a normal distribution with our proposed approach is computationally efficient since it can be done semi-analytically. Furthermore, our approach can provide a reasonable approximation of the tails since both  $\Phi_{\mu^e, \sigma^e}^n$  and tails of true loss distribution converge to an extreme value distribution as  $n$  goes to infinity. We emphasize that our approach is different from fitting a normal distribution directly to scenarios in  $\mathcal{S}$  or  $\mathcal{M}$  by MLE.

---

**Algorithm 3 Fit GMM based on the extreme value theory.**

---

```

1: function BS-EVT( $\mathcal{S}$ )
2:   Divide scenarios in  $\mathcal{S}$  into  $B$  bins of size  $n = N/B$  each
3:    $F_{\mathcal{M}} \leftarrow$  cdf of maxima in each bin
4:   Determine  $(\mu^e, \sigma^e)$  so that  $\Phi_{\mu^e, \sigma^e}^n$  matches  $F_{\mathcal{M}}$ 
5:    $\hat{\pi} \leftarrow \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}$ ,  $\hat{\mu} \leftarrow \begin{pmatrix} \mu^e \\ 2(\mu_{\mathcal{S}} - 0.5\mu^e) \end{pmatrix}$ ,  $\hat{\sigma} \leftarrow \begin{pmatrix} \sigma^e \\ 0 \end{pmatrix}$ 
6:    $\theta \leftarrow (\hat{\pi}, \hat{\mu}, \hat{\sigma})$ 
7:   return  $\mathbb{Q}^\theta$ 

```

---

We end this section by revisiting Example 2, to include the estimation procedures from Sections 4.1 and 4.2. We have already shown in Figure 2 that naively fitting a GMM using MLE underestimates the true risk. In Figure 3, we observe the behavior of the bootstrap (Algorithm 1) in conjunction with the entropic risk matching (Algorithm 2) denoted by **BS-Match**, and the bootstrap (Algorithm 1) in conjunction with the extremes matching (Algorithm 3) denoted by **BS-EVT**. We can observe that the bias correction using **BS-Match** decays at a rate similar to the true bias. However, the bias associated with **BS-MLE** decays at the faster rate compared to the true bias estimate, whereas for **BS-EVT**, it decays slowly. Thus, **BS-MLE** tends to underestimate the risk, while **BS-EVT** tends to overestimate the true risk.

### 4.3 Comparison with benchmarks from the literature

In this section, we review several methods from the literature for estimating entropic risk and show through a numerical example that they significantly underestimate it. We start with a brief summary of each method.


Entropic risk measure can be equivalently written as an optimized certainty equivalent risk measure (Ben-Tal and Teboulle 1986), i.e.,

$$\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) = \inf_t \mathbb{E}_{\mathbb{P}}[h(t, \ell(\boldsymbol{\xi}))], \quad (6)$$

where  $h(t, \ell(\boldsymbol{\xi})) = t + \frac{1}{\alpha} \exp(\alpha(\ell(\boldsymbol{\xi}) - t)) - \frac{1}{\alpha}$ .

**Sample average approximation (SAA):** The SAA estimator is obtained by solving SAA of problem (6):

$$\rho_{\text{SAA}} := \frac{1}{\alpha} \log(\mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha \ell(\boldsymbol{\xi}))]) = \inf_t \mathbb{E}_{\hat{\mathbb{P}}_N}[h(t, \ell(\boldsymbol{\xi}))],$$



**Figure 3: Statistics of bias correction estimated from fitting a GMM by MLE (BS-MLE), entropic risk matching (BS-Match) and tail fitting (BS-EVT) followed by bootstrapping over 100 resampling from the underlying distribution.**

where  $\mathbb{P}$  is replaced with  $\hat{\mathbb{P}}_N$ . It is well-known that decisions based on the SAA can suffer from optimizer’s curse, leading to an optimistic bias (Smith and Winkler 2006). To mitigate the underestimation of the optimal value  $\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$  by the SAA estimator, leave-one out CV and optimizer’s Information Criteria (OIC) are proposed in the literature which are discussed below.

**Leave-one out cross validation (LOOCV):** Let  $\hat{\mathbb{P}}_{N-i}$  denote the empirical distribution without the  $i$ th scenario, and let  $\hat{t}_{-i}$  denote the optimal solution of (6) in which  $\hat{\mathbb{P}}_{N-i}$  is used. The estimator is then defined as:

$$\rho_{\text{LOOCV}} := \frac{1}{N} \sum_{i=1}^N \left( \hat{t}_{-i} + \frac{1}{\alpha} \left( \exp(\alpha(\ell(\hat{\boldsymbol{\xi}}_i) - \hat{t}_{-i})) - 1 \right) \right).$$

Since  $\hat{t}_{-i}$  is a feasible solution of (6), we have  $\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) \leq \mathbb{E}_{\mathbb{P}}[h(\hat{t}_{-i}, \ell(\boldsymbol{\xi}))]$  almost surely with respect to the randomness of  $\hat{t}_{-i}$  for all  $i \in [N]$ . Thus,

$$\begin{aligned} \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) &\leq \mathbb{E} \left[ \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathbb{P}}(h(\hat{t}_{-i}, \ell(\boldsymbol{\xi}))) \right] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{E} [\mathbb{E}_{\mathbb{P}}(h(\hat{t}_{-i}, \ell(\boldsymbol{\xi})))] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[ \mathbb{E}(h(\hat{t}_{-i}, \ell(\hat{\boldsymbol{\xi}}_i)) | \{\hat{\boldsymbol{\xi}}_j\}_{j \in [N]-i}) \right] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}[h(\hat{t}_{-i}, \ell(\hat{\boldsymbol{\xi}}_i))] \\ &= \mathbb{E}[\rho_{\text{LOOCV}}]. \end{aligned}$$

where the first equality follows from linearity of expectation, the second is due to the independence of  $\hat{\boldsymbol{\xi}}_i \sim \mathbb{P}$  and  $\{\hat{\boldsymbol{\xi}}_j\}_{j \in [N]-i}$  and the third follows from the law of iterated expectations. Thus,  $\rho_{\text{LOOCV}}$  is a positively biased estimator of  $\rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$ .



**Optimizer’s information criteria (OIC):** To correct the first-order optimistic bias associated with the SAA, Iyengar et al. (2023) introduced the OIC estimator:

$$\rho_{\text{OIC}} := \rho_{\text{SAA}} + \frac{\text{Var}_{\hat{\mathbb{P}}_N}(\exp(\alpha\ell(\boldsymbol{\xi})))}{N\alpha(\mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha\ell(\boldsymbol{\xi}))])^2},$$

(see Appendix B.2 for detailed calculations).<sup>2</sup>


**Maximum Likelihood Estimation (MLE):** The MLE estimator involves fitting a GMM  $\mathbb{Q}^\theta$  to the loss scenarios using the EM algorithm, and the estimator is given by  $\rho_{\text{MLE}} := (1/\alpha) \log(\mathbb{E}_{\mathbb{Q}^\theta}[\exp(\alpha\ell(\boldsymbol{\xi}))])$ .

**Median-of-means (MoM):** The MoM estimator is constructed by dividing the loss scenarios into  $\lfloor \sqrt{N} \rfloor$  blocks, calculating the entropic risk within each block, and then taking the median of these entropic risk values (Lugosi and Mendelson 2019).

**Bootstrap (BS):** The BS estimator is calculated by repeatedly sampling  $N$  observations with replacement from the empirical distribution of loss scenarios. For each bootstrap sample, the empirical entropic risk is computed. After a large number of repetitions, the average of these entropic risks provides the bootstrap estimator.

**Example 3.** Consider a project selection problem with three projects. Let  $\xi \sim \text{GMM}(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$  with  $Y = 5$  (see Appendix B.5 for parameter values of the GMM with randomly generated weights  $\boldsymbol{\pi}$ ). Suppose the losses associated with the three projects are given by  $0.4\xi$ ,  $0.6\xi$  and  $0.8\xi$ , respectively. Let the risk aversion parameter be  $\alpha = 3$ . The true entropic risk can be analytically expressed as  $\rho_{\mathbb{P}}(\xi) = \frac{1}{3} \log\left(\sum_{y=1}^Y \pi_y \exp\left(3\mu_y + \frac{9}{2}\sigma_y^2\right)\right)$ . To evaluate each estimator, we draw 1000 instances with a sample size of  $N = 10000$  from the GMM  $(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ .

Figure 4 shows that the estimators discussed above significantly underestimate the true entropic risk for projects 2 and 3. It is interesting to note that although the LOOCV estimator overestimates the



**Figure 4: Statistics of the estimates of the true entropic risk obtained from different models for each project.**

true entropic risk, its distribution across the 1000 instances is skewed, with the entropic risk falling

<sup>2</sup>It is to be noted that the OIC estimator is designed to address the optimizer’s curse, while we aim to debias the empirical entropic risk estimator. The OIC estimator relies on the asymptotic normality and linearity assumption of SAA estimator. With heavy tailed losses, these assumptions on SAA estimator do not carry over to practical scenarios with limited data, hence, the underestimation issue persists. Even though the true performance of a decision obtained by solving the SAA problem overestimates the true entropic risk, the OIC estimator based on the empirical distribution underestimates the true risk.



below the true value in over 75% of instances. In contrast, our proposed estimators effectively address the underestimation issue. Specifically, **BS-EVT** overestimates the entropic risk and **BS-Match** reduces the bias of **SAA** estimator significantly.

In the next section, we will show that the proposed procedures for mitigating estimation bias can also be applied to mitigate the optimistic bias when solving entropic risk minimization problems. As discussed earlier, optimistic bias occurs due to lack of data, in which case regularization type techniques (such as distributionally robust optimization) are employed to correct the bias. By providing better estimates of the validation risk, we can more accurately calibrate the hyperparameters compared to traditional CV methods.

## 5 Distributionally robust optimization

Entropic risk minimization considers the following problem:

$$\rho^* = \min_{\mathbf{z} \in \mathcal{Z}} \rho_{\mathbb{P}}(\ell(\mathbf{z}, \boldsymbol{\xi})) := \frac{1}{\alpha} \log \left( \mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))] \right), \quad (7)$$

where  $\mathcal{Z} \subseteq \mathbb{R}^d$  denotes the set of all feasible decisions,  $\boldsymbol{\xi} \in \mathbb{R}^d$  denotes the uncertain vector following the probability distribution  $\mathbb{P}$ , and  $\ell(\mathbf{z}, \boldsymbol{\xi})$  denotes the loss function. For the optimal solution  $\mathbf{z}^*$  of problem (7) to be well defined, we make the following standard assumptions:

**Assumption 3.** We assume that:

- (A.1)  $\mathcal{Z}$  is a compact and convex set.
- (A.2)  $\ell(\mathbf{z}, \boldsymbol{\xi})$  is convex in  $\mathbf{z}$  for almost every  $\boldsymbol{\xi} \in \Xi$ .
- (A.3)  $\ell(\mathbf{z}, \boldsymbol{\xi})$  is  $L$ -Lipschitz continuous in  $\boldsymbol{\xi}$  for all  $\mathbf{z} \in \mathcal{Z}$ .
- (A.4)  $\ell(\mathbf{z}, \boldsymbol{\xi})$  is  $L(\boldsymbol{\xi})$ -Lipschitz continuous in  $\mathbf{z}$  for all  $\boldsymbol{\xi} \in \Xi$  with  $\mathbb{E}_{\mathbb{P}}[L(\boldsymbol{\xi})^q] < \infty$  for all  $q \geq 1$ .
- (A.5)  $|\ell(\mathbf{z}, \boldsymbol{\xi})| \leq \bar{L}(\boldsymbol{\xi})$  for all  $\mathbf{z} \in \mathcal{Z}$  almost surely, with the tail of  $\bar{L}(\boldsymbol{\xi})$  exponentially bounded:

$$\mathbb{P}(\bar{L}(\boldsymbol{\xi}) > a) \leq G \exp(-aaC) \quad \text{for all } a \geq 0,$$

for some constants  $G > 0$  and  $C > 2$ .

As in Assumption 2, the last assumption ensures that the mean and variance of the loss  $\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))$  are finite for each  $\mathbf{z} \in \mathcal{Z}$ . We make the technical assumptions (A.4) and (A.5) to ensure the convergence of SAA solution to the true risk (see Proposition 1).

As the true underlying distribution  $\mathbb{P}$  is typically unknown, it is common practice to replace it with the empirical distribution  $\hat{\mathbb{P}}_N$ , solving the corresponding SAA problem:

$$\rho_{\text{SAA}} = \min_{\mathbf{z} \in \mathcal{Z}} \rho_{\hat{\mathbb{P}}_N}(\ell(\mathbf{z}, \boldsymbol{\xi})) := \frac{1}{\alpha} \log \left( \mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))] \right). \quad (8)$$

Under Assumption 3, it follows that  $\rho_{\text{SAA}} \rightarrow \rho^*$  as  $N$  grows, as shown in Proposition 1.

**Proposition 1.** *Suppose that Assumption 3 holds. Then, for any  $\gamma > 0$ , there exists a constant  $A > 0$  for which,*

$$\mathbb{P} \left( |\rho_{\text{SAA}} - \rho^*| \geq \frac{A}{\sqrt{N} \gamma \alpha \exp(\alpha \rho^*)} \right) \leq \gamma, \quad (9)$$

as long as  $N$  is sufficiently large. Consequently,  $\rho_{\text{SAA}} \rightarrow \rho^*$  in probability.

To prove Proposition 1, we show that  $\exp(\alpha \ell(\boldsymbol{\xi}, \mathbf{z}))$  is  $\kappa(\boldsymbol{\xi})$ -Lipschitz continuous in  $\mathbf{z}$  for all  $\boldsymbol{\xi} \in \Xi$  with  $\kappa(\boldsymbol{\xi}) = \alpha L(\boldsymbol{\xi}) \exp(\alpha \bar{L}(\boldsymbol{\xi}))$ . Then, we apply the uniform convergence results for heavy tailed distributions in Jiang et al. (2020, Theorem 3.2) to show the uniform convergence of empirical utility to true utility for all  $\mathbf{z} \in \mathcal{Z}$ . This theorem only requires that the second moment of  $\kappa(\boldsymbol{\xi})$  is finite instead

of the usual light-tailed assumptions that don't hold for  $\kappa(\boldsymbol{\xi})$ . Subsequently, we use the properties of logarithm function in the neighborhood of zero to show the uniform convergence of empirical risk to optimal risk. This convergence result underpins the prevalent use of the SAA approach for entropic risk minimization problems (Chen et al. 2024a,b). For a detailed examination of the SAA methodology within stochastic programming, see Shapiro et al. (2009).

In the limited data setting, the risk produced by solving problem (8) underestimates the true risk  $\rho^*$  due to overfitting on the empirical distribution. Distributionally robust optimization (DRO) is one of the approaches to mitigate the optimistic bias of SAA by robustifying decisions against perturbations in the empirical distribution (Delage and Ye 2010, Wiesemann et al. 2014, Mohajerin Esfahani and Kuhn 2018, Rahimian and Mehrotra 2022). It is assumed that nature perturbs the empirical distribution within a distributional ambiguity set  $\mathcal{B}(\epsilon)$  containing all distributions  $\mathbb{Q}$  that are at a "distance"  $\epsilon \geq 0$  away from the empirical distribution  $\hat{\mathbb{P}}_N$ , so as to maximize the entropic risk of the decision maker, while decision maker aims to minimize the worst-case risk resulting in the following min-max problem:

$$\rho_{\text{DRO}} := \min_{z \in \mathcal{Z}} \sup_{\mathbb{Q} \in \mathcal{B}(\epsilon)} \frac{1}{\alpha} \log (\mathbb{E}_{\mathbb{Q}}[\exp(\alpha \ell(z, \boldsymbol{\xi}))]). \quad (10)$$

The size of the ambiguity set,  $\epsilon$ , is chosen by the decision maker; however, as we will discuss later, in practice it is treated as a hyperparameter tuned through CV to optimize the performance of the optimization model on unseen data.

In the literature, different ambiguity sets have been considered with the Kullback Leibler (KL)-divergence (Hu and Hong 2012) and Wasserstein ambiguity sets (Mohajerin Esfahani and Kuhn 2018) being the most commonly used (Rahimian and Mehrotra 2022). For KL-divergence-based ambiguity sets, the standard formulation (Hu and Hong 2012) restricts the worst-case distribution to be absolutely continuous with respect to the empirical distribution, limiting its support to the same points as the empirical distribution. This poses a problem because it prevents the representation of worst-case scenarios that typically occur in the tails of the loss distribution. An alternative formulation does allow worst-case distributions with support beyond the empirical distribution, enabling a richer ambiguity set (Chan et al. 2024). However, with an unbounded loss function, this flexibility allows nature to exploit the tail, resulting in infinite loss for the decision maker. Consequently, KL-divergence-based ambiguity sets are ill-suited to our problem, which involves unbounded support and heavy-tailed losses. This also holds for type- $p$  Wasserstein ambiguity set with  $p < \infty$  due to the following result.

**Proposition 2.** *The  $p$ -Wasserstein DRO with entropic risk measure results in unbounded loss if  $p < \infty$ .*

Next, we show that type- $\infty$  Wasserstein ambiguity set is a suitable choice for problem (10). The type- $\infty$  Wasserstein distance is defined as:

$$\mathcal{W}^\infty(\mathbb{P}_1, \mathbb{P}_2) := \inf_{\pi \in \mathcal{M}(\Xi \times \Xi)} \{\text{ess.sup} \|\zeta_1 - \zeta_2\| \pi(d\zeta_1, d\zeta_2)\},$$

where  $\pi$  is a joint distribution of  $\zeta_1$  and  $\zeta_2$  with marginals  $\mathbb{P}_1$  and  $\mathbb{P}_2$ , respectively,  $\text{ess.sup}$  denotes essential supremum, and  $\|\cdot\|$  denotes the norm. Then, type- $\infty$  Wasserstein ambiguity set  $\mathcal{B}_\infty(\epsilon)$  of radius  $\epsilon \geq 0$ , can be defined as follows:

$$\mathcal{B}_\infty(\epsilon) := \left\{ \mathbb{Q} \in \mathcal{M}(\Xi) \mid \mathbb{Q}\{\boldsymbol{\xi} \in \Xi\} = 1, \mathcal{W}_\infty(\mathbb{Q}, \hat{\mathbb{P}}_N) \leq \epsilon \right\}. \quad (11)$$

Bertsimas et al. (2023) have shown that problem (10) can be equivalently written as:

$$\min_{z \in \mathcal{Z}} \sup_{\mathbb{Q} \in \tilde{\mathcal{B}}_\infty(\epsilon)} \frac{1}{\alpha} \log (\mathbb{E}_{\mathbb{Q}}[\exp(\alpha \ell(z, \boldsymbol{\xi}))]) = \min_{z \in \mathcal{Z}} \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i \in [N]} \sup_{\boldsymbol{\xi}: \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\| \leq \epsilon} \exp(\alpha \ell(z, \boldsymbol{\xi})) \right), \quad (12)$$

where the ambiguity set  $\tilde{\mathcal{B}}_\infty(\epsilon)$  is defined as:

$$\tilde{\mathcal{B}}_\infty(\epsilon) := \left\{ \mathbb{Q} \in \mathcal{M}(\Xi) \mid \exists \boldsymbol{\xi}_i \in \Xi, \|\boldsymbol{\xi}_i - \hat{\boldsymbol{\xi}}_i\| \leq \epsilon, \forall i \in [N], \mathbb{Q}(\boldsymbol{\xi}) = \frac{1}{N} \sum_{i=1}^N \delta_{\boldsymbol{\xi}_i}(\boldsymbol{\xi}) \right\}.$$

The equality in (12) follows since the logarithm function is monotone. From the formulation of problem (12), we can clearly see the behavior of the adversarial uncertainty, which allows each scenario  $\hat{\xi}_i$  to be perturbed within a norm ball of radius  $\epsilon$ . Thus, by construction, the worst-case loss is always bounded for finite radius  $\epsilon$  of the uncertainty set.

For piecewise concave loss functions, the following theorem gives an equivalent reformulation of the DRO problem as the finite dimensional convex optimization problem using Fenchel duality (Ben-Tal et al. 2015).

**Theorem 2.** *Let  $\ell(\mathbf{z}, \boldsymbol{\xi}) = \max_{j \in [m]} \ell_j(\mathbf{z}, \boldsymbol{\xi})$  where  $\ell_j(\mathbf{z}, \boldsymbol{\xi})$  is a concave function in  $\boldsymbol{\xi}$  for each  $j \in [m]$  and  $\mathbf{z} \in \mathcal{Z}$ . Then, the DRO problem (12) with type- $\infty$  Wasserstein ambiguity set is equivalent to*

$$\begin{aligned} \min \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ \text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathcal{Z}, \boldsymbol{\varphi}_{ij} \in \mathbb{R}^d \quad \forall i \in [N], j \in [m] \\ & \boldsymbol{\varphi}_{ij}^\top \hat{\xi}_i - \ell_{j*}(\mathbf{z}, \boldsymbol{\varphi}_{ij}) + \epsilon \|\boldsymbol{\varphi}_{ij}\|_* \leq t_i \quad \forall i \in [N], j \in [m], \end{aligned} \quad (13)$$

where  $\ell_{j*}(\mathbf{z}, \boldsymbol{\varphi}_{ij}) := \inf_{\boldsymbol{\xi}} \{\boldsymbol{\varphi}_{ij}^\top \boldsymbol{\xi} - \ell_j(\mathbf{z}, \boldsymbol{\xi})\}$  is the partial concave conjugate of  $\ell_j(\mathbf{z}, \boldsymbol{\xi})$ , and  $\|\cdot\|_*$  denotes the dual norm.

The DRO reformulation and reformulation technique simplify significantly when the loss function is either piecewise linear or linear, rather than piecewise convex in  $\mathbf{z}$  and concave in  $\boldsymbol{\xi}$ . The following corollary presents these special cases.

**Corollary 1.** *Let  $\ell(\mathbf{z}, \boldsymbol{\xi}) := \max_{k \in \mathcal{K}} \{a_k(\mathbf{z}^\top \boldsymbol{\xi}) + b_k\}$  be a piecewise linear function for given parameters  $a_k$  and  $b_k$ . Then, the DRO problem (12) with type- $\infty$  Wasserstein ambiguity set is equivalent to*

$$\begin{aligned} \rho_{DRO} := \min \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N t_i \right) \\ \text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathcal{Z} \\ & \exp \left( \alpha \left( a_k(\mathbf{z}^\top \hat{\xi}_i) + b_k \right) + \epsilon \|a_k \mathbf{z}\|_* \right) \leq t_i \quad \forall i, k \in \mathcal{K} \end{aligned} \quad (14)$$

which for a linear loss  $\ell(\mathbf{z}, \boldsymbol{\xi}) = \mathbf{z}^\top \boldsymbol{\xi}$  can be further simplified to

$$\min_{\mathbf{z} \in \mathcal{Z}} \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha \mathbf{z}^\top \hat{\xi}_i) \right) + \epsilon \|\mathbf{z}\|_*. \quad (15)$$

**Proof.** Here, we provide an alternative proof for the piecewise linear loss functions  $\ell(\mathbf{z}, \boldsymbol{\xi}) := \max_{k \in \mathcal{K}} \{a_k(\mathbf{z}^\top \boldsymbol{\xi}) + b_k\}$  that does not rely on Fenchel duality (Ben-Tal et al. 2015). The supremum of  $\exp(\max_{k \in \mathcal{K}} \{\alpha (a_k(\mathbf{z}^\top \boldsymbol{\xi}) + b_k)\})$  over the set  $\{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \hat{\xi}_i\| \leq \epsilon\}$  is given by:

$$\begin{aligned} & \sup_{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \hat{\xi}_i\| \leq \epsilon} \exp \left( \max_{k \in \mathcal{K}} \{ \alpha (a_k(\mathbf{z}^\top \boldsymbol{\xi}) + b_k) \} \right) \\ &= \sup_{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \hat{\xi}_i\| \leq \epsilon} \max_{k \in \mathcal{K}} \{ \exp(\alpha (a_k(\mathbf{z}^\top \boldsymbol{\xi}) + b_k)) \} \\ &= \max_{k \in \mathcal{K}} \left\{ \exp \left( \alpha \left( a_k(\mathbf{z}^\top \hat{\xi}_i) + b_k \right) + \alpha \sup_{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \hat{\xi}_i\| \leq \epsilon} (a_k \mathbf{z}^\top \boldsymbol{\xi}) \right) \right\} \\ &= \max_{k \in \mathcal{K}} \left\{ \exp \left( \alpha \left( a_k(\mathbf{z}^\top \hat{\xi}_i) + b_k \right) + \alpha \epsilon \|a_k \mathbf{z}\|_* \right) \right\}, \end{aligned}$$

where the first equality follows from interchanging exp and max operations and then using the fact that  $\exp(\cdot)$  is increasing in its arguments, last equality follows from the definition of the dual norm and  $\|\cdot\|_*$  denotes the dual norm of  $\|\cdot\|$ . On combining with the objective function in (12), we obtain:

$$\frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \max_{k \in \mathcal{K}} \left\{ \exp \left( \alpha \left( a_k(\mathbf{z}^\top \hat{\xi}_i) + b_k \right) + \alpha \epsilon \|a_k \mathbf{z}\|_* \right) \right\} \right).$$

So, with a piecewise linear loss function, problem (12) is equivalent to the convex optimization problem in (14). Further, specializing the result to a linear loss function  $\ell(\boldsymbol{\xi}, \boldsymbol{z}) = \boldsymbol{z}^\top \boldsymbol{\xi}$ , problem (12) is equivalent to the regularized risk-averse SAA problem in (15).  $\square$

It is interesting to see that for the linear case, the DRO problem reduces to the regularized SAA problem where the regularization penalty is controlled by the size  $\epsilon$  of the ambiguity set and that the type of penalty depends on the dual of the norm used to define the ambiguity set. To complement these results, Appendix A.8 provides reformulations of the distributionally robust newsvendor and regression problems as exponential cone programs.

Our next theorem formalizes that as the sample size  $N$  tends to infinity, the DRO value  $\rho_{\text{DRO}}$  with a properly chosen radius will converge to the true optimal risk  $\rho^*$  in probability. The proof follows from showing that for Lipschitz continuous (in  $\boldsymbol{z}$ ) loss functions,  $\rho_{\text{SAA}} \leq \rho_{\text{DRO}} \leq \rho_{\text{SAA}} + L\epsilon$  and using Proposition 1 that establishes that  $\rho_{\text{SAA}}$  converges to  $\rho^*$  at the rate  $\mathcal{O}(1/\sqrt{N})$  for locally Lipschitz continuous (in  $\boldsymbol{\xi}$ ) loss functions. Finally, choosing the radius to decay at the rate  $\mathcal{O}(1/\sqrt{N})$  preserves the rate of convergence of SAA.

**Theorem 3.** *Suppose that Assumption 3 holds. Then for any  $\gamma > 0$  and using  $\mathcal{B}_\infty(c/\sqrt{N})$ , for some  $c > 0$ , then there exists a constant  $A > 0$  such that*

$$\mathbb{P}\left(|\rho_{\text{DRO}} - \rho^*| \geq \frac{A}{\sqrt{N}\gamma\alpha \exp(\alpha\rho^*)} + \frac{c}{\sqrt{N}}\right) \leq \gamma,$$

as long as  $N$  is sufficiently large. Consequently,  $\rho_{\text{DRO}} \rightarrow \rho^*$  in probability.

While Theorem 3 provides a rate for  $\epsilon$  that ensures convergence of the DRO risk to the true optimal risk in probability, the values of the constants depend on the unknown underlying probability distribution. In practice,  $\epsilon$  needs to be estimated using CV. However, as discussed in the previous sections, estimating true risk from finite data is challenging. To address this, we next employ the bias-aware estimation procedure described in Section 4.

## 5.1 Bias-aware cross validation

A common approach to select the radius  $\epsilon$  of the ambiguity set is through  $K$ -fold CV. For each  $\epsilon \in \mathcal{E}$ ,  $K$ -fold CV aims to estimate the true performance of policy  $\boldsymbol{z}^*(\hat{\mathbb{P}}_N, \epsilon)$  resulting from problem (13), i.e.,  $\rho_{\mathbb{P}}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_N, \epsilon), \boldsymbol{\xi}))$ , and subsequently select the  $\epsilon$  that minimizes this risk. The approach divides the dataset into  $K$  folds. For each fold, we optimize the DRO model on  $K - 1$  folds and evaluate the solution's performance on the remaining fold, repeating this process for all folds. Specifically, for each candidate value of  $\epsilon$ , the model in problem (13) is solved using the training data  $\hat{\mathbb{P}}_{-k}^K$  from all folds except the  $k$ -th fold to determine  $\boldsymbol{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon)$  which is then evaluated on the validation data to obtain  $\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi}))$ , where  $\hat{\mathbb{P}}_k^K$  denotes the empirical distribution of scenarios in fold  $k$ . The resulting estimator for a given radius  $\epsilon$  is then given  $\rho_{k \sim U(K)}(\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})))$ , where  $U(K)$  is the uniform distribution over the set  $\{1, 2, \dots, K\}$ . Since the goal is to minimize risk, we choose  $\epsilon$  that minimizes the validation risk, i.e.,  $\epsilon^* = \operatorname{argmin}_{\epsilon \in \mathcal{E}} \rho_{k \sim U(K)}(\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})))$ . However, for each value of radius  $\epsilon$  and choice of  $K$ , the following proposition shows that the entropic risk estimator  $\rho_{k \sim U(K)}(\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})))$  based on the  $K$ -fold CV, underestimates the entropic risk of the policy constructed using  $N(1 - \frac{1}{K})$  data points.

**Proposition 3.** *Given  $\epsilon \in \mathcal{E}$ ,*

$$\mathbb{E}[\rho_{k \sim U(K)}(\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})))] < \rho_{\mathbb{P}}(\ell(\boldsymbol{z}^*(\hat{\mathbb{P}}_{N(1-\frac{1}{K})}, \epsilon), \boldsymbol{\xi})). \quad (16)$$

The proof of the above proposition follows from Jensen's inequality and tower property of entropic risk measure which states that  $\rho(\boldsymbol{\zeta}) = \rho(\rho(\boldsymbol{\zeta}_1 | \boldsymbol{\zeta}))$  for random variables  $\boldsymbol{\zeta}_1, \boldsymbol{\zeta}_2$ . Note that this property

is satisfied only by entropic risk measure in the family of law-invariant risk measures (Kupper and Schachermayer 2009). Notice that for large values of  $K < N$ ,  $N(1 - \frac{1}{K})$  approaches  $N$ , thus the right-hand-side of (16) mimics the performance of the solution that uses all  $N$  data points, that is,  $\rho_{\mathbb{P}}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_N, \epsilon), \boldsymbol{\xi}))$ . To mitigate the underestimation of the entropic risk, we propose using the

---

**Algorithm 4 K-fold cross validation.**


---

```

1: function K-FOLD CV( $K, \mathcal{D}_N, \epsilon$ )
2:    $\mathcal{S} \leftarrow \emptyset$ 
3:   for  $k \leftarrow 1$  to  $K$  do
4:      $\mathcal{D}_{-k} \leftarrow \mathcal{D}_N \setminus \mathcal{D}_k$  ▷ Training data (all samples except those in fold  $k$ )
5:      $\hat{\mathbb{P}}_{-k}^K \leftarrow$  empirical distribution of scenarios in  $\mathcal{D}_{-k}$ 
6:     Solve problem (13) with distribution  $\hat{\mathbb{P}}_{-k}^K$  and radius  $\epsilon$  to get  $\mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon)$ 
7:      $\mathcal{S} \leftarrow \mathcal{S} \cup \{\ell(\boldsymbol{\xi}, \mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon)) \mid \boldsymbol{\xi} \in \mathcal{D}_k\}$ 
8:   return  $\mathcal{S}, \rho_{k \sim U(K)}(\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})))$ 

```

---

bias-aware bootstrap procedure described in Algorithm 1 together with either Algorithm 2 based on entropic risk matching, or with Algorithm 3 based on extreme value theory. Algorithm 5 describes our proposed approach for selecting the optimal  $\epsilon$ , with Algorithm 4 describing the  $K$ -fold CV step. One can recover the traditional biased CV procedure by setting  $\delta = 0$  in line 4 of Algorithm 5.

A pictorial representation of the debiasing effect in optimization problems can be seen in Figure 4. If the aim of Example 3 is to select the project with the lowest estimated risk, then most methods would favor project 3. However, by accurately estimating and correcting for bias—using our proposed approaches—project 1 becomes the preferred choice. Similar behavior is also observed when choosing  $\epsilon$  using  $K$ -fold CV, where the problem parallels Example 3 in which the projects can be seen as corresponding to different regularization parameters  $\epsilon$ , with lower values of  $\epsilon$  representing riskier projects. As we will see in the following section, the solution based on our proposed approach significantly outperforms traditional CV procedure.

---

**Algorithm 5 Radius selection for DRO.**


---

```

1: function RADIUS TUNING( $\mathcal{D}_N, K, M$ )
2:   for  $\epsilon \in \mathcal{E}$  do
3:      $\mathcal{S}, \hat{\rho} \leftarrow$  K-fold CV( $K, \mathcal{D}_N, \epsilon$ )
4:      $\delta \leftarrow$  BootstrapBiasCorrection( $\mathcal{S}, M$ ) ▷ Algorithm 1
5:      $\rho(\epsilon) \leftarrow \hat{\rho} + \delta$ 
6:    $\epsilon^* \leftarrow \arg \min_{\epsilon \in \mathcal{E}} \rho(\epsilon)$ 

```

---

## 6 Distributionally robust insurance policy

The US National Flood Insurance Program (NFIP) provides flood coverage at subsidized premiums but faces significant challenges due to the large, correlated losses it insures against. These losses often result in claims exceeding the cumulative premiums collected over time (Marcoux and H Wagner 2023). Consequently, the NFIP currently operates with a deficit exceeding \$20 billion and is compelled to consider raising premiums (Marcoux and H Wagner 2023). However, higher premiums often deter households from purchasing coverage. This reluctance stems from how individuals perceive risk, which is frequently shaped by empirical losses rather than statistical estimates (Kousky and Cooke 2012). As a result, households tend to underestimate the risks associated with rare events. Demand for insurance, therefore, typically spikes only after catastrophic disasters (Gallagher 2014). In other words, individuals who have not experienced a catastrophic flood event are more likely to underestimate the associated risks. Surveys indicate that people exposed to flood risk but without firsthand experience of similar disasters often exhibit overly optimistic views about the threats posed by climate change. This optimism has been linked to houses in high flood-risk areas being overvalued by 6–9% (Bakkensen and Barrage 2022). Furthermore, NFIP premium subsidies reductions, combined with advances in risk

estimation and flood risk mapping, have been shown to decrease house prices in high-risk areas (Hino and Burke 2021). Such behavioral responses are not unique to flood insurance markets. Herrnstadt and Sweeney (2024) use a difference-in-differences method to show that the prices of houses within 500 feet of a gas pipeline in the Bay Area dropped by \$383 per household following the deadly 2010 pipeline explosion in San Francisco. Moreover, residents' perceptions of risk increased significantly above the empirical average for several years after the explosion. Interestingly, however, the prices of properties located 2,000 feet from the pipeline remained unaffected, despite being classified as high-risk. This discrepancy underscores the impact of firsthand experiences of catastrophic events on risk perception.

We consider an insurance pricing problem, with one risk-averse insurer and  $M$  representative risk-averse households. The proposed model can account for correlated losses and asymmetry in perception of risk measured by the empirical loss distributions at household level. Let  $\alpha_h$  denote the risk aversion of household  $h$ , and let  $\alpha_0$  represent the insurer's risk aversion parameter. The uncertain loss faced by household  $h$  is represented by  $\xi_h$ . The insurer offers a policy  $(z_h, \pi_h)$  to household  $h$ , where the indemnity function  $z_h \xi_h$  specifies the coverage provided for their loss  $\xi_h$ , and  $\pi_h$  is the corresponding premium paid by household  $h$ . Consequently, the net loss faced by household  $h$  under this policy is given by  $\pi_h + (1 - z_h)\xi_h$ . Let  $\hat{\mathbb{P}}_{h,N}$  be the empirical distribution of losses faced by household  $h$ . The insurer's demand response model assumes that household  $h$  will accept the policy  $(z_h, \pi_h)$  if the empirical entropic risk with insurance is less than the entropic risk without insurance. This condition is expressed by the following constraint:

$$\rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\pi_h + (1 - z_h)\xi_h) \leq \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\xi_h). \quad (17)$$

The insurer aims to minimize risk across the policies offered to all households. Accordingly, the insurer's true entropic risk is given by  $\rho_{\mathbb{P}}^{\alpha_0}(\mathbf{z}^\top \boldsymbol{\xi} - \mathbf{1}^\top \boldsymbol{\pi})$ , where  $\mathbf{1}$  is the vector of ones of the appropriate dimension. Since the true joint distribution,  $\mathbb{P}$ , of losses across all households with marginals  $\mathbb{P}_h$ , is unknown, the insurer replaces it with the empirical distribution  $\hat{\mathbb{P}}_N$  and solves the following optimization problem to determine the policies offered to the  $M$  households:

$$\begin{aligned} \min \quad & \rho_{\hat{\mathbb{P}}_N}^{\alpha_0}(\mathbf{z}^\top \boldsymbol{\xi} - \mathbf{1}^\top \boldsymbol{\pi}) \\ \text{s.t.} \quad & \boldsymbol{\pi} \in \mathbb{R}_+^M, \mathbf{z} \in [0, 1]^M \\ & \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\pi_h + (1 - z_h)\xi_h) \leq \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\xi_h) \quad \forall h \in [M]. \end{aligned} \quad (18)$$

As discussed in the previous section, decisions based on the empirical data can be optimistically biased. To address this, the insurer solves the following distributionally robust insurance pricing problem, which minimizes the worst-case entropic risk:

$$\begin{aligned} \min \quad & \sup_{\mathbb{Q} \in \mathcal{B}_\infty(\epsilon)} \rho_{\mathbb{Q}}^{\alpha_0}(\mathbf{z}^\top \boldsymbol{\xi} - \mathbf{1}^\top \boldsymbol{\pi}) \\ \text{s.t.} \quad & \boldsymbol{\pi} \in \mathbb{R}_+^M, \mathbf{z} \in [0, 1]^M \\ & \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\pi_h + (1 - z_h)\xi_h) \leq \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\xi_h) \quad \forall h \in [M], \end{aligned}$$

where  $\mathbb{Q}$  lies in the type- $\infty$  Wasserstein ambiguity set given in (11). Since the loss function is linear, it follows from Corollary 1 that the problem can be reformulated as the following regularized exponential cone program:

$$\begin{aligned} \min \quad & \rho_{\hat{\mathbb{P}}_N}^{\alpha_0}(\mathbf{z}^\top \boldsymbol{\xi} - \mathbf{1}^\top \boldsymbol{\pi}) + \epsilon \|\mathbf{z}\|_* \\ \text{s.t.} \quad & \boldsymbol{\pi} \in \mathbb{R}_+^M, \mathbf{z} \in [0, 1]^M \\ & \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\pi_h + (1 - z_h)\xi_h) \leq \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\xi_h) \quad \forall h \in [M]. \end{aligned} \quad (19)$$

The proposed setting is motivated by Bernard et al. (2020) who assumes that both the insurer and households are expected utility maximizers, with complete information on the true loss distributions for each household and their risk aversion parameters. Our approach relaxes the assumption of known loss distributions by providing only samples of the loss distribution to both the insurer and the households.



While Bernard et al. (2020) impose some additional assumptions to analytically characterize pricing and coverage decisions under partially correlated risks, our data-driven method formulates the problem as a tractable exponential cone program (18) which needs to be solved numerically. Additionally, the insurer’s robustified (DRO) problem (19) retains this tractability. At optimality, the constraints hold with equality due to the monotonicity of entropic risk measure, that is,  $\rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\pi_h + (1 - z_h)\xi_h) = \rho_{\hat{\mathbb{P}}_{h,N}}^{\alpha_h}(\xi_h)$  which can be equivalently written as:

$$\pi_h = \frac{1}{\alpha_h} \log \left( \frac{\mathbb{E}_{\hat{\mathbb{P}}_{h,N}}[\exp(\alpha_h \xi_h)]}{\mathbb{E}_{\hat{\mathbb{P}}_{h,N}}[\exp(\alpha_h (1 - z_h)\xi_h)]} \right). \quad (20)$$

In Example 1, we computed the optimal premium price under full coverage  $z_h = 1$  in which case the premium in (20) equals the entropic risk. Furthermore, the demand response model (17), which links premiums to coverage, accounts for the asymmetry in risk perception between households—who rely on  $\hat{\mathbb{P}}_{h,N}$ —and the insurer, who uses  $\hat{\mathbb{P}}_N$ . This flexible framework adapts to households’ evolving responses as additional information is incorporated through observed catastrophic events, thereby updating the empirical distributions. It is worth noting that our study could easily be adapted to accommodate alternative demand response models as long as the insurers valuations of the insurance is a concave function of coverage. Moreover, the regulatory constraints enforcing minimum coverage requirements for each household could easily be integrated.

## 6.1 Numerical experiments

The numerical experiments conducted in this section demonstrate the effectiveness of our proposed distributionally robust insurance pricing model under various conditions. Our main goal is to evaluate how different calibration methods for the radius  $\epsilon$  influence the insurer’s out-of-sample entropic risk and investigate the structure of the optimal policies  $(\mathbf{z}, \boldsymbol{\pi})$  offered to the households. Each household’s loss follows a Gamma distribution  $\Gamma(\kappa_h, \lambda_h)$ , with shape  $\kappa_h$  and scale  $\lambda_h$  parameters specific to each household  $h \in [M]$ . The correlation among the losses of different households is modelled using a Gaussian copula, with  $\Sigma = r\mathbf{1}\mathbf{1}^\top + (1 - r)I$ , where  $r$  controls the amount of correlation among the different households,  $\mathbf{1}$  is a vector of all ones, and  $I$  is the identity matrix.

In the following experiments, we consider  $M = 5$  households with risk aversion parameters  $\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5 = 2.9, 2.7, 2.5, 2.3, 2.1$  and set the insurer’s risk aversion parameter to  $\alpha_0 = 2$ . In Algorithm 5, we select the radius of the ambiguity set  $\epsilon$  from set  $\mathcal{E}$  which contains twenty equally-spaced values in the interval  $[0, 6]$ . To generate an instance, we sample  $N$  loss scenarios for each household and evaluate the out-of-sample performance, by generating  $10^6$  i.i.d. data points (test data). This procedure is repeated over 100 instances to obtain the statistics presented in the subsequent sections.

We consider five different calibration methods. Models **BS-Match** and **BS-EVT** use the 5-fold CV Algorithm 5 together with Algorithm 2 and Algorithm 3 within the bootstrapping procedure, respectively, i.e.,  $K = 5$ . **CV** model corresponds to the traditional 5-fold CV where we set  $\delta = 0$  in step 4 of Algorithm 5. The model labeled as **Oracle** uses the test data for calibrating the radius  $\epsilon$  at the validation step. Finally, model **SAA** solves problem (18) and does not involve any calibration.

All experiments described in the following sections were conducted in Python. The MOSEK 10.1 solver was used to solve exponential cone programs, while the entropic risk matching was performed on a GPU using the POT library (Flamary et al. 2021).

### 6.1.1 Case with mild correlation

In the first experiment, we examine the effect of sample size on the out-of-sample risk observed by the insurer. In the base scenario, all households have a common Gamma-distributed marginal loss distribution,  $\Gamma(10, 0.45)$ , with a correlation coefficient  $r = 0.5$ . This configuration allows us to focus on the effects of sample size  $N$  and correlation coefficient  $r$  on the insurer’s decisions, controlling for

variability from differing marginal distributions. Similar insights hold for cases with heterogeneous marginal distributions among households, discussed in Appendix B.6. The results are summarized in Figures 5–8. Figure 5a shows that both BS-Match and BS-EVT consistently outperform CV and SAA across different sample sizes, while achieving an out-of-sample entropic risk similar to the oracle-based calibration. This can be explained by looking at the ambiguity radius  $\epsilon^*$  chosen by each method. In Figure 5b, we observe that CV typically chooses  $\epsilon^*$  values that are significantly lower than the optimal choice, while the BS-Match and BS-EVT choices are closer to optimal. This discrepancy results from CV’s estimation procedure, which underestimates the true entropic risk for each  $\epsilon$  (as discussed in Section 4), leading to select an overly optimistic  $\epsilon^*$  in step 6 of Algorithm 5. This can be seen in Figure 6 where we plot the variation in the entropic risk estimator with the radius  $\epsilon$  for each model with  $N = 1000$  (similar plots are obtained in Appendix B.7 for  $N \in \{500, 5000, 10000\}$ ). In contrast, BS-Match and BS-EVT better estimate the trend in the variation of true entropic risk with  $\epsilon$ , thus making a more informed choice for  $\epsilon^*$ .

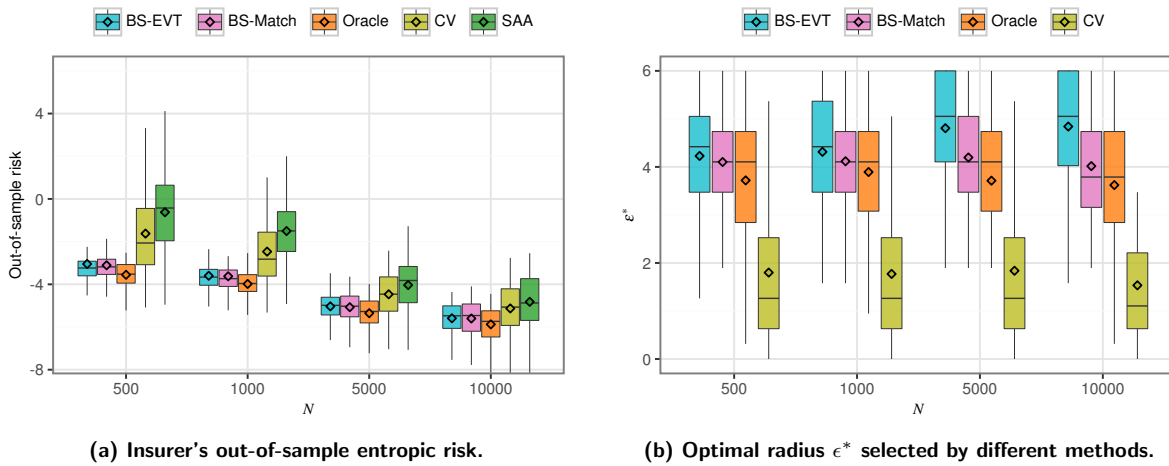


Figure 5: Comparison of the effects of training sample size  $N$  on out-of-sample entropic risk (left) and optimal radius  $\epsilon^*$  (right). Boxplots present the statistics after 100 resampling of datasets, and diamonds presents the mean for each  $N$ .




Figure 6: Statistics of entropic risk estimators for different radius and  $N = 1000$  after 100 resampling of datasets.

Figure 7a depicts the estimation of the optimal entropic risk computed by each method. The results have a similar interpretation as Figure 4 in Section 4.3. We observe a significant underestimation of the true entropic by CV and SAA, while the estimation of the BS-Match and BS-EVT stay close to the estimation of the optimal calibration Oracle. We observe that as the sample size increases, the optimal risk decreases and so is the case of each method’s optimal risk estimation. This is because it is optimal for households to pay higher premiums due to the increase in their estimate of the risk of their respective loss.



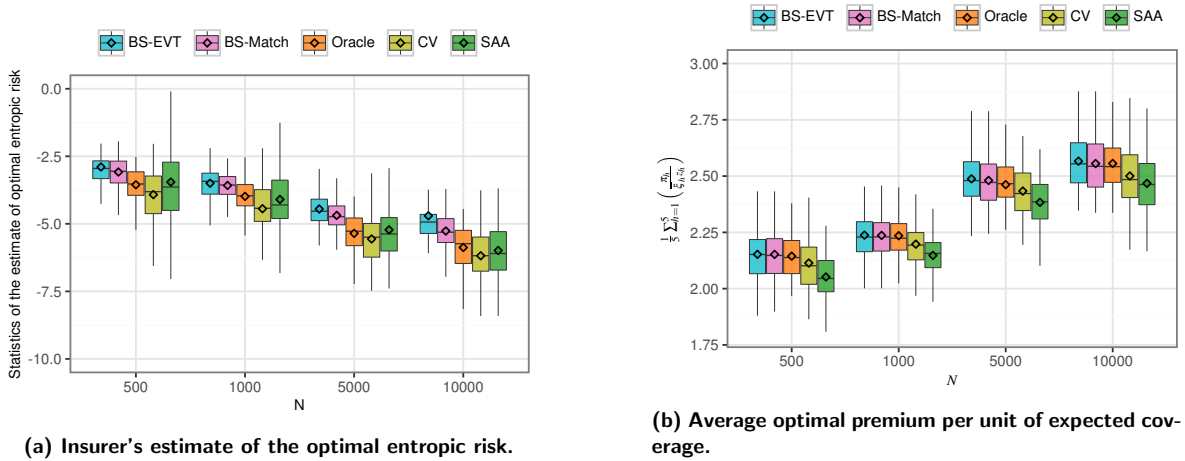


Figure 7: Comparison of the effects of number of training samples  $N$  on insurer's estimate of optimal entropic risk and average premium per unit of expected coverage.

Consequently, the insurer can charge higher premiums per unit expected coverage as a function of  $N$ , see Figure 7b where the average premium per unit expected coverage across the 5 households is given  $(1/5) \sum_{h=1}^5 \pi_h / (\bar{\xi}_h z_h)$  with  $\bar{\xi}_h$  as the expected loss of household  $h$ . We observe an increase in this ratio as  $N$  increases, which can be explained from the observation both the insurer and households become more capable of accurately estimating their true risk, thus enabling the insurer to extract higher premiums from households for the same coverage level. Moreover, CV and SAA charge lower premium per unit of expected coverage compared to our proposed approaches because they underestimate the risk. Namely, both methods are overly optimistic regarding how to correct the estimation error due to sampling, effectively using an  $\epsilon$  that is too low. To illustrate the variation in the households' willingness to pay for insurance as a function of number of training samples, we use as a proxy, the proportion of instances where the premium is at least twice the expected coverage. For BS-Match, Figure 8 illustrates that highly risk-averse households pay higher premiums per unit of expected coverage more frequently even when the number of training samples is low. Additionally, as  $N$  increases, the proportion approaches 100% for all households.




Figure 8: Effect of the number of training samples on the proportion of instances where premium exceeds twice the expected coverage for BS-Match. Households become less risk averse as we go from the left of the panel to the right one.

### 6.1.2 Effect of correlations

In the second experiment, we fix the sample size at  $N = 1000$ . Similar to the previous experiment, we assume that all households share a common marginal loss distribution modeled by a Gamma distribution,  $\Gamma(10, 0.45)$ . However, in this experiment, we vary the parameter  $r$ , which controls the pairwise correlation of losses between households. The correlation coefficient ranges from 0 (independent losses) to 1 (comonotone losses). Figures 9–12 summarize the results. First, note that similar insights about the effectiveness of our proposed approaches can also be observed in this setting. Indeed, Figure 9b shows that both CV and SAA are over-optimistic, leading them to select lower  $\epsilon^*$  values than those chosen by our proposed approaches. With regards to the behavior in terms of  $r$ , Figure 9a demonstrates

that the out-of-sample entropic risk initially increases with the correlation coefficient  $r$ , but eventually stabilizes. This trend is intuitive: higher correlation among households' losses means that extreme loss events are more likely to occur simultaneously, increasing the insurer's risk exposure.

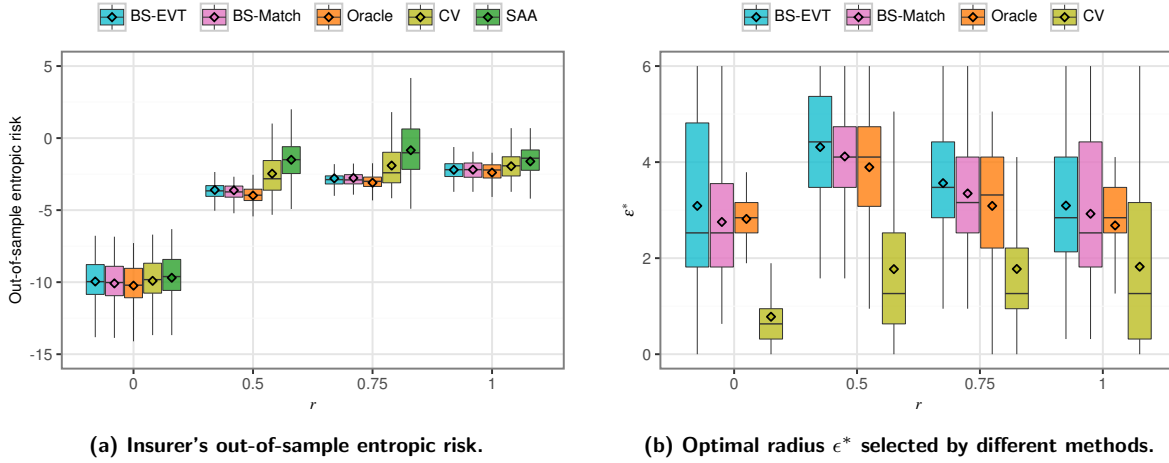


Figure 9: Comparison of the effects of correlation coefficient  $r$  on out-of-sample entropic risk (left) and optimal radius  $\epsilon^*$  (right).

Figure 10a shows the estimates of the optimal entropic risk produced by each model. BS-EVT and BS-Match overestimate the optimal entropic risk of the insurer compared to Oracle at low correlation levels. Figure 10b shows that the average optimal premium per unit of expected coverage also increases with the correlation coefficient. This reflects the insurer's response to higher risk by charging higher premiums to compensate for the increased likelihood of large, simultaneous payouts. However, there is a diminishing return effect; beyond a certain point, further increases in correlation do not lead to significantly higher premiums per unit of expected coverage. As the correlation between household losses increases, the benefits of risk pooling diminish, so the insurer reduces coverage levels significantly to reduce the risk exposure, see Figure 11a and 11b. While Figures 11a and 11b show that households receive more coverage with SAA than BS-Match, this high coverage exposes the insurer to higher risks than the optimal coverage in the case of highly correlated losses.

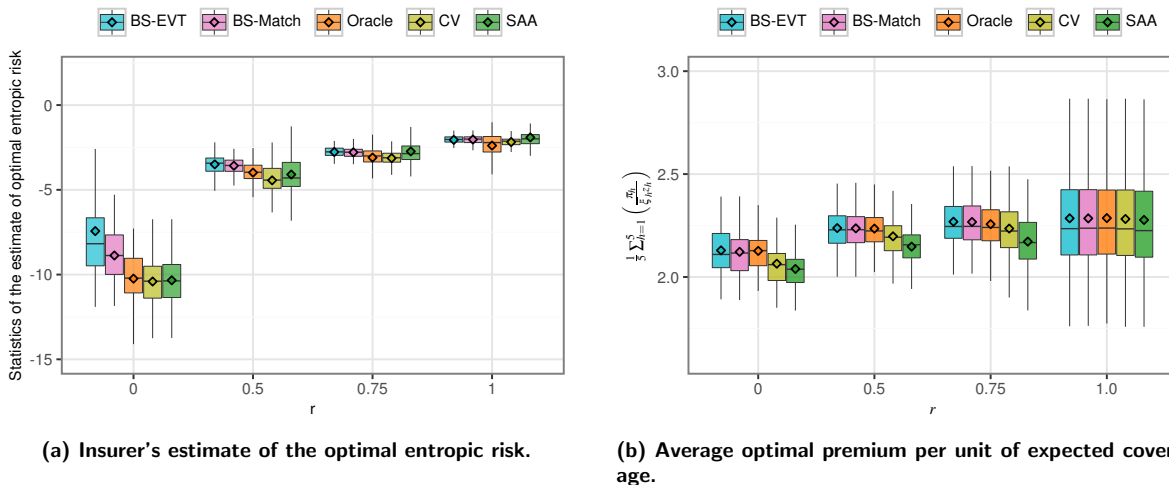


Figure 10: Comparison of the effects of correlation coefficient  $r$  on estimate of optimal entropic risk and premium per unit expected coverage.

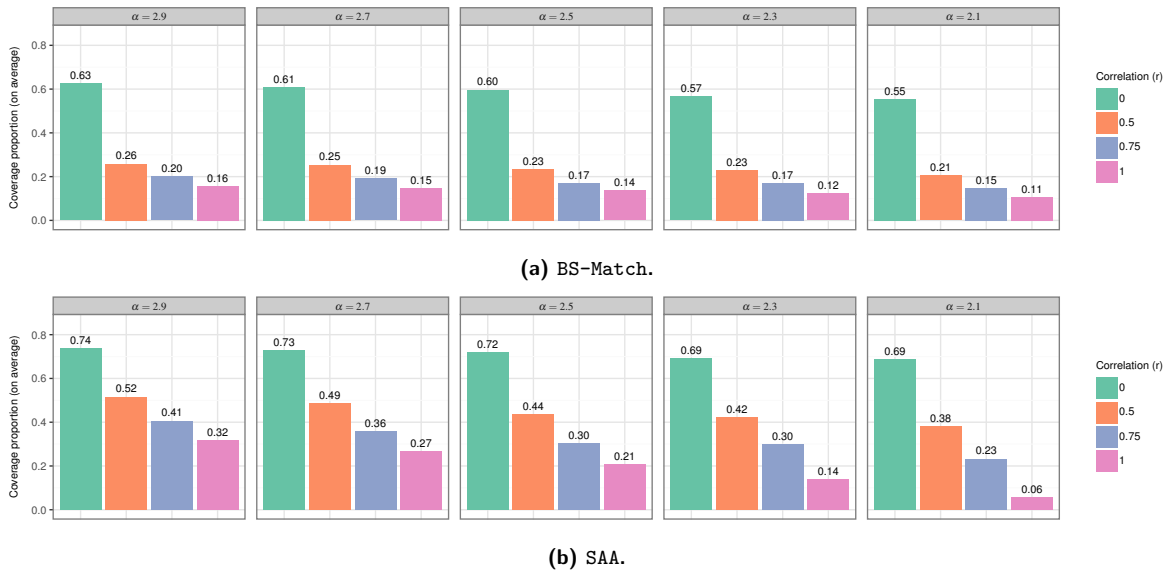


Figure 11: Effect of correlation on the coverage proportion offered to households averaged across 100 instances. Each panel represents a different household. Households become less risk-averse as we go from the left panel to the right one.

Figure 12 demonstrates that as the correlation  $r$  across households increases, the proportion of instances with premiums exceeding twice the expected coverage also increase, and this effect is more pronounced for more risk-averse households. The high risk-averse households secure greater coverage by paying high premiums to reduce their risk exposure.




Figure 12: Effect of correlation  $r$  on proportion of instances where premium exceeds two times the expected coverage. Each panel represents a different household. Households become less risk-averse as we go from the left panel to the right one.

## 7 Conclusions

In this paper, we propose two practical approaches to mitigate the bias inherent in the empirical entropic risk estimator, which arises when the empirical distribution is used in place of the true distribution in entropic risk calculations. We propose to correct the bias using bootstrapping and prove that such a bias-corrected entropic risk estimator converges almost surely to the true risk. However, when bootstrapping is based on a GMM fitted using MLE, the bias-corrected estimator continues to underestimate the true risk with finite samples. Our proposed approaches mitigate this underestimation problem by estimating the bias in samples drawn from a “bias-aware” GMM fitted to the data.

To address the associated optimistic bias when optimizing the entropic risk measure, we propose to use type- $\infty$  Wasserstein ambiguity sets in DRO models. For piecewise-concave loss functions, we provide equivalent reformulations of DRO problems as finite-dimensional convex optimization problems

using Fenchel duality. We show that traditional K-fold CV used to tune the radius of the ambiguity set can result in the underestimation of risk and the proposed approaches can mitigate the underestimation problem. Specifically, we introduce a distributionally robust insurance pricing problem with a risk-averse insurer and risk-averse homeowners whose demand for insurance depends on their own empirical distribution of losses. Numerical experiments demonstrate that our bias mitigation procedures help identify the appropriate radius for the ambiguity set, resulting in premium pricing and coverage policies that achieve lower out-of-sample entropic risks when compared to CV methods that do not account for bias in the validation scenarios. The numerical results also validate the effectiveness of our DRO framework in improving insurance pricing strategies under uncertainty. The experiments underscore the critical role of accurate risk estimation in insurance pricing, especially for catastrophic events exacerbated by climate change. Insurers must consider not only the empirical data but also the uncertainty inherent in limited samples and the potential for extreme losses. Our findings suggest that incorporating distributional robustness into the pricing model allows insurers to balance the trade-off between premiums charged and coverage offered for potential losses. This approach can help mitigate the financial vulnerability exemplified by programs like the NFIP, which faces substantial deficits due to under-priced premiums and low uptake.

There are several avenues for extending this research. First, the proposed bias correction procedures can be adapted to estimate other risk measures, such as Conditional Value at Risk. Importantly, our methods do not depend heavily on the specific form of the entropic risk measure, aside from the fact that tail scenarios exert disproportionate influence on it. Exploring the effectiveness of these bias corrections in multi-stage settings, where entropic risk measures are widely used, such as control theory and reinforcement learning, would be valuable. Furthermore, investigating the potential benefits of bias correction in other contexts—such as portfolio optimization and inventory management—represents a promising direction for future research. Finally, leveraging advancements in solving exponential cone programs (see Ye and Xie 2021, Chen et al. 2024a) to address large-scale entropic risk minimization problems, potentially with integer constraints, offers a valuable opportunity for enhancing computational efficiency.

## A Additional proofs and results

### A.1 Proof of Lemma 1

**Proof.** We have that:

$$\begin{aligned}
\mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))] &= -\int_{-\infty}^0 \mathbb{P}(\exp(\alpha\ell(\boldsymbol{\xi})) \leq x)dx + \int_0^{\infty} \mathbb{P}(\exp(\alpha\ell(\boldsymbol{\xi})) > x)dx \\
&\leq \int_0^{\infty} \mathbb{P}(\exp(\alpha\ell(\boldsymbol{\xi})) > x)dx \\
&= \int_0^{\infty} \mathbb{P}(\exp(\alpha\ell(\boldsymbol{\xi})) > \exp(\alpha y))\alpha \exp(\alpha y)dy \\
&= \int_0^{\infty} \mathbb{P}(\ell(\boldsymbol{\xi}) > y)\alpha \exp(\alpha y)dy \\
&\leq \int_0^{\infty} \mathbb{P}(|\ell(\boldsymbol{\xi})| > y)\alpha \exp(\alpha y)dy \\
&\leq \alpha \int_0^{\infty} G \exp(-(C-1)\alpha y)dy = \frac{G}{C-1},
\end{aligned}$$

where first equality follows from representing expectation of a random variable using its cumulative distribution function (cdf), the last inequality follows from Assumption 1 and  $C > 2$  is used to obtain the final result.

From Jensen's inequality, we also have that:

$$\begin{aligned}
\log(\mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))]) &\geq \mathbb{E}[\log(\exp(\alpha\ell(\boldsymbol{\xi})))] \\
&= \mathbb{E}[\alpha\ell(\boldsymbol{\xi})] \\
&= -\int_{-\infty}^0 \mathbb{P}(\alpha\ell(\boldsymbol{\xi}) \leq x)dx + \int_0^{\infty} \mathbb{P}(\alpha\ell(\boldsymbol{\xi}) > x)dx \\
&\geq -\int_{-\infty}^0 \mathbb{P}(\alpha\ell(\boldsymbol{\xi}) \leq x)dx \\
&\geq -\int_{-\infty}^0 \mathbb{P}(|\alpha\ell(\boldsymbol{\xi})| \geq -x)dx \\
&= -\int_0^{\infty} \mathbb{P}(|\alpha\ell(\boldsymbol{\xi})| \geq y)dy \\
&\geq -\int_0^{\infty} \mathbb{P}(|\ell(\boldsymbol{\xi})| > y/(2\alpha))dy \\
&\geq -\int_0^{\infty} G \exp(-Cy/2)dy \\
&= -2G/C,
\end{aligned}$$

where the last inequality follows from Assumption 1. Hence, we have that  $\mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))] \geq \exp(-2G/C)$ .

Similarly,

$$\begin{aligned}
\mathbb{E}[\exp(2\alpha\ell(\boldsymbol{\xi}))] &\leq \int_0^{\infty} \mathbb{P}(\exp(2\alpha\ell(\boldsymbol{\xi})) > \exp(2\alpha y))2\alpha \exp(2\alpha y)dy \\
&\leq \int_0^{\infty} \mathbb{P}(|\ell(\boldsymbol{\xi})| > y)2\alpha \exp(2\alpha y)dy \\
&\leq 2\alpha \int_0^{\infty} G \exp(-(C-2)\alpha y)dy = \frac{2G}{C-2},
\end{aligned}$$

where we used  $C > 2$  to obtain the last equality, see Assumption 1. Hence,

$$0 \leq \text{Var}(\exp(\alpha\ell(\boldsymbol{\xi}))) = \mathbb{E}[\exp(2\alpha\ell(\boldsymbol{\xi}))] - \mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))]^2 \leq \mathbb{E}[\exp(2\alpha\ell(\boldsymbol{\xi}))] \leq \frac{2G}{C-2}.$$

□

## A.2 Proof of Theorem 1

**Proof.** We want to show that  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) + \delta_N(\mathbb{Q}_N) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$  almost surely. To do so, we will show that both  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$  and  $\delta_N(\mathbb{Q}_N) \rightarrow 0$  almost surely. Indeed, if both  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$  and  $\delta_N(\mathbb{Q}_N) \rightarrow 0$  almost surely, then we can use the fact that the sum of two convergent sequences converges to the sum of their limits to conclude that:

$$\begin{aligned}
\mathbb{P}(\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) + \delta_N(\mathbb{Q}_N) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))) &\geq \mathbb{P}(\{\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))\} \cap \{\delta_N(\mathbb{Q}_N) \rightarrow 0\}) \\
&= 1 - \mathbb{P}(\{\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \not\rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))\} \cup \{\delta_N(\mathbb{Q}_N) \not\rightarrow 0\}) \\
&\geq 1 - (1 - \mathbb{P}(\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})))) - (1 - \mathbb{P}(\delta_N(\mathbb{Q}_N) \rightarrow 0)) \\
&= 1,
\end{aligned}$$

where the second inequality follows from the union bound. **Step 1:**  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$  **almost surely.** Given that  $\mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))]$  is finite (see Lemma 1) and each  $\boldsymbol{\xi}_i$  is i.i.d., the strong law of large numbers tells us that:

$$\frac{1}{N} \sum_{i=1}^N \exp(\alpha\ell(\hat{\boldsymbol{\xi}}_i)) \rightarrow \mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))] \text{ almost surely.} \quad (21)$$

Since logarithm function is continuous over the strictly positive values, using the continuous mapping theorem (Van der Vaart 2000), we obtain:

$$(21) \implies \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha \ell(\hat{\boldsymbol{\xi}}_i)) \right) \rightarrow \frac{1}{\alpha} \log (\mathbb{E}[\exp(\alpha \ell(\boldsymbol{\xi}))]) \text{ almost surely.}$$

Hence,  $\rho_{\hat{\mathbb{P}}_N}(\ell(\boldsymbol{\xi})) \rightarrow \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi}))$  almost surely.

**Step 2:**  $\delta_N(\mathbb{Q}_N) \rightarrow 0$  **almost surely.** We will show that  $\delta_N(\mathbb{Q}_N) \rightarrow 0$  almost surely by showing that  $\exp(-\alpha \delta_N(\mathbb{Q}_N)) \rightarrow 1$  almost surely.

Let's consider any sequence  $\{\bar{\mathbb{Q}}_N\}_{N=1}^{\infty}$  that is uniformly exponentially bounded, with each element of the sequence satisfying Assumption 2. For each  $N$ , we have:

$$\begin{aligned} \exp(-\alpha \delta_N(\bar{\mathbb{Q}}_N)) &= \exp(\alpha \text{median}(\rho_{\hat{\mathbb{Q}}_{N,N}}(\zeta) - \rho_{\bar{\mathbb{Q}}_N}(\zeta) | \bar{\mathbb{Q}}_N)) \\ &= \text{median}(\exp(\alpha \rho_{\hat{\mathbb{Q}}_{N,N}}(\zeta) - \alpha \rho_{\bar{\mathbb{Q}}_N}(\zeta)) | \bar{\mathbb{Q}}_N) \\ &= \text{median} \left( \exp \left( \log \left( \frac{(1/N) \sum_{i=1}^N \exp(\alpha \zeta_i)}{\mathbb{E}_{\bar{\mathbb{Q}}_N}[\exp(\alpha \zeta)]} \right) \right) \middle| \bar{\mathbb{Q}}_N \right) \\ &= \text{median}[X_N | \bar{\mathbb{Q}}_N], \end{aligned}$$

where  $X_N := \frac{(1/N) \sum_{i=1}^N \exp(\alpha \zeta_i)}{\mathbb{E}_{\bar{\mathbb{Q}}_N}[\exp(\alpha \zeta)]}$ . The second equality follows by the monotonicity of the exponential function, and the third equality follows by the definition of the entropic risk and the properties of the logarithm function. Note that  $\{\zeta_i\}_{i=1}^N$  are drawn i.i.d. from  $\bar{\mathbb{Q}}_N$ .

To analyze the  $\text{median}[X_N | \bar{\mathbb{Q}}_N]$ , we first compute the mean and variance of  $X_N$ . It is easy to see that  $\mathbb{E}_{\bar{\mathbb{Q}}_N}[X_N] = 1$ , while the variance of  $X_N$  can be bounded as follows:

$$\begin{aligned} \text{Var}_{\bar{\mathbb{Q}}_N}(X_N) &= \frac{\text{Var}_{\bar{\mathbb{Q}}_N} \left( \sum_{i=1}^N \exp(\alpha \zeta_i) \right)}{(N \mathbb{E}_{\bar{\mathbb{Q}}_N}[\exp(\alpha \zeta)])^2} \\ &= \frac{\text{Var}_{\bar{\mathbb{Q}}_N}(\exp(\alpha \zeta))}{N (\mathbb{E}_{\bar{\mathbb{Q}}_N}[\exp(\alpha \zeta)])^2} \\ &\leq \frac{2G \exp(\frac{4G}{C})}{N(C-2)}, \end{aligned}$$

where the second equality follows from the fact that  $\{\zeta_i\}_{i=1}^N$  are i.i.d. The inequality follows since  $\{\mathbb{Q}_N\}_{N=1}^{\infty}$  satisfy Assumption 2, thus Lemma 1 provides bounds for both  $\text{Var}_{\bar{\mathbb{Q}}_N}(\exp(\alpha \zeta))$  and  $\mathbb{E}_{\bar{\mathbb{Q}}_N}[\exp(\alpha \zeta)]$ , resulting in the bound  $\frac{\text{Var}_{\bar{\mathbb{Q}}_N}(\exp(\alpha \zeta))}{(\mathbb{E}_{\bar{\mathbb{Q}}_N}[\exp(\alpha \zeta)])^2} \leq \frac{2G \exp(\frac{4G}{C})}{C-2}$ .

We next show that  $\text{median}[X_N | \bar{\mathbb{Q}}_N]$  is bounded. To this end, consider the Chebyshev's inequality:

$$\mathbb{P}_{\bar{\mathbb{Q}}_N} \left( |X_N - 1| \geq k \sqrt{\text{Var}_{\bar{\mathbb{Q}}_N}(X_N)} \right) \leq \frac{1}{k^2}.$$

Substituting the bound for the  $\text{Var}_{\bar{\mathbb{Q}}_N}(X_N)$  and setting  $k = 2$ , which implies an upper tail probability bound of 25%, results in

$$\mathbb{P}_{\bar{\mathbb{Q}}_N} \left( |X_N - 1| \geq 2\Delta/\sqrt{N} \right) \leq \frac{1}{4},$$

where  $\Delta := \sqrt{2G \exp(\frac{4G}{C})/\sqrt{C-2}}$ . Thus, we conclude that  $\text{median}[X_N | \bar{\mathbb{Q}}_N] \in [1 - 2\Delta/\sqrt{N}, 1 + 2\Delta/\sqrt{N}]$  since otherwise it would imply that 50% of the probability is outside this interval (either on the right or the left), which would contradict the fact that the total probability outside the interval is below 1/4.

Finally, we show that as  $N$  tends to infinity,  $\text{median}[X_N|\bar{\mathbb{Q}}_N]$  converges to 1 almost surely. For any  $\epsilon > 0$ , there exists an  $N_0 = 4\Delta^2/\epsilon^2$  such that for all  $N > N_0$ ,  $|\text{median}[X_N|\bar{\mathbb{Q}}_N] - 1| \leq \epsilon$  which implies  $\lim_{N \rightarrow \infty} \text{median}[X_N|\bar{\mathbb{Q}}_N] = 1$ .

We conclude from the above analysis that any sequence  $\{\bar{\mathbb{Q}}_N\}_{N=1}^\infty$  that is uniformly exponentially bounded as prescribed in Assumption 2 must be such that  $\exp(-\alpha\delta_N(\bar{\mathbb{Q}}_N)) = \text{median}[X_N|\bar{\mathbb{Q}}_N]$  converges to 1 almost surely. Since Assumption 2 imposes that such sequences are almost surely obtained, we must have that  $\delta_N(\bar{\mathbb{Q}}_N) \rightarrow 0$  almost surely. □

### A.3 Proof of Proposition 1

**Proof.** We will show that  $\rho_{\text{SAA}}$  converges to  $\rho^*$  by demonstrating that  $\rho_{\hat{\mathbb{P}}_N}(\ell(\mathbf{z}, \boldsymbol{\xi}))$  converges to  $\rho_{\mathbb{P}}(\ell(\mathbf{z}, \boldsymbol{\xi}))$  for all  $\mathbf{z} \in \mathcal{Z}$ , i.e., a uniform rate of convergence. The proof is conducted in two steps. In the first step, we apply Jiang et al. (2020, Theorem 3.2) to show that the empirical exponential utility,  $\mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))]$ , converges uniformly to the true exponential utility,  $\mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))]$ . In the second step, we leverage the properties of the logarithm function to extend this uniform convergence from the empirical exponential utility to the entropic risk measure.

To apply Jiang et al. (2020, Theorem 3.2), we need to verify several properties for the exponential utility minimization problem. It is easy to see that (i)  $\mathcal{Z}$  is compact by Assumption 3(A.1); (ii) the exponential utility  $\mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))]$  is continuous in  $\mathbf{z}$  and (iii) the empirical distribution  $\hat{\mathbb{P}}_N$  is constructed using  $N$  i.i.d samples from  $\mathbb{P}$ ; (iv) Assumption 3(A.5) implies that  $\ell(\mathbf{z}, \boldsymbol{\xi})$  satisfies Assumption 1 for all  $\mathbf{z} \in \mathcal{Z}$  so that Lemma 1 confirms that  $\mathbb{E}[(\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi})))^2] < \infty$  for all  $\mathbf{z} \in \mathcal{Z}$ . Lastly, (v) the following inequalities verify that  $\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))$  is H-calm from above with Lipschitz modulus  $\kappa(\boldsymbol{\xi}) := \alpha L(\boldsymbol{\xi}) \exp(\alpha\bar{L}(\boldsymbol{\xi}))$  and order one:

$$\begin{aligned} \exp(\alpha\ell(\mathbf{z}_2, \boldsymbol{\xi})) - \exp(\alpha\ell(\mathbf{z}_1, \boldsymbol{\xi})) &\leq (\partial_{\mathbf{z}} \exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))|_{\mathbf{z}=\mathbf{z}_2})^\top (\mathbf{z}_2 - \mathbf{z}_1) \\ &= \alpha \exp(\alpha\ell(\mathbf{z}_2, \boldsymbol{\xi})) (\partial_{\mathbf{z}} \ell(\mathbf{z}_2, \boldsymbol{\xi}))^\top (\mathbf{z}_2 - \mathbf{z}_1) \\ &\leq \alpha \exp(\alpha\ell(\mathbf{z}_2, \boldsymbol{\xi})) \|\partial_{\mathbf{z}} \ell(\mathbf{z}_2, \boldsymbol{\xi})\|_* \|\mathbf{z}_2 - \mathbf{z}_1\| \\ &\leq \alpha \exp(\alpha\bar{L}(\boldsymbol{\xi})) L(\boldsymbol{\xi}) \|\mathbf{z}_2 - \mathbf{z}_1\|. \end{aligned}$$

The first inequality follows from the convexity of the exponential utility, where  $\partial_{\mathbf{z}}$  denotes the subgradient operator with respect to  $\mathbf{z}$ , and the equality follows by applying the chain rule. The second inequality follows from the definition of the dual norm. Since by Assumption 3  $\ell(\mathbf{z}, \boldsymbol{\xi})$  is locally Lipschitz continuous with Lipschitz constant  $L(\boldsymbol{\xi})$ , and the dual norm of the subgradient of  $\ell(\mathbf{z}, \boldsymbol{\xi})$  is bounded by the Lipschitz constant  $L(\boldsymbol{\xi})$  (Shalev-Shwartz et al. 2012, Lemma 2.6), we obtain the final inequality. Finally, we can confirm that the second moment of  $\kappa(\boldsymbol{\xi})$  is bounded. Namely, letting  $0 < \varsigma < (C/2) - 1$ , based on Hölder's inequality, we have that:

$$\mathbb{E}[\kappa(\boldsymbol{\xi})^2] = \alpha^2 \mathbb{E}[L(\boldsymbol{\xi})^2 \exp(2\alpha\bar{L}(\boldsymbol{\xi}))] \leq \alpha^2 \mathbb{E}[L(\boldsymbol{\xi})^{2(1+1/\varsigma)}]^{1/(1+\varsigma)} \mathbb{E}[\exp(2(1+\varsigma)\alpha\bar{L}(\boldsymbol{\xi}))]^{1/(1+\varsigma)}$$

We can further show using Assumption 3(A.4) that :

$$\begin{aligned} \mathbb{E}[\exp(2(1+\varsigma)\alpha\bar{L}(\boldsymbol{\xi}))] &\leq \int_0^\infty \mathbb{P}(\exp(2(1+\varsigma)\alpha\bar{L}(\boldsymbol{\xi})) > x) dx \\ &= \int_0^\infty \mathbb{P}(\exp(2(1+\varsigma)\alpha\bar{L}(\boldsymbol{\xi})) > \exp(2(1+\varsigma)\alpha y)) 2(1+\varsigma)\alpha \exp(2(1+\varsigma)\alpha y) dy \\ &= \int_0^\infty \mathbb{P}(\bar{L}(\boldsymbol{\xi}) > y) 2(1+\varsigma)\alpha \exp(2(1+\varsigma)\alpha y) dy \\ &\leq 2(1+\varsigma)\alpha \int_0^\infty G \exp(-(C - 2(1+\varsigma))\alpha y) dy = \frac{2G(1+\varsigma)}{C - 2(1+\varsigma)} < \infty, \end{aligned}$$


Moreover,  $\mathbb{E}[L(\boldsymbol{\xi})^{2(1+1/\varsigma)}]$  is finite based on Assumption 3(A.5), which allows us to conclude that  $\mathbb{E}[\kappa(\boldsymbol{\xi})^2] < \infty$ .

We now have all the components to apply Jiang et al. (2020, Theorem 3.2) which states that for any  $\gamma > 0$ , there exists an  $A > 0$ , independent of  $N$ , such that:

$$\mathbb{P}\left(\sup_{\mathbf{z} \in \mathcal{Z}} \left| \mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))] - \mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))] \right| \geq A/\sqrt{N\gamma}\right) \leq \gamma, \quad (22)$$

for sufficiently large  $N$ , thus concluding that the empirical exponential utility converges uniformly to the true exponential utility, with a rate of convergence of  $\mathcal{O}(1/\sqrt{N})$ .

Next, we show that empirical risk converges uniformly to the true risk by exploiting the properties of the log function, namely that  $\log(1 + \epsilon) \leq \epsilon$  for all  $\epsilon \geq 0$ , and  $\log(1 - \epsilon) \geq -\epsilon/(1 - 1/e)$  for all  $\epsilon \in [0, 1 - 1/e]$ . This result is summarized in Lemma 3 and can be pictorially verified in Figure 13.



(a)  $\log(1 + \epsilon) \leq \epsilon, \forall \epsilon \geq 0$ .

(b)  $\log(1 - \epsilon) \geq -\epsilon/(1 - 1/e) \forall \epsilon \in [0, 1 - 1/e]$ .

Figure 13: Plot of the inequalities that the log function satisfies around 0.

In the subsequent steps of the proof, we consider  $N$  to be large enough such that  $\frac{A}{\sqrt{N\gamma} \exp(\alpha \rho^*)} \leq 1 - 1/e$ , with  $\rho^* := \min_{\mathbf{z} \in \mathcal{Z}} (1/\alpha) \log(\mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))])$  and focus on scenarios for which the absolute difference between the empirical utility and true utility in (22) is upper bounded by  $A/(\sqrt{N\gamma})$  for all  $\mathbf{z} \in \mathcal{Z}$ . Specifically, we will show that it implies that the absolute difference between the empirical risk and true risk is upper bounded by  $A/(\alpha\sqrt{N\gamma} \exp(\alpha \rho^*))$  for all  $\mathbf{z} \in \mathcal{Z}$ . To simplify notation, we further denote by  $\hat{u}_N(\mathbf{z}) := \mathbb{E}_{\hat{\mathbb{P}}_N}(\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi})))$ ,  $u(\mathbf{z}) := \mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))]$ ,  $u^* := \min_{\mathbf{z} \in \mathcal{Z}} \mathbb{E}_{\mathbb{P}}(\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))) = \exp(\alpha \rho^*)$ . For any  $\gamma > 0$ , from (22), we have:

$$\begin{aligned} & |\hat{u}_N(\mathbf{z}) - u(\mathbf{z})| < \frac{A}{\sqrt{N\gamma}} & \forall \mathbf{z} \in \mathcal{Z} \\ \Rightarrow & \frac{|\hat{u}_N(\mathbf{z}) - u(\mathbf{z})|}{u^*} < \frac{A}{\sqrt{N\gamma} \exp(\alpha \rho^*)} & \forall \mathbf{z} \in \mathcal{Z} \\ \Rightarrow & \frac{|\hat{u}_N(\mathbf{z}) - u(\mathbf{z})|}{u(\mathbf{z})} < \frac{A}{\sqrt{N\gamma} \exp(\alpha \rho^*)} & \forall \mathbf{z} \in \mathcal{Z} \\ \Rightarrow & \frac{\hat{u}_N(\mathbf{z})}{u(\mathbf{z})} < 1 + \frac{A}{\sqrt{N\gamma} \exp(\alpha \rho^*)} & \forall \mathbf{z} \in \mathcal{Z} \\ \Rightarrow & \frac{1}{\alpha} \log\left(\frac{\hat{u}_N(\mathbf{z})}{u(\mathbf{z})}\right) < \frac{1}{\alpha} \log\left(1 + \frac{A}{\sqrt{N\gamma} \exp(\alpha \rho^*)}\right) & \forall \mathbf{z} \in \mathcal{Z}, \end{aligned}$$

where in the first implication we divided both sides by  $u^*$ , while the second implication follows since  $u^* \leq \min_{\mathbf{z} \in \mathcal{Z}} u(\mathbf{z})$ . The last implication follows by applying the logarithm on both sides of the inequality and dividing by  $\alpha$ . The last expression can be further simplified since by the properties of the logarithm function, we have  $\log(1 + \epsilon) \leq \epsilon$  for all  $\epsilon \geq 0$ , hence:

$$\frac{1}{\alpha} \log(\hat{u}_N(\mathbf{z})) - \frac{1}{\alpha} \log(u(\mathbf{z})) < \frac{A}{\alpha\sqrt{N\gamma} \exp(\alpha \rho^*)} < \frac{A}{(1 - 1/e)\alpha\sqrt{N\gamma} \exp(\alpha \rho^*)}, \forall \mathbf{z} \in \mathcal{Z}. \quad (23)$$



Similarly, for any  $\gamma > 0$ , from (22) we have

$$\begin{aligned} & |\hat{u}_N(\mathbf{z}) - u(\mathbf{z})| < \frac{A}{\sqrt{N}\gamma} & \forall \mathbf{z} \in \mathcal{Z} \\ \implies & \frac{|\hat{u}_N(\mathbf{z}) - u(\mathbf{z})|}{u(\mathbf{z})} < \frac{A}{\sqrt{N}\gamma \exp(\alpha\rho^*)} & \forall \mathbf{z} \in \mathcal{Z} \\ \implies & -\frac{A}{\sqrt{N}\gamma \exp(\alpha\rho^*)} < \frac{\hat{u}_N(\mathbf{z})}{u(\mathbf{z})} - 1, & \forall \mathbf{z} \in \mathcal{Z} \\ \implies & \frac{1}{\alpha} \log\left(1 - \frac{A}{\sqrt{N}\gamma \exp(\alpha\rho^*)}\right) < \frac{1}{\alpha} \log\left(\frac{\hat{u}_N(\mathbf{z})}{u(\mathbf{z})}\right) & \forall \mathbf{z} \in \mathcal{Z}, \end{aligned}$$

where the first and last implication follow as before. The last expression can be further simplified given that  $\frac{A}{\sqrt{N}\gamma \exp(\alpha\rho^*)} \leq 1 - 1/e$  for sufficiently large  $N$ , which allows us to apply the property of the logarithm function,  $\log(1 - \epsilon) \geq -\epsilon/(1 - 1/e)$  for all  $\epsilon \in [0, 1 - 1/e]$ , to obtain:

$$-\frac{A}{(1 - 1/e)\alpha\sqrt{N}\gamma \exp(\alpha\rho^*)} < \frac{1}{\alpha} \log(\hat{u}_N(\mathbf{z})) - \frac{1}{\alpha} \log(u(\mathbf{z})) \quad \forall \mathbf{z} \in \mathcal{Z}. \quad (24)$$

We now combine (23) and (24) as follows:

$$\left| \frac{1}{\alpha} \log(\hat{u}_N(\mathbf{z})) - \frac{1}{\alpha} \log(u(\mathbf{z})) \right| < \frac{A}{(1 - 1/e)\alpha\sqrt{N}\gamma \exp(\alpha\rho^*)} \quad \forall \mathbf{z} \in \mathcal{Z}. \quad (25)$$

From (22), we obtain:

$$\begin{aligned} \mathbb{P}\left(\sup_{\mathbf{z} \in \mathcal{Z}} \left| \rho_{\hat{\mathbb{P}}_N}(\mathbf{z}) - \rho_{\mathbb{P}}(\mathbf{z}) \right| < \frac{A}{(1 - 1/e)\alpha\sqrt{N}\gamma \exp(\alpha\rho^*)}\right) &\geq \mathbb{P}\left(\sup_{\mathbf{z} \in \mathcal{Z}} |\hat{u}_N(\mathbf{z}) - u(\mathbf{z})| < \frac{A}{\sqrt{N}\gamma}\right) \\ &> 1 - \gamma, \end{aligned}$$

which implies that:

$$\mathbb{P}\left(\left|\min_{\mathbf{z} \in \mathcal{Z}} \rho_{\hat{\mathbb{P}}_N}(\mathbf{z}) - \min_{\mathbf{z} \in \mathcal{Z}} \rho_{\mathbb{P}}(\mathbf{z})\right| < \frac{A}{(1 - 1/e)\alpha\sqrt{N}\gamma \exp(\alpha\rho^*)}\right) > 1 - \gamma.$$

The latter can be rewritten using our notation as:

$$\mathbb{P}\left(|\rho_{\text{SAA}} - \rho^*| \geq B/(\sqrt{N}\gamma \alpha \exp(\alpha\rho^*))\right) \leq \gamma,$$

with  $B := A/(1 - 1/e) > 0$ .

To prove the convergence in probability, we simply consider any  $\gamma > 0$  and  $\Delta > 0$  and confirm that:

$$\mathbb{P}\left(|\rho_{\text{SAA}} - \rho^*| \geq \Delta\right) \leq \gamma,$$

as long as  $N$  is large enough for  $\frac{A}{\sqrt{N}\gamma \exp(\alpha\rho^*)} \leq (1 - 1/e) \min(\Delta, 1)$ . Indeed, for such  $N$ , we necessarily have that:

$$\mathbb{P}\left(|\rho_{\text{SAA}} - \rho^*| \geq \Delta\right) \leq \mathbb{P}\left(|\rho_{\text{SAA}} - \rho^*| \geq \frac{A}{(1 - 1/e)\sqrt{N}\gamma \exp(\alpha\rho^*)}\right) \leq \gamma.$$

□

## A.4 Proof of Theorem 3

**Proof.** We will show that  $\rho_{\text{DRO}}$  converges to  $\rho^*$  by examining the relationship between  $\rho_{\text{DRO}}$  and  $\rho_{\text{SAA}}$ , and  $\rho_{\text{SAA}}$  and  $\rho^*$ . The latter relationship is established in Proposition 1.

Concerning the relationship between  $\rho_{\text{SAA}}$  and  $\rho_{\text{DRO}}$  defined using the ambiguity set  $\mathcal{B}_\infty(\epsilon)$  with any  $\epsilon \geq 0$ , since  $\hat{\mathbb{P}}_N \in \mathcal{B}_\infty(\epsilon)$ , the following inequality holds:

$$\sup_{\mathbb{Q} \in \mathcal{B}_\infty(\epsilon)} \frac{1}{\alpha} \log(\mathbb{E}_{\mathbb{Q}}[\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))]) \geq \frac{1}{\alpha} \log(\mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha\ell(\mathbf{z}, \boldsymbol{\xi}))]) \quad \forall \mathbf{z} \in \mathcal{Z},$$

which implies that  $\rho_{\text{DRO}} \geq \rho_{\text{SAA}}$ . Moreover, we know from Bertsimas et al. (2023) that

$$\sup_{\mathbb{Q} \in \mathcal{B}_\infty(\epsilon)} \frac{1}{\alpha} \log(\mathbb{E}_{\mathbb{Q}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))]) = \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i \in [N]} \sup_{\boldsymbol{\xi}: \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\| \leq \epsilon} \exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi})) \right). \quad (26)$$

In addition, from the  $L$ -Lipschitz continuity of  $\ell$  in  $\boldsymbol{\xi}$  for all  $\mathbf{z} \in \mathcal{Z}$  (see Assumption 3(A.3)), we have

$$\begin{aligned} & |\ell(\mathbf{z}, \boldsymbol{\xi}) - \ell(\mathbf{z}, \hat{\boldsymbol{\xi}}_i)| \leq L \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\| && \forall \boldsymbol{\xi}, \forall i \\ \implies & |\ell(\mathbf{z}, \boldsymbol{\xi}) - \ell(\mathbf{z}, \hat{\boldsymbol{\xi}}_i)| \leq L\epsilon && \forall \boldsymbol{\xi} \in \{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\| \leq \epsilon\}, \forall i \\ \implies & \sup_{\boldsymbol{\xi}: \|\boldsymbol{\xi}_i - \hat{\boldsymbol{\xi}}_i\| \leq \epsilon} \exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi})) \leq \exp(\alpha(\ell(\mathbf{z}, \hat{\boldsymbol{\xi}}_i) + L\epsilon)) && \forall i, \end{aligned}$$

where the first implication follows from the definition of the ambiguity set. Substituting the resulting inequality in (26) results in

$$\begin{aligned} \sup_{\mathbb{Q} \in \mathcal{B}_\infty(\epsilon)} \frac{1}{\alpha} \log(\mathbb{E}_{\mathbb{Q}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))]) &\leq \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i \in [N]} \left( \exp(\alpha \ell(\mathbf{z}, \hat{\boldsymbol{\xi}}_i) + \alpha L\epsilon) \right) \right) \\ &= \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i \in [N]} \exp(\alpha \ell(\mathbf{z}, \hat{\boldsymbol{\xi}}_i)) \right) + L\epsilon. \end{aligned}$$

From the above inequality, it follows that  $\rho_{\text{DRO}} \leq \rho_{\text{SAA}} + L\epsilon$ . Combining with  $\rho_{\text{DRO}} \geq \rho_{\text{SAA}}$ , we conclude that  $\rho_{\text{SAA}} \leq \rho_{\text{DRO}} \leq \rho_{\text{SAA}} + L\epsilon$ .

We can now establish a high confidence bound on  $|\rho_{\text{DRO}} - \rho^*|$  that converges to zero at the rate of  $\mathcal{O}(1/\sqrt{N})$  when using  $\mathcal{B}_\infty(c/\sqrt{N})$  for some  $c > 0$ . Namely, given some  $\gamma > 0$ , we let  $\phi(N, \gamma) := \frac{A}{\sqrt{N}\gamma\alpha \exp(\alpha\rho^*)} = \mathcal{O}(1/\sqrt{N})$  as defined in Proposition 1. We can then show:

$$\begin{aligned} \mathbb{P} \left( |\rho_{\text{DRO}} - \rho^*| \geq \phi(N, \gamma) + Lc/\sqrt{N} \right) &\leq \mathbb{P} \left( |\rho_{\text{DRO}} - \rho_{\text{SAA}}| + |\rho_{\text{SAA}} - \rho^*| \geq \phi(N, \gamma) + Lc/\sqrt{N} \right) \\ &\leq \mathbb{P}(|\rho_{\text{SAA}} - \rho^*| \geq \phi(N, \gamma)) \\ &\leq \gamma, \end{aligned}$$

where the first inequality follows from the triangle inequality, the second from  $|\rho_{\text{DRO}} - \rho_{\text{SAA}}| \leq Lc/\sqrt{N}$ , and the third from Proposition 1 as long as  $N$  is large enough.

To prove the convergence in probability, we simply consider any  $\gamma > 0$  and  $\Delta > 0$  and confirm that:

$$\mathbb{P}(|\rho_{\text{DRO}} - \rho^*| \geq \Delta) \leq \gamma,$$

as long as  $N$  is large enough for Proposition 1 to apply and  $\phi(N, \gamma) + Lc/\sqrt{N} \leq \Delta$ . Indeed, for such  $N$ , we necessarily have that:

$$\mathbb{P}(|\rho_{\text{DRO}} - \rho^*| \geq \Delta) \leq \mathbb{P} \left( |\rho_{\text{DRO}} - \rho^*| \geq \phi(N, \gamma) + Lc/\sqrt{N} \right) \geq \gamma. \quad \square$$

## A.5 Proof of Proposition 2

**Proof.** Let  $u(\mathbf{z})$  denote the worst-case expected utility for a given decision  $\mathbf{z}$ , that is,

$$u(\mathbf{z}) := \sup_{\mathbb{Q} \in \mathcal{B}(\epsilon)} u_{\mathbb{Q}}(\ell(\mathbf{z}, \boldsymbol{\xi})) = \mathbb{E}_{\mathbb{Q}}[\exp(\alpha \ell(\mathbf{z}, \boldsymbol{\xi}))], \quad (27)$$

where the ambiguity set  $\mathcal{B}(\epsilon)$  of radius  $\epsilon \geq 0$  is defined as:

$$\mathcal{B}(\epsilon) := \left\{ \mathbb{Q} \in \mathcal{M}(\Xi) \mid \mathbb{Q} \{ \boldsymbol{\xi} \in \Xi \} = 1, \mathcal{W}^p(\mathbb{Q}, \hat{\mathbb{P}}_N) \leq \epsilon \right\}.$$

The type- $p$  Wasserstein distance,  $\mathcal{W}^p(\mathbb{P}_1, \mathbb{P}_2)$ , is defined as:

$$\mathcal{W}^p(\mathbb{P}_1, \mathbb{P}_2) = \inf_{\pi \in \mathcal{M}(\Xi \times \Xi)} \left\{ \left( \int_{\Xi} \int_{\Xi} \|\xi_1 - \xi_2\|^p \pi(d\xi_1, d\xi_2) \right)^{1/p} \right\},$$

where  $\pi$  is a joint distribution of  $\zeta_1$  and  $\zeta_2$  with marginals  $\mathbb{P}_1$  and  $\mathbb{P}_2$ , respectively. From Gao and Kleywegt (2023, Lemma 2, Proposition 2), the worst-case utility  $u(\mathbf{z})$  for any  $\mathbf{z} \in \mathcal{Z}$  is infinite with  $p$ -Wasserstein ( $p < \infty$ ) ambiguity set since the loss function  $\exp(\alpha \ell(\mathbf{z}, \xi))$  does not satisfy the growth condition, that is, there does not exist any  $\xi_0 \in \Xi$  and constants  $L > 0, M > 0$  such that  $\exp(\alpha \ell(\mathbf{z}, \xi)) \leq L \|\xi - \xi_0\|^p + M$  holds for all  $\xi \in \Xi$ .

Next, we prove by contradiction that the worst-case entropic risk  $\rho(\mathbf{z}) := \sup_{\mathbb{Q} \in \mathcal{B}(\epsilon)} \rho_{\mathbb{Q}}(\ell(\mathbf{z}, \xi))$  is unbounded for all  $\mathbf{z} \in \mathcal{Z}$ . Suppose that  $\rho(\mathbf{z})$  is bounded, implying that there exists an  $M < \infty$  such that:

$$\rho(\mathbf{z}) = \sup_{\mathbb{Q} \in \mathcal{B}(\epsilon)} \frac{1}{\alpha} \log(u_{\mathbb{Q}}(\mathbf{z})) \leq M, \quad (28)$$

which can be equivalently written as:

$$u_{\mathbb{Q}}(\ell(\mathbf{z}, \xi)) \leq \exp(\alpha M) \text{ for all } \mathbb{Q} \in \mathcal{B}(\epsilon), \quad (29)$$

since the exponential function is a monotonically increasing. This implies that  $u(\mathbf{z})$  is bounded, which contradicts the result in Gao and Kleywegt (2023). Hence, the worst-case entropic risk  $\rho(\mathbf{z})$  is infinite for all  $\mathbf{z} \in \mathcal{Z}$ .  $\square$

## A.6 Proof of Theorem 2

**Proof.** Problem (12) can be equivalently written as:

$$\begin{aligned} \min \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ \text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathcal{Z} \\ & t_i \geq \sup_{\xi \in \Xi_i^i} \ell(\mathbf{z}, \xi) \quad \forall i \in [N], \end{aligned}$$

where  $\Xi_\epsilon^i = \{\xi : \|\xi - \hat{\xi}_i\| \leq \epsilon\}$ . Since  $\ell(\mathbf{z}, \xi) = \max_{j \in [m]} \ell_j(\mathbf{z}, \xi)$ , we obtain:

$$\begin{aligned} \min_{\mathbf{t}} \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ \text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathcal{Z}, \\ & t_i \geq \sup_{\xi \in \Xi_\epsilon^i} \ell_j(\mathbf{z}, \xi) \quad \forall i \in [N], j \in [m]. \end{aligned} \quad (30)$$

The relative interior of the intersection of set  $\Xi_\epsilon^i$  and domain of  $\ell_j(\mathbf{z}, \xi)$  is non-empty for all  $\epsilon \geq 0$ . So, we can use Fenchel duality theorem (Ben-Tal et al. 2015) to obtain:

$$\sup_{\xi \in \Xi_\epsilon^i} \ell_j(\mathbf{z}, \xi) = \inf_{\varphi_{ij}} \delta^*(\varphi_{ij} | \Xi_\epsilon^i) - \ell_{j*}(\mathbf{z}, \varphi_{ij}), \quad (31a)$$

where  $\varphi_{ij} \in \mathbb{R}^d$ ,  $\ell_{j*}(\mathbf{z}, \varphi_{ij}) := \inf_{\xi} \{\varphi_{ij}^\top \xi - \ell_j(\mathbf{z}, \xi)\}$  is the partial concave conjugate of  $\ell_j(\mathbf{z}, \xi)$  and  $\delta^*(\varphi_{ij} | \Xi_\epsilon^i)$  is the support function of  $\Xi_\epsilon^i$ , i.e.,

$$\delta^*(\varphi_{ij} | \Xi_\epsilon^i) = \sup_{\xi \in \Xi_\epsilon^i} \varphi_{ij}^\top \xi = \varphi_{ij}^\top \hat{\xi}_i + \sup_{\zeta : \|\zeta\| \leq \epsilon} \varphi_{ij}^\top \zeta = \varphi_{ij}^\top \hat{\xi}_i + \epsilon \|\varphi_{ij}\|_*, \quad (31b)$$

where the last equality follows by the definition of the dual norm. Substituting (31) in (30) results in the following finite dimensional conic program:

$$\begin{aligned}
\min \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\
\text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathcal{Z}, \boldsymbol{\varphi}_{ij} \in \mathbb{R}^d \quad \forall i \in [N], j \in [m] \\
& \boldsymbol{\varphi}_{ij}^\top \hat{\boldsymbol{\xi}}_i + \epsilon \|\boldsymbol{\varphi}_{ij}\|_* - \ell_{j^*}(\mathbf{z}, \boldsymbol{\varphi}_{ij}) \leq t_i \quad \forall i \in [N], j \in [m].
\end{aligned}$$

□

## A.7 Proof of Proposition 3

**Proof.**

$$\begin{aligned}
\mathbb{E}[\rho_{k \sim U(K)}(\rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})))]) &= \mathbb{E} \left[ \frac{1}{\alpha} \log \left( \frac{1}{K} \sum_{k=1}^K \exp \left( \alpha \rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})) \right) \right) \right] \\
&\leq \frac{1}{\alpha} \log \left( \mathbb{E} \left[ \frac{1}{K} \sum_{k=1}^K \exp \left( \alpha \rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})) \right) \right] \right) \\
&= \frac{1}{\alpha} \log \left( \frac{1}{K} \sum_{k=1}^K \mathbb{E} \left[ \exp \left( \alpha \rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_k^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-k}^K, \epsilon), \boldsymbol{\xi})) \right) \right] \right) \\
&= \frac{1}{\alpha} \log \left( \mathbb{E} \left[ \exp \left( \alpha \rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_1^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-1}^K, \epsilon), \boldsymbol{\xi})) \right) \right] \right) \\
&= \rho \left( \rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_1^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-1}^K, \epsilon), \boldsymbol{\xi})) \right) \\
&= \rho \left( \rho \left( \rho_{\boldsymbol{\xi} \sim \hat{\mathbb{P}}_1^K}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-1}^K, \epsilon), \boldsymbol{\xi})) \middle| \hat{\mathbb{P}}_{-1}^K \right) \right) \\
&= \rho \left( \rho_{\boldsymbol{\xi} \sim \mathbb{P}}(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{-1}^K, \epsilon), \boldsymbol{\xi})) \right) \\
&= \rho(\ell(\mathbf{z}^*(\hat{\mathbb{P}}_{N-N/K}, \epsilon), \boldsymbol{\xi})),
\end{aligned}$$

where expectations and  $\rho$ 's are with respect to randomness in the data  $\mathcal{D}_N$ , except for the last equation where the randomness is in both the data and a new sample  $\boldsymbol{\xi} \sim \mathbb{P}$ . The first inequality follows from concavity of log function and Jensen's inequality, then we exploit the fact that each  $(\hat{\mathbb{P}}_k^K, \hat{\mathbb{P}}_{-k}^K)$  pair is identically distributed to  $(\hat{\mathbb{P}}_1^K, \hat{\mathbb{P}}_{-1}^K)$ , and finally, we use the tower property of the entropic risk measure. □

## A.8 Additional results

**Lemma 2.** *The following conditions are equivalent:*

1. *There exist some constants  $G > 0$  and  $C > 2$  such that  $\mathbb{P}(|\ell(\boldsymbol{\xi})| > a) \leq G \exp(-a\alpha C)$ ,  $\forall a \geq 0$ ,*
2. *The moment-generating function of  $\ell(\boldsymbol{\xi})$  satisfies  $\mathbb{E}[\exp(t\ell(\boldsymbol{\xi}))] \in \mathbb{R}$  for all  $t \in [-\alpha C, \alpha C]$ , for some  $C > 2$ .*

**Proof.** The property  $\mathbb{P}(|\ell(\boldsymbol{\xi})| > a) \leq G \exp(-a\alpha C)$ ,  $\forall a \geq 0$  for some  $G > 0$  and  $C > 2$ , implies that when  $t \in [-\alpha C, \alpha C]$ ,  $C > 2$ , we have that:

$$\begin{aligned}
0 &\leq \mathbb{E}[\exp(t\ell(\boldsymbol{\xi}))] \leq \mathbb{E}[\exp(|t||\ell(\boldsymbol{\xi})|)] \\
&\leq \int_0^\infty \mathbb{P}(\exp(|t||\ell(\boldsymbol{\xi})|) > x) dx \\
&= \int_0^\infty \mathbb{P}(|\ell(\boldsymbol{\xi})| > y) |t| \exp(|t|y) dy
\end{aligned}$$

$$\begin{aligned} &\leq \int_0^\infty G|t| \exp(-(\alpha C - |t|)y) dy \\ &= \frac{G|t|}{\alpha C - |t|} \leq \frac{G\alpha C}{\alpha C - |t|}, \end{aligned}$$

where we use similar arguments to the ones used to bound  $\mathbb{E}[\exp(\alpha\ell(\boldsymbol{\xi}))]$  above in the proof of Lemma 1 and exploit the fact that  $|t| \leq \alpha C$ .

Alternatively,  $\mathbb{E}[\exp(t\ell(\boldsymbol{\xi}))] \in \mathbb{R}$  for all  $t \in [-\alpha C, \alpha C]$  for some  $C > 2$  implies that for any  $a \geq 0$

$$\mathbb{P}(\ell(\boldsymbol{\xi}) > a) = \mathbb{P}(\exp(\alpha C \ell(\boldsymbol{\xi})) > \exp(\alpha a C)) \leq \frac{\mathbb{E}[\exp(\alpha C \ell(\boldsymbol{\xi}))]}{\exp(\alpha a C)} = G^+ \exp(-\alpha a C),$$

where we have used Markov's inequality to obtain the upper bound and  $G^+ := \mathbb{E}[\exp(\alpha C \ell(\boldsymbol{\xi}))] \in \mathbb{R}$ . A similar argument holds for  $\mathbb{P}(-\ell(\boldsymbol{\xi}) > a) \leq G^- \exp(-\alpha a C)$  with  $G^- := \mathbb{E}[\exp(-\alpha C \ell(\boldsymbol{\xi}))] \in \mathbb{R}$ . Thus, by the union bound, we have:

$$\mathbb{P}(|\ell(\boldsymbol{\xi})| > a) \leq (G^+ + G^-) \exp(-\alpha a C), \quad \forall a \geq 0.$$

□

**Corollary 2.** *The DRO newsvendor problem with the cost function given by  $\ell(\xi, z) = wz + b(\xi - z)_+ + h(z - \xi)_+$  and type- $\infty$  Wasserstein ambiguity set is equivalent to:*

$$\begin{aligned} &\text{minimize} && \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ &\text{subject to} && (\hat{\xi}_i + \epsilon)b - z(w - b) \leq t_i \quad \forall i \in [N] \\ &&& (\epsilon - \hat{\xi}_i)h + z(w + h) \leq t_i \quad \forall i \in [N]. \end{aligned} \tag{32}$$

**Proof.** For the newsvendor problem, the cost function is given by  $\ell(z, \xi) = wz + b(\xi - z)_+ + h(z - \xi)_+$ . Problem (30) is equivalent to:

$$\begin{aligned} &\min && \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ &\text{s.t.} && \mathbf{t} \in \mathbb{R}^N, z \in \mathbb{R} \\ &&& t_i \geq wz + b\xi - bz \quad \forall \xi \in \Xi_\epsilon^i, \forall i \in [N] \\ &&& t_i \geq wz + hz - h\xi \quad \forall \xi \in \Xi_\epsilon^i, \forall i \in [N], \end{aligned} \tag{33}$$

where  $\Xi_\epsilon^i = \{\xi : \|\xi - \hat{\xi}_i\| \leq \epsilon\}$ . Define  $\ell^+(z, \xi) := wz + b\xi - bz$  and  $\ell^-(z, \xi) = wz + hz - h\xi$  for which the partial concave conjugates are given by:

$$\begin{aligned} \ell_*^+(z, \varphi^+) &= \begin{cases} z(b - w) & \text{if } \varphi^+ = b \\ -\infty & \text{otherwise,} \end{cases} \\ \ell_*^-(z, \varphi^-) &= \begin{cases} -z(w + h) & \text{if } \varphi^- = -h \\ -\infty & \text{otherwise.} \end{cases} \end{aligned}$$

Thus, from Theorem 2, problem (30) reduces to (32). □

**Corollary 3.** *The DRO regression problem with loss function  $\ell(\mathbf{z}, \boldsymbol{\xi}) = |\boldsymbol{\xi}_{d+1} - \boldsymbol{\xi}_{1:d}^\top \mathbf{z}|$ , and type- $\infty$  Wasserstein ambiguity set is equivalent to:*

$$\min_{\mathbf{z} \in \mathbb{R}^d} \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha |\hat{\boldsymbol{\xi}}_{i,d+1} - \hat{\boldsymbol{\xi}}_{i,1:d}^\top \mathbf{z}|) \right) + \epsilon \|[-1 \ \mathbf{z}^\top]^\top\|_*. \tag{34}$$

**Proof.** For the regression problem, the cost function is given by  $\ell(\mathbf{z}, \boldsymbol{\xi}) = |\boldsymbol{\xi}_{d+1} - \boldsymbol{\xi}_{1:d}^\top \mathbf{z}|$ . Problem (30) is equivalent to:

$$\begin{aligned} \min \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ \text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathbb{R}^d \\ & t_i \geq \boldsymbol{\xi}_{d+1} - \boldsymbol{\xi}_{1:d}^\top \mathbf{z} \quad \forall \boldsymbol{\xi} \in \Xi_\epsilon^i, \forall i \in [N] \\ & t_i \geq \boldsymbol{\xi}_{1:d}^\top \mathbf{z} - \boldsymbol{\xi}_{d+1} \quad \forall \boldsymbol{\xi} \in \Xi_\epsilon^i, \forall i \in [N], \end{aligned} \quad (35)$$

where  $\Xi_\epsilon^i = \{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\| \leq \epsilon\}$ . Define  $\ell^+(\mathbf{z}, \boldsymbol{\xi}) := \boldsymbol{\xi}_{d+1} - \boldsymbol{\xi}_{1:d}^\top \mathbf{z}$  and  $\ell^-(\mathbf{z}, \boldsymbol{\xi}) := \boldsymbol{\xi}_{1:d}^\top \mathbf{z} - \boldsymbol{\xi}_{d+1}$ . Then, the partial concave conjugates of  $\ell^+(\mathbf{z}, \boldsymbol{\xi})$  and  $\ell^-(\mathbf{z}, \boldsymbol{\xi})$  are given by:

$$\ell_*^+(\mathbf{z}, \boldsymbol{\varphi}^+) = \begin{cases} 0 & \text{if } \varphi_{d+1}^+ = 1 \text{ and } \varphi_{1:d}^+ = -\mathbf{z} \\ -\infty & \text{otherwise,} \end{cases} \quad (36a)$$

$$\ell_*^-(\mathbf{z}, \boldsymbol{\varphi}^-) = \begin{cases} 0 & \text{if } \varphi_{d+1}^- = -1 \text{ and } \varphi_{1:d}^- = \mathbf{z} \\ -\infty & \text{otherwise.} \end{cases} \quad (36b)$$

Substituting (36) in problem (30) reduces to:

$$\begin{aligned} \min \quad & \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha t_i) \right) \\ \text{s.t.} \quad & \mathbf{t} \in \mathbb{R}^N, \mathbf{z} \in \mathbb{R}^d \\ & t_i \geq \hat{\boldsymbol{\xi}}_{i,d+1} - \hat{\boldsymbol{\xi}}_{i,1:d}^\top \mathbf{z} + \epsilon \|[1 - \mathbf{z}^\top]^\top\|_* \quad \forall i \in [N] \\ & t_i \geq \hat{\boldsymbol{\xi}}_{i,1:d}^\top \mathbf{z} - \hat{\boldsymbol{\xi}}_{i,d+1} + \epsilon \|[ -1 \ \mathbf{z}^\top ]^\top\|_* \quad \forall i \in [N]. \end{aligned}$$

Finally:

$$\min_{\mathbf{z} \in \mathbb{R}^d} \frac{1}{\alpha} \log \left( \frac{1}{N} \sum_{i=1}^N \exp(\alpha |\hat{\boldsymbol{\xi}}_{i,d+1} - \hat{\boldsymbol{\xi}}_{i,1:d}^\top \mathbf{z}|) \right) + \epsilon \|[ -1 \ \mathbf{z}^\top ]^\top\|_*. \quad \square$$

**Lemma 3.** *The following inequalities follow from the properties of the logarithm function:*

$$\begin{aligned} \log(1 + \epsilon) &\leq \epsilon && \text{if } \epsilon \geq 0, \\ \log(1 - \epsilon) &\geq -\epsilon/(1 - 1/e) && \text{if } \epsilon \in [0, 1 - 1/e]. \end{aligned}$$

**Proof.** The logarithm function is concave, thus it follows that

$$\log(1 + \epsilon) \leq \log(1) + \epsilon \left. \frac{d \log(x)}{dx} \right|_{x=1} = \epsilon.$$

Moreover, by concavity of  $\log(1 - \epsilon)$  for  $\epsilon \in [0, 1 - 1/e]$ , we have:

$$\log(1 - \epsilon) \geq \log(1) + \epsilon \frac{\log(1/e) - \log(1)}{1 - 1/e} = -\frac{\epsilon}{1 - 1/e}.$$

Figure 13 gives a pictorial representation of the result. □

## B Additional details

### B.1 Derivation of influence function

The influence function measures the sensitivity of a statistic to small changes in the data. The influence function of a statistic  $T$  at a point  $\zeta$  for a distribution  $\mathbb{P}$  is given by:

$$\text{IF}(\zeta) = \lim_{\epsilon \rightarrow 0^+} \frac{T((1 - \epsilon)\mathbb{P} + \epsilon\delta_\zeta) - T(\mathbb{P})}{\epsilon},$$

where  $\delta_\zeta$  is the Dirac distribution at the point  $\zeta$ . The influence function of entropic risk measure is given by:

$$\begin{aligned} \text{IF}(\boldsymbol{\xi}) &= \lim_{\varepsilon \rightarrow 0^+} \frac{\log((1-\varepsilon)\mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\boldsymbol{\xi}))] + \varepsilon \exp(\alpha\ell(\boldsymbol{\xi}))) - \log(\mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\boldsymbol{\xi}))])}{\alpha\varepsilon} \\ &= \lim_{\varepsilon \rightarrow 0^+} \frac{\log\left(\frac{(1-\varepsilon) + \varepsilon \frac{\exp(\alpha\ell(\boldsymbol{\xi}))}{\mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\boldsymbol{\xi}))]}}{\alpha\varepsilon}\right)}{\alpha\varepsilon} = -\frac{1}{\alpha} \left(1 - \frac{\exp(\alpha\ell(\boldsymbol{\xi}))}{\mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\boldsymbol{\xi}))]}\right). \end{aligned}$$

When loss function  $\ell(\xi) := \xi$  with  $\xi \sim \Gamma(\kappa, \lambda)$ , the influence function can be obtained in closed form as follows:

$$\text{IF}(\hat{\xi}) = -\frac{1}{\alpha} + \frac{\exp(\alpha\hat{\xi})}{\alpha(1-\lambda\alpha)^{-\kappa}}.$$

Figure 14 presents the histogram of 500 loss values generated from the  $\Gamma(10, 0.24)$  distribution. For each bin, we calculate the average influence function over samples and normalize these values to lie in the interval  $[0, 1]$ , where darker shades represent higher average influence function. It can be seen that tail events have highest impact on the entropic risk but occur with very low probability. Therefore, these high-impact scenarios are likely not to be included in a finite sample, resulting in the underestimation of the entropic risk.




Figure 14: Histogram of the 500 loss samples,  $\xi$ , from a  $\Gamma(10, 0.24)$  distribution and the average IF of the entropic risk is computed by taking the mean of IF over scenarios in each bin with the darker blue shade representing higher (normalized) average IF.

## B.2 Bias mitigation using OIC

From Theorem 1 in Iyengar et al. (2023), it follows that for a loss function  $h(t, \ell(\boldsymbol{\xi}))$  with decision  $t$ :

$$\mathbb{E}[\mathbb{E}_{\mathbb{P}}(h(\hat{t}, \ell(\boldsymbol{\xi})))] = \mathbb{E}[h(\hat{t}, \ell(\hat{\boldsymbol{\xi}}_i))] - \underbrace{\frac{1}{N} \mathbb{E}_{\mathbb{P}}[\nabla_t h(t^*, \ell(\boldsymbol{\xi})) \text{IF}(t^*)]}_{\delta_{\text{OIC}}} + o\left(\frac{1}{N}\right),$$

where expectation is with respect to the randomness of  $\mathcal{D}_N$ . For  $h(t, \ell(\boldsymbol{\xi})) = t + \frac{1}{\alpha} \exp(\alpha(\ell(\boldsymbol{\xi}) - t)) - \frac{1}{\alpha}$ , we know that  $t^* = \rho_{\mathbb{P}}(\ell(\boldsymbol{\xi})) = \mathbb{E}_{\mathbb{P}}[\exp(\alpha\ell(\boldsymbol{\xi}))]$ ,  $\nabla_t h(t^*, \ell(\boldsymbol{\xi})) = 1 - \exp(\alpha(\ell(\boldsymbol{\xi}) - t^*))$  and  $\nabla_{t,t} h(t^*, \ell(\boldsymbol{\xi})) = \alpha \exp(\alpha(\ell(\boldsymbol{\xi}) - t^*))$ . The influence function in the expression of  $\delta_{\text{OIC}}$  is obtained as follows:

$$\text{IF}(t^*) = -(\mathbb{E}_{\mathbb{P}}[\nabla_{t,t}^2 h(t^*, \ell(\boldsymbol{\xi}))])^{-1} \nabla_t h(t^*, \ell(\boldsymbol{\xi})) = -\frac{1 - \exp(\alpha(\ell(\boldsymbol{\xi}) - t^*))}{\mathbb{E}_{\mathbb{P}}[\alpha \exp(\alpha(\ell(\boldsymbol{\xi}) - t^*))]}$$

Next, we substitute the value of  $\text{IF}(t^*)$  and  $\nabla_t h(t^*, \ell(\boldsymbol{\xi}))$  to obtain the bias of a decision  $t^*$ :

$$\begin{aligned} \delta_{\text{OIC}} &= -\frac{1}{N} \mathbb{E}_{\mathbb{P}}[\nabla_t h(t^*, \ell(\boldsymbol{\xi})) \text{IF}(t^*)] \\ &= \frac{1}{N} \left( \frac{\mathbb{E}_{\mathbb{P}}[1 - \exp(\alpha(\ell(\boldsymbol{\xi}) - t^*))]^2}{\alpha \mathbb{E}_{\mathbb{P}}[\exp(\alpha(\ell(\boldsymbol{\xi}) - t^*))]} \right) \\ &= \frac{1}{N} \left( \frac{\mathbb{E}_{\mathbb{P}}[\mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\boldsymbol{\xi}))] - \exp(\alpha \ell(\boldsymbol{\xi}))]^2}{\alpha (\mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\boldsymbol{\xi}))])^2} \right) \\ &= \frac{\text{Var}_{\mathbb{P}}(\exp(\alpha \ell(\boldsymbol{\xi})))}{N \alpha (\mathbb{E}_{\mathbb{P}}[\exp(\alpha \ell(\boldsymbol{\xi}))])^2}. \end{aligned}$$

Since  $\mathbb{P}$  is not known, Iyengar et al. (2023) replace  $\mathbb{P}$  with  $\hat{\mathbb{P}}$  to obtain their estimator,  $\rho_{\text{OIC}} := \hat{t} + \text{Var}_{\hat{\mathbb{P}}_N}(\exp(\alpha \ell(\boldsymbol{\xi}))) / (N \alpha (\mathbb{E}_{\hat{\mathbb{P}}_N}[\exp(\alpha \ell(\boldsymbol{\xi}))])^2)$ .

### B.3 Fitting a GMM

**Entropic risk matching.** After each update of the parameters of a GMM using the gradient descent procedure described in Algorithm 2, the parameters are projected back into the feasible region for a valid GMM. This ensures that the mixing weights  $\hat{\boldsymbol{\pi}}_{t+1}$  remain valid probabilities (which is achieved using a softmax function), and that the standard deviations  $\hat{\boldsymbol{\sigma}}_{t+1}$  are positive (enforced by taking the maximum of  $\exp(-5)$  and  $\hat{\boldsymbol{\sigma}}_{t+1}$ ). The number of components  $Y$  in the GMM is selected by CV based on the Wasserstein distance between the distribution of the entropic risk of samples drawn from fitted GMM and the distribution of entropic risk constructed from the scenarios in the validation set. In all numerical experiments, we set the maximum iterations  $T = 30000$  and tolerance  $\epsilon = \exp(-9)$ .

**Matching the extremes.** The distribution of the maxima of  $n$  i.i.d samples from a normal distribution  $\mathcal{N}(\mu, \sigma)$  is given by  $(\Phi_{\mu, \sigma})^n$  where  $\Phi_{\mu, \sigma}$  is the cdf of a normally distributed random variable with mean  $\mu$  and  $\sigma$ . We find the parameters of the normal distribution by matching the 50th and 90th quantiles of  $F_{\mathcal{M}}(\cdot)$  to the corresponding quantiles of  $(\Phi_{\mu, \sigma})^n$ :

$$\begin{aligned} \mu + \sigma \Phi_{0,1}^{-1}(0.5^{1/n}) &= F_{\mathcal{M}}(0.5) \\ \mu + \sigma \Phi_{0,1}^{-1}(0.9^{1/n}) &= F_{\mathcal{M}}(0.9), \end{aligned}$$

where the  $p$ -th quantile for  $Y \sim \mathcal{N}(\mu, \sigma)$  is given by  $\mu + \sigma \Phi_{0,1}^{-1}(p)$ . Solving the above two equations in  $(\mu, \sigma)$  gives an approximate distribution for the tails of the underlying distribution, which depends on the true distribution, the total number of samples  $N$ , and the number of bins  $B$ .

To balance the trade-off between the number of bins and the sample size in each bin, we set  $B = \sqrt{N}$ , which is a reasonable compromise. A large number of bins provides more independent realizations, reducing estimation bias, while a sufficiently large sample size in each bin ensures that the maxima accurately represent the extremes of the distribution.

### B.4 Differential sampling from GMM

In Algorithm 2, differentiable samples are generated using Algorithm 6. This algorithm leverages the reparameterization trick (Kingma et al. 2015) for continuous distributions and the Gumbel-Softmax trick (Jang et al. 2017, Maddison et al. 2017) for discrete distributions. The Gumbel-Softmax trick allows approximate, differentiable sampling of mixture components, while Gaussian samples are obtained by combining deterministic transformations of the parameters with random noise. As a result, gradients can flow through both the discrete and continuous sampling steps. This enables the minimization of the Wasserstein distance between the empirical and model-based entropic risk distributions.



**Algorithm 6 Differentiable Sampling from GMM.**


---

```

1: function SAMPLEGMM( $n, \theta, \tau$ )
2:   Initialize an empty set of samples  $\mathcal{S}$ 
3:   Extract mixture weights  $\boldsymbol{\pi}$ , means  $\boldsymbol{\mu}$ , and standard deviations  $\boldsymbol{\sigma}$  from  $\theta$ 
4:   Let  $Y$  be the number of mixture components (length of  $\boldsymbol{\pi}$ )
5:   for  $i = 1$  to  $n$  do
6:     Generate Gumbel noise  $g \in \mathbb{R}^Y$ 
7:     Calculate logits:  $\text{logits} = \log(\boldsymbol{\pi}) + g$ 
8:     Compute softmax weights:  $\boldsymbol{w} = \text{softmax}(\text{logits}/\tau)$ 
9:     Generate standard normal noises  $\epsilon$  with  $\epsilon_k \sim \mathcal{N}(0, 1)$ 
10:    Compute component samples:  $z_k = \mu_k + \sigma_k \cdot \epsilon_k$  for  $k = 1$  to  $Y$ 
11:    Compute final sample:  $\text{sample} = \sum_{k=1}^Y w_k \cdot z_k$ 
12:    Append sample to  $\mathcal{S}$ 
13:  return  $\mathcal{S}$ 

```

---

**B.5 Parameters in Example 3**

The parameters of the GMM in Example 3 are given by:

$$\boldsymbol{\pi} = \begin{bmatrix} 0.16 \\ 0.28 \\ 0.23 \\ 0.20 \\ 0.13 \end{bmatrix}, \quad \boldsymbol{\mu} = \begin{bmatrix} -19.5 \\ -19.0 \\ -18.5 \\ -18.0 \\ -17.5 \end{bmatrix}, \quad \boldsymbol{\sigma} = \begin{bmatrix} 4/25 \\ 1/4 \\ 4/9 \\ 1 \\ 4 \end{bmatrix}.$$

The expected value of  $\xi$  is  $-18.57$  and standard deviation is  $1.65$ .

**B.6 Households have different marginal distribution**

In this experiment, we examined the scenario where each household has a distinct Gamma-distributed loss function. The scale,  $\kappa$  and location parameters,  $\lambda$ , of the  $\Gamma$ -distribution of the loss of the five households are given by  $(8, 0.41)$ ,  $(8.5, 0.42)$ ,  $(9, 0.43)$ ,  $(9.5, 0.44)$ , and  $(10, 0.45)$ , respectively. The correlation coefficient was set to  $r = 0.5$ , indicating a moderate positive correlation among the losses. We compared the performance of the proposed methods—**BS-EVT** and **BS-Match**—with the traditional CV approach and SAA.

The findings in this case mirror those of the first experiment where households have the same marginal loss distribution: As depicted in Figure 15a, **BS-EVT** and **BS-Match** consistently outperform the traditional CV method and SAA across different sample sizes. Furthermore, increasing  $N$  leads to a reduction in the out-of-sample entropic risk for all methods, with **BS-EVT** and **BS-Match** showing the most significant improvements. This is due to the fact that **CV** and **SAA** are overly optimistic and choose smaller radius as compared to **Oracle**, whereas the proposed methods **BS-EVT** and **BS-Match** choose close to optimal radius, see Figure 15b.

**B.7 Estimate of entropic risk**

Figures 16a–16c present the statistics of the estimate of the out-of-sample risk for different  $N \in \{500, 5000, 10000\}$  and radius  $\epsilon$  of the ambiguity set in the interval  $[0, 6]$ . Similar to Figure 6, it can be seen that **CV** underestimates the optimal entropic risk for each  $\epsilon$ . However, **BS-Match** and **BS-EVT** make better estimation of the variation in the true entropic risk with  $\epsilon$ , thereby enabling a more informed choice of  $\epsilon^*$ .

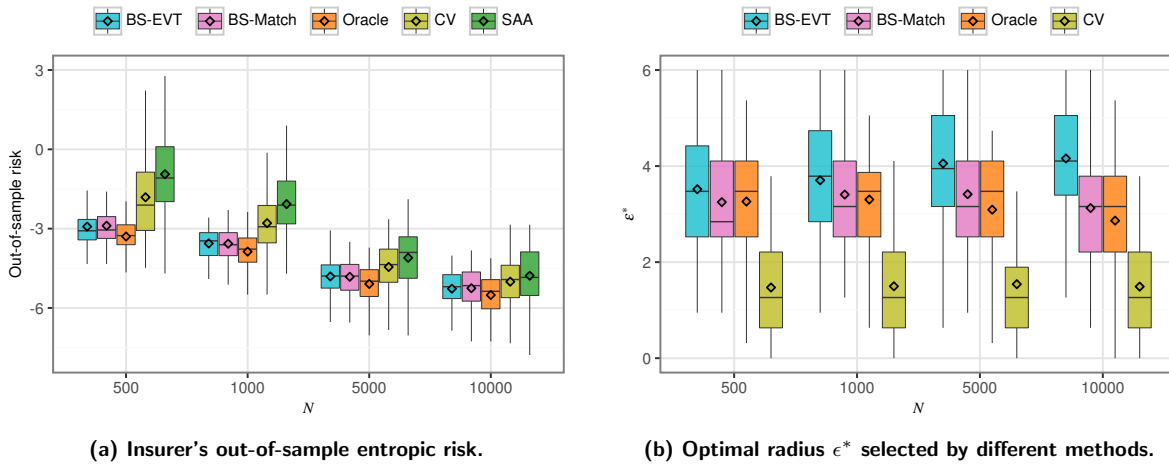


Figure 15: Comparison of the effects of training sample size  $N$  on entropic risk (left) and radius  $\epsilon^*$  (right). Each household observes samples from different  $\Gamma$  marginals and the correlation coefficient  $r = 0.5$ .

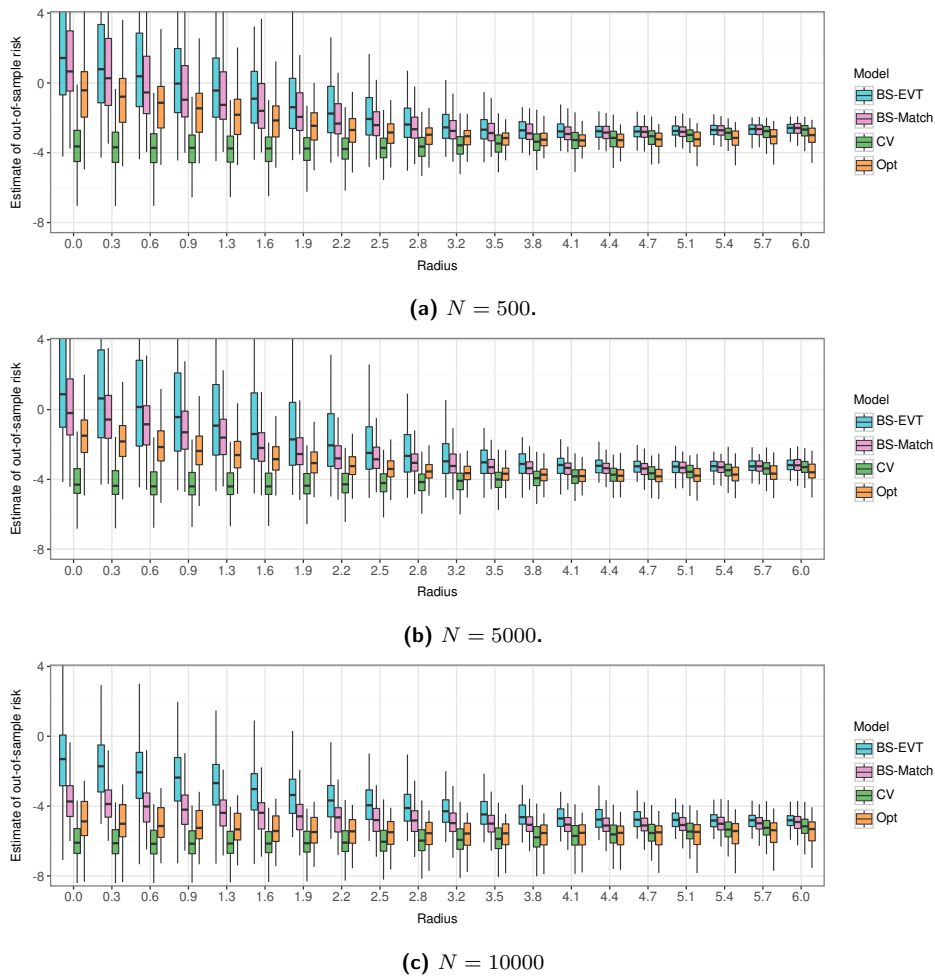


Figure 16: Estimate of entropic risk for different radius and  $N$ .

## References

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5):941–973.

- Arrow, K. J. (1971). *Essays in the Theory of Risk Bearing*. Markham Publishing Co., Chicago.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical finance*, 9(3):203–228.
- Bakkensen, L. A. and Barrage, L. (2022). Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. *The Review of Financial Studies*, 35(8):3666–3709.
- Bartl, D. and Mendelson, S. (2022). On Monte-Carlo methods in convex stochastic optimization. *The Annals of Applied Probability*, 32(4):3146–3198.
- Beirami, A., Razaviyayn, M., Shahrapour, S., and Tarokh, V. (2017). On optimal generalizability in parametric learning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30, pages 3456–3466. Curran Associates, Inc.
- Ben-Tal, A., den Hertog, D., and Vial, J.-P. (2015). Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming*, 149(1):265–299.
- Ben-Tal, A. and Teboulle, M. (1986). Expected utility, penalty functions, and duality in stochastic nonlinear programming. *Management Science*, 32(11):1445–1466.
- Bernard, C., Liu, F., and Vanduffel, S. (2020). Optimal insurance in the presence of multiple policyholders. *Journal of Economic Behavior & Organization*, 180:638–656.
- Bernard, C. and Tian, W. (2010). Insurance market effects of risk management metrics. *The Geneva Risk and Insurance Review*, 35(1):47–80.
- Bertsimas, D., Shtern, S., and Sturt, B. (2023). A data-driven approach to multistage stochastic linear optimization. *Management Science*, 69(1):51–74.
- Bousquet, O. and Elisseeff, A. (2000). Algorithmic stability and generalization performance. In Leen, T., Dietterich, T., and Tresp, V., editors, *Advances in Neural Information Processing Systems*, volume 13, pages 196–202. MIT Press, Cambridge, MA.
- Brandtner, M., Kürsten, W., and Rischau, R. (2018). Entropic risk measures and their comparative statics in portfolio selection: Coherence vs. convexity. *European Journal of Operational Research*, 264(2):707–716.
- Bäuerle, N. and Jaśkiewicz, A. (2024). Markov decision processes with risk-sensitive criteria: an overview. *Mathematical Methods of Operations Research*, 99(1):141–178.
- Cai, J., Tan, K. S., Weng, C., and Zhang, Y. (2008). Optimal reinsurance under VaR and CTE risk measures. *Insurance: Mathematics and Economics*, 43(1):185–196.
- Catoni, O. (2012). Challenging the empirical mean and empirical variance: A deviation study. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 48(4):1148–1185.
- Chan, G., Van Parys, B., and Bennouna, A. (2024). From distributional robustness to robust statistics: A confidence sets perspective. *arXiv preprint arXiv:2410.14008*. Accessed on Dec 19, 2024. <https://arxiv.org/abs/2410.14008>.
- Chen, L., He, L., and Zhou, Y. (2024a). An exponential cone programming approach for managing electric vehicle charging. *Operations Research*, 72(5):2215–2240.
- Chen, L., Ramachandra, A., Rujeerapaiboon, N., and Sim, M. (2024b). Robust Data-Driven CARA Optimization. *arXiv preprint arXiv:2107.06714v3*. Accessed on Dec 19, 2024. <https://arxiv.org/pdf/2107.06714>.
- Chen, L. and Sim, M. (2024). Robust CARA optimization. *Operations Research*. Forthcoming.
- Cheung, K. C., Sung, K. C. J., and Yam, S. C. P. (2014). Risk-minimizing reinsurance protection for multivariate risks. *The Journal of Risk and Insurance*, 81(1):219–236.
- Choi, S. and Ruszczyński, A. (2011). A multi-product risk-averse newsvendor with exponential utility function. *European Journal of Operational Research*, 214(1):78–84.
- DasGupta, A. (2008). The Bootstrap. In *Asymptotic Theory of Statistics and Probability*, pages 461–497. Springer, New York, NY.
- de Haan, L. and Ferreira, A. (2006). *Extreme Value Theory: An Introduction*. Springer, Berlin.
- Delage, E. and Ye, Y. (2010). Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems. *Operations Research*, 58(3):595–612.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- Donti, P., Amos, B., and Kolter, J. Z. (2017). Task-based end-to-end model learning in stochastic optimization. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30, pages 5484–5494. Curran Associates, Inc.
- Elmachtoub, A. N. and Grigas, P. (2021). Smart “predict, then optimize”. *Management Science*, 68(1):9–26.
- Fei, Y., Yang, Z., Chen, Y., and Wang, Z. (2021). Exponential Bellman Equation and Improved Regret Bounds for Risk-Sensitive Reinforcement Learning. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P.,

- and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems*, volume 34, pages 20436–20446.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., et al. (2021). POT: Python Optimal Transport. *Journal of Machine Learning Research*, 22(78):1–8.
- Fu, L. and Moncher, R. B. (2004). Severity Distributions for GLMs: Gamma or Lognormal? Evidence from Monte Carlo Simulations. *Casualty Actuarial Society Discussion Paper Program*, pages 149–230.
- Föllmer, H. and Schied, A. (2002). Convex measures of risk and trading constraints. *Finance and Stochastics*, 6:429–447.
- Föllmer, H. and Schied, A. (2016). *Stochastic finance: an introduction in discrete time*. Walter de Gruyter, Berlin, 4th edition.
- Gallagher, J. (2014). Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States. *American Economic Journal: Applied Economics*, 6(3):206–233.
- Gao, R. and Kleywegt, A. (2023). Distributionally Robust Stochastic Optimization with Wasserstein Distance. *Mathematics of Operations Research*, 48(2):603–655.
- Gerber, H. U. (1974). On additive premium calculation principles. *ASTIN Bulletin: The Journal of the IAA*, 7(3):215–222.
- Grigas, P., Qi, M., and Shen, M. (2023). Integrated Conditional Estimation-Optimization. arXiv preprint arXiv:2110.12351v4. Accessed on Dec 19, 2024. <https://arxiv.org/pdf/2110.12351>.
- Gupta, V., Huang, M., and Rusmevichientong, P. (2024). Debiasing in-sample policy performance for small-data, large-scale optimization. *Operations Research*, 72(2):848–870.
- Hampel, F. R. (1974). The Influence Curve and its Role in Robust Estimation. *Journal of the American Statistical Association*, 69(346):383–393.
- Hau, J. L., Petrik, M., and Ghavamzadeh, M. (2023). Entropic Risk Optimization in Discounted MDPs. In Ruiz, F., Dy, J., and van de Meentyear, J.-W., editors, *International Conference on Artificial Intelligence and Statistics*, volume 206, pages 47–76. PMLR.
- Herrnstadt, E. and Sweeney, R. L. (2024). Housing market capitalization of pipeline risk: Evidence from a shock to salience and awareness. *Land Economics*. Forthcoming.
- Hino, M. and Burke, M. (2021). The effect of information about climate risk on property values. *Proceedings of the National Academy of Sciences*, 118(17):1–9.
- Howard, R. A. and Matheson, J. E. (1972). Risk-sensitive Markov decision processes. *Management science*, 18(7):356–369.
- Hu, Z. and Hong, L. J. (2012). Kullback-Leibler Divergence Constrained Distributionally Robust Optimization. *Optimization Online*. Accessed on Dec 19, 2024. <https://optimization-online.org/2012/11/3677/>.
- Ito, S., Yabe, A., and Fujimaki, R. (2018). Unbiased Objective Estimation in Predictive Optimization. In Dy, J. and Krause, A., editors, *International Conference on Machine Learning*, pages 2176–2185. PMLR.
- Iyengar, G., Lam, H., and Wang, T. (2023). Optimizer’s information criterion: Dissecting and correcting bias in data-driven optimization. arXiv preprint arXiv:2306.10081. Accessed on Dec 19, 2024. <https://arxiv.org/pdf/2306.10081>.
- Jang, E., Gu, S., and Poole, B. (2017). Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*, volume 3, pages 1920–1931. Curran Associates, Inc.
- Jiang, J., Chen, Z., and Yang, X. (2020). Rates of convergence of sample average approximation under heavy tailed distributions. *Optimization Online*. Accessed on Dec 19, 2024. <https://optimization-online.org/wp-content/uploads/2020/06/7849.pdf>.
- Kaluszka, M. (2004a). An extension of Arrow’s result on optimality of a stop loss contract. *Insurance: Mathematics and Economics*, 35(3):527–536.
- Kaluszka, M. (2004b). Mean-Variance Optimal Reinsurance Arrangements. *Scandinavian Actuarial Journal*, 2004(1):28–41.
- Kim, J. H. T. (2010). Bias correction for estimated distortion risk measure using the bootstrap. *Insurance: Mathematics and Economics*, 47(2):198–205.
- Kim, J. H. T. and Hardy, M. R. (2007). Quantifying and Correcting the Bias in Estimated Risk Measures. *ASTIN Bulletin: The Journal of the IAA*, 37(2):365–386.
- Kingma, D. P., Salimans, T., and Welling, M. (2015). Variational Dropout and the Local Reparameterization Trick. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Kolouri, S., Pope, P. E., Martin, C. E., and Rohde, G. K. (2019). Sliced Wasserstein Auto-Encoders. In *International Conference on Learning Representations*, volume 5, pages 3481–3499. Curran Associates, Inc.
- Kousky, C. and Cooke, R. (2012). Explaining the Failure to Insure Catastrophic Risks. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 37:206–227.

- Kuiper, W., Yang, W., Hassan, A., Ng, Y., Bidkhori, H., Blanchet, J., and Tarokh, V. (2024). Distributionally Robust Optimization as a Scalable Framework to Characterize Extreme Value Distributions. In Kiyavash, N. and Mooij, J. M., editors, *The Conference on Uncertainty in Artificial Intelligence*, volume 244. PMLR.
- Kupper, M. and Schachermayer, W. (2009). Representation results for law invariant time consistent functions. *Mathematics and Financial Economics*, 2(3):189–210.
- L.A., P. and Bhat, S. P. (2022). A Wasserstein Distance Approach for Concentration of Empirical Risk Estimates. *Journal of Machine Learning Research*, 23(238):1–61.
- Lam, H. and Mottet, C. (2017). Tail analysis without parametric models: A worst-case perspective. *Operations Research*, 65(6):1696–1711.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. (2023). On Tilted Losses in Machine Learning: Theory and Applications. *Journal of Machine Learning Research*, 24(142):1–79.
- Lim, A. E. B. and Shanthikumar, J. G. (2007). Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research*, 55(2):198–214.
- Linnerooth-Bayer, J. and Hochrainer-Stigler, S. (2015). Financial instruments for disaster risk management and climate change adaptation. *Climatic Change*, 133(1):85–100.
- Lugosi, G. and Mendelson, S. (2019). Mean Estimation and Regression Under Heavy-Tailed Distributions: A Survey. *Foundations of Computational Mathematics*, 19(5):1145–1190.
- Maddison, C. J., Mnih, A., and Teh, Y. W. (2017). The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *International Conference on Learning Representations (ICLR)*, volume 3.
- Marcoux, K. and H Wagner, K. R. (2023). Fifty Years of US Natural Disaster Insurance Policy. *Handbook of Insurance*. Accessed on Dec 19, 2024. <https://www.krhwagner.com/papers/NaturalDisasterInsuranceHandbook.pdf>.
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.
- Markowitz, H. (2014). Mean–variance approximations to expected utility. *European Journal of Operational Research*, 234(2):346–355.
- McNeil, A. J., Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.
- Mohajerin Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1–2):115–166.
- Nair, J., Wierman, A., and Zwart, B. (2022). *The Fundamentals of Heavy Tails: Properties, Emergence, and Estimation*. Cambridge University Press.
- Nass, D., Belousov, B., and Peters, J. (2019). Entropic risk measure in policy search. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1101–1106, Macau, China. IEEE Press.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136.
- Rahimian, H. and Mehrotra, S. (2022). Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization*, 3:1–85.
- Sadana, U., Chenreddy, A., Delage, E., Forel, A., Frejinger, E., and Vidal, T. (2025). A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289.
- Saldi, N., Başar, T., and Raginsky, M. (2020). Approximate Markov-Nash equilibria for discrete-time risk-sensitive mean-field games. *Mathematics of Operations Research*, 45(4):1596–1620.
- Shalev-Shwartz, S. et al. (2012). Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. (2009). *Lectures on Stochastic Programming: Modeling and Theory*. Society for Industrial and Applied Mathematics.
- Siegel, A. F. and Wagner, M. R. (2023). Technical Note—Data-Driven Profit Estimation Error in the Newsvendor Model. *Operations Research*, 71(6):2146–2157.
- Smith, J. E. and Winkler, R. L. (2006). The Optimizer’s Curse: Skepticism and Postdecision Surprise in Decision Analysis. *Management Science*, 52(3):311–322.
- Smith, K. M. and Chapman, M. P. (2023). On Exponential Utility and Conditional Value-at-Risk as risk-averse performance criteria. *IEEE Transactions on Control Systems Technology*, 31(6):2555–2570.
- Svensson, L. E. O. and Werner, I. M. (1993). Nontraded assets in incomplete markets: Pricing and portfolio choice. *European Economic Review*, 37(5):1149–1168.
- Troop, D., Godin, F., and Yu, J. Y. (2021). Bias-corrected peaks-over-threshold estimation of the cvar. In *de Campos, C. and Maathuis, M. H., editors, Uncertainty in Artificial Intelligence*, volume 161, pages 1809–1818. PMLR.
- Van der Vaart, A. W. (2000). *Asymptotic Statistics*. Cambridge University Press.

- Von Neumann, J. and Morgenstern, O. (1944). Theory of games and economic behavior. Princeton University Press.
- Wiesemann, W., Kuhn, D., and Sim, M. (2014). Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376.
- Ye, Q. and Xie, W. (2021). Second-order conic and polyhedral approximations of the exponential cone: Application to mixed-integer exponential conic programs. arXiv preprint arXiv:2106.09123. Accessed on 19 Dec, 2024. <https://arxiv.org/pdf/2106.09123>.