Managing Distributional Ambiguity in Stochastic Optimization through a Statistical Upper Bound Framework

Shixin Liu*and Jian Hu†

Department of Industrial and Manufacturing Systems Engineering, University of Michigan - Dearborn, Dearborn, MI 48128

November 14, 2024

Abstract

Stochastic optimization is often hampered by distributional ambiguity, where critical probability distributions are poorly characterized or unknown. Addressing this challenge, we introduce a new framework that targets the minimization of a statistical upper bound for the expected value of uncertain objectives, facilitating more statistically robust decision-making. Central to our approach is the Average Percentile Upper Bound (APUB), a novel construct that simultaneously delivers a statistically rigorous upper bound for the population mean and a meaningful risk metric for the sample mean. The integration of APUB into stochastic optimization not only fortifies the process against distributional ambiguity but also reinforces key data-driven decision-making attributes, such as reliability, consistency, and comprehensibility. Notably, APUB-enriched optimization problems feature tractability, with particular advantages in two-stage stochastic optimization with random recourse. Empirical demonstrations on two-stage product mix and multi-product newsvendor benchmark problems reveal the benefit of the APUB optimization framework, in comparison with conventional techniques such as sample average approximation and distributionally robust optimization.

Key Words: Distributional Ambiguity, Stochastic optimization, Upper Confidence Bound, Asymptotic Correctness, Asymptotic Consistency, Bootstrap Sampling Approximation

*liushixi@umich.edu

†jianhu@umich.edu

1 Introduction

In both engineering and management, effective decision-making often hinges on the ability to navigate through uncertain parameters. Stochastic optimization offers a modeling framework representing the uncertain parameters as random variables, complete with their respective probabilistic information or distribution. Despite this theoretical foundation, real-world applications grapple with challenges such as limited data, incomplete information, and the inherent complexity of the systems being modeled. As a result, the exact estimation of probability distributions becomes a formidable task. This discrepancy between theoretical models and practical constraints introduces distributional ambiguity — a pervasive challenge in decision-making scenarios. This creates a substantial hurdle known as the Optimizers' Curse (Smith and Winkler, 2006), where optimizing under the influence of this ambiguity may lead to suboptimal solutions. Addressing the Optimizers' Curse becomes imperative in ensuring the practical efficacy of decision-making processes in the face of real-world complexities.

In this paper, we develop a novel approach that integrates statistical inference into stochastic optimization to address the challenges posed by distributional ambiguity. Within the field of statistical inference, frequentists tackle distributional ambiguity by considering sample uncertainty, leading to the concept of the sampling distribution of a point estimator. Utilizing repeated sampling, frequentist methods aim to comprehend the inherent variability in point estimators. This approach facilitates the quantification of uncertainty surrounding an estimator's performance and provides a foundation for constructing upper confidence bounds within a frequentist framework.

A general expectation minimization problem can be formulated as follows:

$$\min_{x\in\mathcal{X}} \mathbb{E}[F(x,\xi)], \tag{EM-M}$$

where $\mathcal{X}$ is a decision region, ξ is a vector-valued random parameter associated with probability measure $\mathbb{P}$, and F is a multivariate cost function. Suppose we lack knowledge of $\mathbb{P}$ but possess a random sample and its associated empirical distribution $\widehat{\mathbb{P}}_N$. With a sufficiently large sample size N, the Sample Average Approximation (SAA) model,

$$\min_{x\in\mathcal{X}} \mathbb{E}\left[F(x,\xi)\ \middle|\ \widehat{\mathbb{P}}_N\right], \tag{SAA-M}$$

can serve a good estimate of EM-M. However, the estimation may exhibit significant bias when N is small. This research aims to develop a $100(1-\alpha)\%$ upper confidence bound for the expectation $\mathbb{E}[F(x,\xi)]$ based on the empirical distribution $\widehat{\mathbb{P}}_N$. Denote by $\mathbb{U}^\alpha\left[\mathbb{E}[F(x,\xi)]\middle|\widehat{\mathbb{P}}_N\right]$ this upper confidence bound and generally describe our data-driven model as follows:

$$\min_{x\in\mathcal{X}} \mathbb{U}^\alpha\left[\mathbb{E}[F(x,\xi)]\middle|\widehat{\mathbb{P}}_N\right]. \tag{UB-M}$$

UB-M incorporates an upper confidence bound into stochastic optimization. This method effectively addresses challenges stemming from distributional ambiguity by leveraging robust statistical techniques. These techniques, specifically designed to accommodate estimation errors, provide a range of values within which the true population mean is likely to reside. By presenting this range, particularly the upper bound, we not only quantify uncertainty but also enhance the reliability of our statistical inferences. Nevertheless, conventional upper confidence bounds, while based on firm statistical concepts, may be non-convex or overly conservative due to a lack of data sensitivity. We will address these limitations by proposing a practical statistical upper bound applicable to optimization.

In the existing literature, the recognition of distributional ambiguity highlights the necessity of deploying distributionally robust optimization (DRO) strategies (Rahimian and Mehrotra (2019); Lin et al. (2022) and references therein). These strategies are crafted to guide decision-making processes that excel across a spectrum of plausible distributional assumptions, avoiding reliance on a single assumed distribution. Two predominant DRO approaches involve representing moment-based and discrepancy-based ambiguity sets for the distributions. The moment-based approach characterizes uncertainty by imposing constraints on the moments (such as mean and variance) of the distribution (Calafiore and Ghaoui, 2006; Delage and Ye, 2010; Wiesemann et al., 2014). On the other hand, the discrepancy-based approach focuses on measuring the difference between the true distribution and a candidate distribution within the set. Examples of discrepancy-based ambiguity sets include those based on ϕ-divergence (Read and Cressie, 2012; Ben-Tal et al., 2013; Bayraksan and Love, 2015) and the Wasserstein metric (Mohajerin Esfahani and Kuhn, 2018; Blanchet and Murthy, 2019; Xie, 2020; Duque et al., 2022; Gao and Kleywegt, 2023). By embracing the inherent ambiguity associated with underlying probability distributions, DRO empowers decision-makers to formulate strategies that demonstrate resilience and effectiveness under diverse scenarios. Notably, the Wasserstein metric-based DRO has gained popularity in various fields due to its appealing properties, which include finite-sample guarantees and asymptotic consistency.

The advantages of using DRO include: (i) DRO frameworks are designed to hedge against model uncertainty. They account for variations in the data distribution, making solutions more reliable under different future scenarios. (ii) Because DRO considers a range of possible distributions, it often results in solutions that generalize better to unseen data, as opposed to solutions based solely on empirical risk minimization. (iii) DRO can be tailored to various levels of conservatism depending on the choice of ambiguity set (the set of distributions considered). This flexibility allows tuning the robustness to specific application needs. (iv) By explicitly incorporating ambiguity into the optimization problem, DRO provides a systematic way to handle cases where data is scarce or noisy, improving decision-making under uncertainty. (v) DRO often comes with strong theoretical guarantees regarding performance bounds, which enhances confidence in the solutions produced. While DRO offers significant benefits, it also presents several limitations. The computational intensity of DRO models escalates with the model size, making large-scale optimization daunting and time-consuming. DRO's pivotal strategy of hedging against the worst-case distribution can also lead to overly conservative outcomes, which, while safe, may not be cost-optimal or practical under typical operating conditions. Furthermore, the practicality of DRO hinges on the assumption of a well-defined ambiguity set; however, constructing such a set with accuracy can be a complex endeavor in data-scarce environments. Additionally, the nature of DRO as a 'black-box' means that it typically does not provide insight into the specific probability distributions it encapsulates, which can be a significant limitation for decision-makers seeking transparent and interpretable models for policy development and risk management in practice.

In contrast to the conventional practice of defining an ambiguity set of distributions in Distributionally Robust Optimization (DRO), UB-M proposes an innovative methodology that intersects two significant areas: statistical inference and stochastic programming. A notable advantage of UB-M is its introduction of the following well-regarded merits of statistical inference into the optimization process: (i) Upper confidence bounds provide a range of values that likely contain the true population mean, giving a clear measure of the uncertainty associated with estimates derived from limited data. (ii) Upper confidence bounds explicitly account for sampling and estimation errors, providing a more realistic picture of the potential variability in the estimates. (iii) Upper confidence bounds can be adjusted for different confidence levels (e.g., 90%, 95%), allowing researchers and practitioners to balance between precision and confidence based on the specific requirements of the situation. (iv) The adaptive nature of upper confidence bounds avoids over-conservatism. As more data becomes available, the intervals can be recalculated, continuously refining the estimates and improving the robustness of decisions based on them. (v) By communicating the degree of uncertainty, upper confidence bounds make the results more transparent and easier to interpret for stakeholders, fostering trust in the analysis. (vi) Upper confidence upper can be applied across different statistical and stochastic models, making them a versatile tool in various domains and research settings.

By minimizing an upper confidence bound rather than relying on a single point estimate, UB-M enables more informed and robust decision-making. In this study, we address the following three perspectives in optimization:

1. **Design a statistically robust yet consistent data-driven mechanism.** The primary objective is to create a versatile mechanism that takes advantage of the available data to inform modeling choices, encouraging performance and reliability without succumbing to undue conservatism. This approach aspires to establish consistency in decision-making, where robustness is harmonized with adaptability to data availability.

2. **Interpret outcomes within statistical frameworks for decision-making.** We intend to construct a bridge between complex statistical models and practical decision-making processes. Our aim is to introduce reliability levels cognizant of statistical significance and confidence intervals, ultimately rendering these models more intuitive and usable.

3. **Enhance computational efficiency for large-scale optimization.** Acknowledging the importance of computational prowess, this line of inquiry targets the advancement of computational methodologies for large-scale stochastic programming. This research is oriented toward improving algorithmic efficiency and practical applicability for sizable optimization challenges. Especially, we will develop an L-shaped method to solve two-stage stochastic programming problems with random recourse. Notably, these problems often present challenges

that are insurmountable by conventional DRO techniques.

1.1 Literature review of statistical confidence interval

One-sided and two-sided confidence intervals, a well-explored domain in statistical theory, are renowned for their robust asymptotic correctness, accuracy, and consistency, directly linked to sample size. A confidence interval is deemed first-order accurate if its confidence level error is within the inverse of the square root of the sample size, and second-order accurate if within the inverse of the sample size (Vaart, 1998). The inherent convergence with the sample size is a notable feature of those upper confidence bounds. Moreover, statistical upper bounds provide additional advantages rooted in their well-established theoretical foundations. The definitions and concepts underpinning those bounds have undergone thorough scrutiny and refinement in the field of statistics. This rigorous academic exploration has yielded a comprehensive and uniform framework, facilitating a nuanced understanding and practical application of those bounds. Consequently, incorporating a statistical upper bound into stochastic optimization offers not only robustness and reliability but also interpretation of our proposed optimization framework.

The most classical frequentist asymptotic approach, which utilizes the sample mean, standard deviation, and normal approximation, has been extensively discussed (see Devore (2009) and references therein). This method is particularly favored in practice for sample sizes larger than 30, offering a straightforward yet effective means of estimating confidence intervals (Hazra, 2017). It hinges on the Central Limit Theorem (CLT), which posits that the distribution of a sample mean approximates a normal distribution as the sample size increases, regardless of the population's distribution. Efron's bootstrap confidence interval (Efron, 1981), another method for constructing statistical upper bounds, employs the $100(1-\alpha)$-th percentile of the bootstrap distribution of a sample mean. Both these methods achieve first-order accuracy. Building upon these concepts, (Efron, 1987) developed bias-corrected and accelerated bootstrap confidence interval ($\mathrm{BC_a}$) achieving a second-order accuracy. This approach fine-tunes the nominal level as a function of the sample size. This advancement marks a significant step towards more precise confidence interval estimation, especially in cases where first-order methods might not suffice due to smaller sample sizes or more complex data structures. In machine learning, the construction of upper confidence bounds often leverages concentration inequalities to ensure robustness against data variability. Hoeffding-type bounds (Auer et al., 2002) are popular for their simplicity and effectiveness in bounding the sum of bounded random variables, particularly useful in scenarios with limited prior knowledge about data distributions. Empirical Bernstein-type bounds (Mnih et al., 2008) offer improvements over Hoeffding's approach by incorporating sample variance, making them more adaptable to data with varying degrees of variability. This type of bound is particularly advantageous in dealing with heteroskedastic data, where the variance is not constant. The self-normalized bounds (Abbasi-yadkori et al., 2011), on the other hand, are designed to handle the challenges of auto-correlated data, common in time-series analysis. These bounds normalize the sum of random variables by their cumulative variance, offering a more dynamic approach to uncertainty quantification in sequential decision processes. While these bounds are grounded in robust statistical principles, they each have limitations: they may be non-convex or not sufficiently data-driven, leading to challenges in optimization or an overly conservative nature.

1.2 Contributions and Organization of this Paper

The primary contributions of this paper are summarized as follows:

- We introduce the Average Percentile Upper Bound (APUB), a novel statistical construct that serves as an upper bound for population means and a risk metric for sample means. The robust statistical foundations of APUB are established through rigorous proofs of its asymptotic correctness and consistency, offering a reliable basis for its applications in data-driven decision-making.

- The innovative integration of APUB into stochastic optimization frameworks mitigates the ambiguity stemming from sparse data in probability distributions. The theoretical properties of APUB are adeptly applied to a new optimization framework, simultaneously ensuring model reliability and interpretability while reducing over-conservatism. This synergy narrows the gap between statistical upper bounds and stochastic optimization, fostering theoretical advancement and practical utility.

- The development of a bootstrap sampling approximation method tailored for solving APUB-embedded optimization frameworks is addressed, with a particular focus on two-stage linear stochastic optimization with random recourse. Practically, we present a cut generation algorithm for solving large-scale cases and an L-shaped method for two-stage stochastic programming problems with random recourse.

The rest of this paper is organized as follows: Section 2 introduces the concept of APUB. Specifically, Section 2.1 formally defines APUB and explores its statistical implications, while Section 2.2 provides a thorough examination of the asymptotic properties of APUB, including the proofs of asymptotic correctness and consistency. Section 3 details the integration of APUB into stochastic optimization. Section 3.1 establishes the asymptotic consistency of our optimization framework, and Section 3.2 addresses its asymptotic correctness. Section 4 develops a bootstrap sampling approximation approach to solve our optimization framework, particularly in the context of the two-stage stochastic programming. We develop practical solution methods for solving large-scale cases. Section 5 presents a comprehensive numerical analysis, applying the proposed framework across a variety of classical stochastic optimization scenarios. In Section A.5, we rigorously prove the the asymptotic correctness of the optimization framework. The paper concludes with Section 6, summarizing key findings and contributions.

2 Average Percentile Upper Bound

We establish formal definitions for the key concepts employed in this paper. Consider an induced probability space $(\Xi, \mathfrak{B}, \mathbb{P})$, where $\Xi \subseteq \mathbb{R}^{d_\xi}$ is the support of a random vector, $\mathfrak{B}$ is the Borel σ-algebra, and $\mathbb{P}$ is a probability measure. Let $(\xi_1, \ldots, \xi_N) \sim \mathbb{P}$ indicate an independent and identically distributed (i.i.d) random sample with a size of N generated from $(\Xi, \mathfrak{B}, \mathbb{P})$. The empirical distribution associated with the random sample is represented as

$$\widehat{\mathbb{P}}_N := \frac{1}{N} \sum_{n=1}^{N} \delta_{\xi_n},$$

where δ_{ξ_n} is the Dirac delta function at ξ_n. As N increases to infinity, we have a sample path $(\xi_1, \xi_2, \ldots)$. Without loss of generality, we ignore the decision variable x and focus our discussion on a measurable cost function $F : \Xi \to \mathbb{R}$ in this section. Denote by $\mu := \mathbb{E}[F(\xi)]$ the population mean and by $\sigma^2 := \mathbb{E}[(F(\xi) - \mu)^2]$ the population variance. We assume μ and σ to be finite in the late statement.Let $\widehat{\mu}_N := \mathbb{E}[F(\zeta)|\widehat{\mathbb{P}}_N]$ be the sample mean and $\widehat{\sigma}_N^2 := \mathbb{E}[(F(\zeta) - \widehat{\mu}_N)^2|\widehat{\mathbb{P}}_N]$ be the asymptotic sample variance, where $\zeta \sim \widehat{\mathbb{P}}_N$.

2.1 Concept of Average Percentile Upper Bound

Using the bootstrap percentile method, Efron (1981) presents a $100(1-\alpha)\%$ bootstrap-based upper confidence bound for μ as

$$\inf \left\{ t \in \mathbb{R} \ : \Pr\left(\mu^* \le t \mid \mathbb{P}_*\right) \ge 1 - \alpha \right\},$$

where $\mathbb{P}_*$ is a bootstrap distribution and μ^* is a bootstrap estimator of μ. The above equation can be explicitly represented as the limit of

$$U_M^\alpha := \inf \left\{ t \in \mathbb{R} \ : \ \frac{1}{M} \sum_{m=1}^{M} \mathbb{1} \left\{ \frac{1}{N} \sum_{n=1}^{N} F(\zeta_{m,n}) \le t \right\} \ge 1 - \alpha \right\},$$

where $\mathbb{1}\{\cdot\}$ is the indicator function and $(\zeta_{m,1}, \ldots, \zeta_{m,N}) \sim \widehat{\mathbb{P}}_N$, for $m = 1, \ldots, M$, are bootstrap samples. The following proposition depicts the limit of U_M^α.

Proposition 2.1. *Let*

$$\mathbb{U}^\alpha_{\mathit{Efron}}[\mu|\widehat{\mathbb{P}}_N] := \inf \left\{ t \in \mathbb{R} : \ \Pr\left(\frac{1}{N} \sum_{n=1}^{N} F(\zeta_n) \le t \ \middle| \ \widehat{\mathbb{P}}_N \right) \ge 1 - \alpha \right\},$$

where $\zeta_n \sim \widehat{\mathbb{P}}_N$ for $n = 1, \ldots, N$. Then, as $M \to \infty$, $U_M^\alpha \to \mathbb{U}^\alpha_{\mathit{Efron}}[\mu|\widehat{\mathbb{P}}_N]$ w.p.1 (for ζ).

While $\mathbb{U}^\alpha_{\text{Efron}}[\mu|\widehat{\mathbb{P}}_N]$, the percentile-based upper bound, is satisfactory in many statistical analyses, it is non-convex and difficult to control/optimize for highly skewed distributions, which are regarded as inferior properties in the realm of optimization. Therefore, we extend Efron's upper bound by averaging over the values to the right of the $100(1 - \alpha)$-th percentile.

Definition 2.2. *The average percentile upper bound for μ with a nominal level $(1-\alpha)$ is denoted as*

$$\mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N] := \frac{1}{\alpha}\int_0^{\alpha} \mathbb{U}^{\tau}_{Efron}[\mu|\widehat{\mathbb{P}}_N]d\tau. \tag{APUB}$$

We can interpret Efron's upper bound, alternatively in the realm of risk management and decision making, as an approximation of the Value at Risk (VaR) of $\widehat{\mu}_N$ by substituting $\widehat{\mathbb{P}}_N$ for the true distribution $\mathbb{P}$ of the random variable ξ_n in the following VaR equation:

$$\mathrm{VaR}_{\alpha}(\widehat{\mu}_N) = \inf\left\{t \in \mathbb{R}:\ \Pr\left(\frac{1}{N}\sum_{n=1}^{N} F(\xi_n) \le t\right) \ge 1-\alpha\right\}.$$

Analogously, APUB approximates the Conditional Value at Risk (CVaR) of $\widehat{\mu}_N$,

$$\mathrm{CVaR}_{\alpha}(\widehat{\mu}_N) = \frac{1}{\alpha}\int_0^{\alpha} \mathrm{VaR}_{\tau}(\widehat{\mu}_N)d\tau.$$

APUB serves a dual purpose: as an upper bound for the population mean in statistics and as an approximate risk measure for the sample mean in risk assessment. As a risk measure, it primarily focuses on approximating the tail distribution of the potential estimation error of the population mean, which could result from an inadequacy of sample points. Furthermore, APUB complies with fundamental properties of a coherent risk measure, such as sub-additivity, homogeneity, convexity, translational invariance, and monotonicity. These characteristics make APUB a good candidate to be applied to stochastic optimization under distributional ambiguity, particularly in scenarios requiring solvability, such as two-stage stochastic optimization with random recourse. Analogous to Theorem 1 in Rockafellar and Uryasev (2000) (Theorem B.1), the following proposition provides an alternative representation for APUB.

Proposition 2.3.

$$\begin{aligned}\mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N] &= \inf_{t\in\mathbb{R}}\ \left\{t + \frac{1}{\alpha}\mathbb{E}\left[\left[\frac{1}{N}\sum_{n=1}^{N} F(\zeta_n) - t\right]_+ \ \Bigg|\ \widehat{\mathbb{P}}_N\right]\right\} \\ &= \inf_{t\in\mathbb{R}}\ \left\{t + \frac{1}{\alpha}\boldsymbol{\int}\left[\frac{1}{N}\sum_{n=1}^{N} F(\zeta_n) - t\right]_+ \prod_{n=1}^{N}\widehat{\mathbb{P}}_N(d\zeta_n)\right\},\end{aligned}$$

where the bold integral symbol means an N-fold integral over the N-fold product of $\widehat{\mathbb{P}}_N$.

Remark 2.4. *By Proposition 2.3, we have that $\mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N]$ monotonically decreases in $\alpha \in (0,1]$ w.p.1. This implies that, for $\alpha \in (0,1]$,*

$$\mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N] \ge \mathbb{U}^{1}_{APUB}[\mu|\widehat{\mathbb{P}}_N] = \mathbb{E}\left[\frac{1}{N}\sum_{n=1}^{N} F(\zeta_n)\ \Bigg|\ \widehat{\mathbb{P}}_N\right] = \frac{1}{N}\sum_{n=1}^{N}\mathbb{E}\left[F(\zeta_n)\ \Big|\ \widehat{\mathbb{P}}_N\right] = \widehat{\mu}_N, \quad \textit{w.p.1.}$$

The quality of a statistical upper bound refers to the rate of its true coverage probability increasing beyond the nominal level $(1-\alpha)$ as the sample size grows. Example 2.5 illustrates the following two attractive asymptotic characteristics of APUB. A theoretical discussion is given in Section 2.2.

1. *Asymptotic Correctness* (defined in Vaart (1998, Section 23.3)): A statistical upper bound $\mathbb{U}^\alpha[\mu|\widehat{\mathbb{P}}_N]$ for μ is correct at level $(1-\alpha)$ up to κth-order if its coverage probability

$$\Pr\left(\mu \leq \mathbb{U}^\alpha[\mu|\widehat{\mathbb{P}}_N]\right) \geq (1-\alpha) + O(N^{-\kappa/2}).$$

If the equality holds, one says that $\mathbb{U}^\alpha[\mu|\widehat{\mathbb{P}}_N]$ is κth-order accurate (Hall, 1986). Efron's upper bound is first-order accurate and APUB is first-order correct. Notably, the terms 'asymptotic correctness' and 'asymptotic accuracy' are different. Specifically, the asymptotic correctness implies that, when the sample size N is sufficiently large, the nominal level serves as a reliable lower bound for the coverage probability.

2. *Asymptotic consistency*: APUB converges to the population mean w.p.1 as the sample size increases to ∞. This attribute ensures that APUB is a consistent estimator for the population mean.

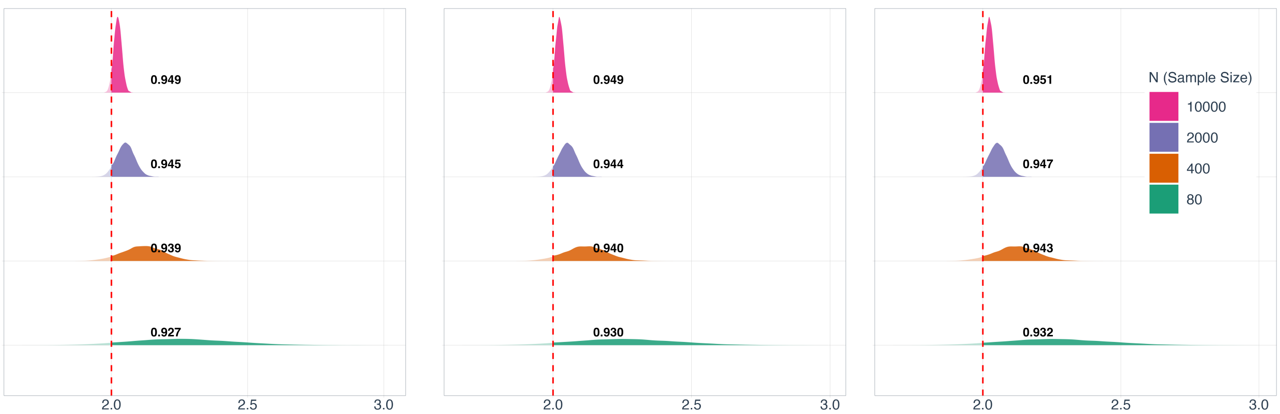

(a) Standard Upper Bound (b) Efron's Upper Bound (c) APUB

Figure 1: The comparison between APUB, Efron's upper bound, and the standard large-sample upper bound.

Example 2.5. *Let $\xi \sim$ Gamma$(2,1)$ and $F(\xi) = \xi$. So the population mean $\mu = 2$. We compare APUB with Efron's percentile-based upper bound and the standard large-sample upper bound given as $\widehat{\mu}_N + z_\alpha \widehat{\sigma}_N/\sqrt{N}$, where z_α denotes z critical value. In order to estimate the probability density functions (pdf) of three upper bounds, we performed a Monte Carlo simulation with $\alpha = 0.05$ while allowing the sample sizes, N, to vary from 80 to 10,000.*

As illustrated in Figure 1, the coverage probability is essentially the area to the right of the vertical dotted line at $\mu = 2$ in a pdf. Our results show that as N increases, the coverage probabilities for both the large-sample and Efron's upper bounds get closer to $(1-\alpha) = 0.95$. This demonstrates the asymptotic accuracy of these two types of bounds. However, this is in stark contrast to APUB which doesn't have asymptotic accuracy. In fact, as N approaches ∞, the coverage probability of APUB can grow beyond $(1-\alpha)$. Moreover, this growth rate is observed to be more rapid than the other two bounds, which brings attention to the unique nature of APUB.

Furthermore, all three upper bounds exhibit asymptotic consistency. As N increases, they all converge to $\mu = 2$ w.p.1. This essentially means they become more precise as more data is collected. By examining the pdf curves, it is apparent that they narrow and concentrate more intensely around μ, which visually indicates this trend.

2.2 Asymptotic Characteristics of APUB

We now theoretically discuss the asymptotic correctness and consistency of APUB. The following proposition shows its asymptotic correctness.

Proposition 2.6. *Suppose that the skewness $\mathbb{E}[F(\xi) - \mu]^3/\sigma^3 < \infty$. Then, for a fixed nominal level $1-\alpha$, $\mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N]$ is first-order asymptotically correct, i.e.,*

$$\Pr\left(\mu \leq \mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N]\right) \geq (1-\alpha) + O(N^{-1/2}).$$

Proposition 2.6 establishes the asymptotic correctness of APUB, which guarantees that the nominal level is a conservative boundary for the actual coverage probability. This attribute confirms that APUB is an effective upper-bound statistic, especially valuable for its robust response to distributional ambiguity encountered with limited sample data. Considering the uncertainty diminishes along with an increase in the sample size, we next show that APUB is a consistent estimator for the population mean.

Theorem 2.7. *For any $\alpha \in (0, 1]$, as $N \to \infty$,*

$$\mathbb{U}^{\alpha}_{APUB}[\mu|\widehat{\mathbb{P}}_N] \to \mu, \text{ w.p.1.}$$

3 Optimization with APUB

In this section we apply APUB to stochastic optimization problems. In the context of optimization, we have a decision region $\mathcal{X} \subseteq \mathbb{R}^{d_x}$ and let the cost function $F(x,\xi) : \mathcal{X} \times \Xi \mapsto \mathbb{R}$ be $\mathfrak{B}$-measurable for all $x \in \mathcal{X}$. Denote the mean and standard deviation of $F(x,\xi)$ by $\mu(x)$ and $\sigma(x)$ respectively. The UB-M framework using APUB is written as

$$\min_{x\in\mathcal{X}} \mathbb{U}^{\alpha}_{\text{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]. \tag{APUB-M}$$

Let $\widehat{\vartheta}^{\alpha}_N$ be the optimal value of APUB-M and $\widehat{\mathcal{S}}^{\alpha}_N$ denote the set of its optimal solutions. Also, denote by ϑ^* the optimal value of EM-M and by $\mathcal{S}$ the set of its optimal solutions. By Remark 2.4, we know that $\widehat{\vartheta}^{\alpha}_N$ decreases in $\alpha \in (0, 1]$ w.p.1 and $\widehat{\vartheta}^{1}_N$ is the optimal value of SAA-M.

By applying the statistical technique, we expect to significantly reduce the impact of distributional ambiguity. This methodology has the following benefits, poised to yield insightful solutions to our three research perspectives mentioned in Section 1.

(i) The coverage probability of APUB can be employed to evaluate enhancements in the resilience and reliability of APUB-M when contending with distributional ambiguity. In practice, the probability is approximated by the nominal level $(1-\alpha)$. The nominal level of APUB, which adheres to asymptotic correctness, provides a reliable lower bound for the coverage probability.

(ii) APUB exemplifies asymptotic consistency, rendering it ideal for a data-driven method. As ambiguity in distribution decreases with larger sample sizes, the influence of APUB consequently lessens. As a result, APUB-M avoids conservatism.

(iii) APUB can be interpreted as a risk measure, namely, the CVaR of a sample mean. CVaR is widely recognized as one of the most popular risk measures by many disciplines (Artzner et al., 1999; Rockafellar and Uryasev, 2000). In this way, APUB-M is designed to model risk aversion, aiming to mitigate the inherent risk in SAA-M caused by the limited number of sample points. APUB adheres to fundamental properties expected of a coherent risk measure, including sub-additivity, homogeneity, convexity, translational invariance, and monotonicity. These inherent characteristics render APUB-M not only coherent but also readily solvable, particularly in the context of two-stage stochastic programming with random recourse.

We now present some mild assumptions as follows.

Assumption A. *There exists a compact set $\mathcal{K} \subseteq \mathcal{X}$ such that:*

(A1) $\mathcal{S} \subseteq \mathcal{K}$*;*

(A2) $\widehat{\mathcal{S}}_N^{\alpha} \subseteq \mathcal{K}$ *w.p.1 for sufficiently large N and $\alpha \in (0,1]$.*

Assumption A is frequently encountered in the literature pertaining to the asymptotic analysis of the SAA method (Birge and Louveaux, 2011; Shapiro et al., 2021). This assumption posits that it is adequate to confine the examination of decision properties to the compact set $\mathcal{K}$. For the purposes of the discussion in the remainder of Section 3, we proceed under the premise that the decision space is indeed $\mathcal{K}$, a simplification that does not limit the generality of our analysis.

Assumption B. *There exists an open convex hull $\mathcal{N}$ containing $\mathcal{K}$ such that:*

(B1) $F(x,\xi)$ *is convex on $\mathcal{N}$ for each $\xi \in \Xi$;*

(B2) $\mu(x)$ *and* $\sigma(x)$ *are finite for all $x \in \mathcal{N}$.*

In Sections 3.1 and 3.2, we examine asymptotic characteristics of APUB-M, focusing on its data-driven nature which includes aspects such as consistency, reliability, and ease of interpretation.

3.1 Asymptotic Consistency

In optimization, asymptotic consistency refers to the convergence of the optimal value and optimal solution set of APUB-M with their counterparts in EM-M w.p.1 as the sample size increases. The following theorem exhibits the asymptotic consistency of APUB-M.

Theorem 3.1. *Suppose Assumptions A and B hold. Then, for any given $\alpha \in (0, 1]$, as $N \to \infty$,*

$$\widehat{\vartheta}_N^\alpha \to \vartheta^*, \qquad \text{and} \qquad \mathbb{D}(\widehat{\mathcal{S}}_N^\alpha, \mathcal{S}) := \sup_{y \in \widehat{\mathcal{S}}_N^\alpha} \inf_{z \in \mathcal{S}} \|y - z\| \to 0 \qquad \text{w.p.1.}$$

Remark 3.2. *Unlike DRO approaches that require additional parameter adjustments based on the sample size to achieve data-driven objectives, the sample size itself is the unique factor to determine the convergence of APUB-M. This characteristic offers a more consistent data-driven approach in practice. As ambiguity in distribution decreases with larger sample sizes, the influence of APUB consequently lessens. As a result, APUB-M avoids over-conservatism.*

3.2 Asymptotic Correctness

Mohajerin Esfahani and Kuhn (2018) introduce the concept of reliability for a certain optimal solution in DRO approaches. The reliability refers to the probability that the optimal value of a DRO model exceeds the expected cost of the system at its optimal solution in true scenarios. We extend this concept to the entire optimal solution set, which in our case is called the coverage probability of the general UB-M framework. Denote a probability function of a given subset $S \subseteq \mathcal{X}$ as

$$\beta(\vartheta, S) := \Pr\left(\vartheta \geq \max_{x \in S} \mu(x)\right). \tag{1}$$

Let $\overline{\vartheta}_N^\alpha$ and $\overline{\mathcal{S}}_N^\alpha$ be the optimal value and optimal solution set of UB-M, respectively. The coverage probability of UB-M is $\beta(\overline{\vartheta}_N^\alpha, \overline{\mathcal{S}}_N^\alpha)$, which measures the chance that $\overline{\vartheta}_N^\alpha$ can serve as an upper bound of the actual performance of UB-M across all optimal solutions. In the following, we define the asymptotic correctness of UB-M.

Definition 3.3. *UB-M is asymptotically correct if the limit of its coverage probability serves as a lower bound for the nominal level as*

$$\lim_{N \to \infty} \beta(\overline{\vartheta}_N^\alpha, \overline{\mathcal{S}}_N^\alpha) \geq (1 - \alpha)$$

Defined on the entire optimal solution set, the concept of asymptotic correctness is stricter than the reliability concerning a certain optimal solution. If a UB-M framework is asymptotically correct, we have that, for any $\bar{x}_N \in \overline{\mathcal{S}}_N^\alpha$,

$$\lim_{N \to \infty} \beta(\overline{\vartheta}_N^\alpha, \{\bar{x}_N\}) \geq \lim_{N \to \infty} \beta(\overline{\vartheta}_N^\alpha, \overline{\mathcal{S}}_N^\alpha) \geq (1 - \alpha).$$

In the subsequent statement, we refer to $\beta(\overline{\vartheta}_N^\alpha, \{\bar{x}_N\})$ as the coverage probability of UB-M concerning $\bar{x}$, or simply the coverage probability at $\bar{x}_N$. Thus, we can say that the asymptotic correctness of UB-M guarantees the asymptotic correctness at any optimal solution. The following theorem shows the asymptotic correctness of APUB-M.

Theorem 3.4. *Suppose that Assumption A and B hold. Assume that (i) there exists $x_0 \in \mathcal{K}$ such that $\mathbb{E}\left[|F(x_0, \xi)|^3\right] < \infty$; (ii) for any $x, y \in \mathcal{K}$, $|F(x, \xi) - F(y, \xi)| < L(\xi)\|x - y\|$, where $\mathbb{E}\left[|L(\xi)|^3\right] <$*

∞; and (iii) $\sigma(x)$ and $\sigma^{-1}(x)$ are bounded on $x \in \mathcal{K}$. Then, APUB-M is asymptotically correct for $\alpha \in (0,1]$, i.e.,

$$\lim_{N\to\infty} \beta(\widehat{\vartheta}_N^\alpha, \widehat{\mathcal{S}}_N^\alpha) \geq (1-\alpha).$$

Remark 3.5. *The attribute of asymptotic correctness lends APUB-M interpretability in the context of statistics. This means that the decision-maker can intuitively set the desired reliability level of APUB-M by selecting an appropriate nominal level. Section 5 provides a numerical demonstration of how this model interpretability confers an advantage.*

Remark 3.6. *Recall ϑ^* is the optimal value of EM-M. Since $\vartheta^* \leq \mu(x)$ for all $x \in \widehat{\mathcal{S}}_N^\alpha$, we obtain that*

$$\lim_{N\to\infty} \Pr\left(\widehat{\vartheta}_N^\alpha \geq \vartheta^*\right) \geq \lim_{N\to\infty} \beta(\widehat{\vartheta}_N^\alpha, \widehat{\mathcal{S}}_N^\alpha) \geq (1-\alpha).$$

This implies that the nominal level approximately represents the lower bound of the probability that $\widehat{\vartheta}_N^\alpha$ serves as an upper bound for ϑ^*.

4 Solution Method Based on Bootstrap Sampling Approximation

By Proposition 2.3, we rewrite APUB-M as

$$\min_{(x,t)\in\mathcal{X}\times\mathbb{R}} \quad t + \frac{1}{\alpha}\int \left[\frac{1}{N}\sum_{n=1}^N F(x,\zeta_n) - t\right]_+ \prod_{n=1}^N \widehat{\mathbb{P}}_N(d\zeta_n). \tag{2}$$

Recall that the empirical distribution $\widehat{\mathbb{P}}_N$ is associated with the original random sample $(\xi_1,\ldots,\xi_N)$. We generate a bootstrap sample point $(\zeta_1,\ldots,\zeta_N)$ from $\widehat{\mathbb{P}}_N$, Let us count the number of times the specific original sample point ξ_n appears in $(\zeta_1,\ldots,\zeta_N)$, which is denoted by V_n. Note that $V_n \geq 0$ and $\sum_{n=1}^N V_n = N$. We can interpret $(V_1,\ldots,V_N)$ as a random vector with a multinomial distribution with index N and parameter vector $(1/N,\ldots,1/N)$. Its probability mass function (pmf), in the support $\mathfrak{V} := \{(v_1,\ldots,v_N) \in \mathbb{Z}_+^N : \sum_{n=1}^N v_n = N\}$, is written as

$$\mathfrak{M}_N(v_1,\ldots,v_N) = \frac{N!}{v_1!\cdots v_N!}\left(\frac{1}{N}\right)^{v_1}\cdots\left(\frac{1}{N}\right)^{v_N} = \frac{N!}{N^N v_1!\cdots v_N!}. \tag{3}$$

This observation motives us to reformulate (2) as

$$\min_{(x,t)\in\mathcal{X}\times\mathbb{R}} \quad t + \frac{1}{\alpha}\sum_{(v_1,\ldots,v_N)\in\mathfrak{V}} \left[\frac{1}{N}\sum_{n=1}^N v_n F(x,\xi_n) - t\right]_+ \mathfrak{M}_N(v_1,\ldots,v_N). \tag{4}$$

The multinomial distribution denoted as in (3) comprises $\binom{2N-1}{N}$ random scenarios. Addressing its large-scale nature becomes imperative. To tackle the complexity arising from the sheer number of scenarios, we can leverage the sampling approximation method to solve

$$\min_{(x,t)\in\mathcal{X}\times\mathbb{R}} \quad t + \frac{1}{\alpha M}\sum_{m=1}^M \left[\frac{1}{N}\sum_{n=1}^N V_{m,n} F(x,\xi_n) - t\right]_+, \tag{BP-APUB-M}$$

where $(V_{m,1}, \ldots, V_{m,N})$, for $m = 1, \ldots, M$, are a random sample of a size M generated from the multinomial distribution. Achieving a satisfactory approximation, BP-APUB-M actually needs a significantly smaller number $M \ll \binom{2N-1}{N}$ of sample points in practice. For detailed discussions on the asymptotic behaviors of the SAA method, which has been demonstrated to exhibit fast convergence rates under mild assumptions, we refer readers to Shapiro et al. (2021) and the references therein.

Algorithm 1 Cut Generation Algorithm for BP-APUB-M.

Step 1: Initialization

Let $k = 0$. Find a feasible solution $\hat{x}_0 \in \mathcal{X}$ and let $\hat{\eta}_0 = -\infty$.

Step 2: Solve Master problem

If $k = 0$, go to Step 3. Otherwise, find an optimal solution $(\hat{x}_k, \hat{\eta}_k)$ of the following problem:

$$\min_{x,\eta} \ \eta \tag{5a}$$

$$\text{s.t. } \eta \geq w_j(x), \quad j = 1, \ldots, k, \tag{5b}$$

where $w_j(\cdot)$ is defined in (7).

Step 3: Generate Cuts

Sort the following sequence

$$r_m(\hat{x}_k) := \frac{1}{N}\sum_{n=1}^{N} V_{m,n} F(\hat{x}_k, \xi_n), \quad m = 1, \ldots, M, \tag{6}$$

in ascending order and label the sorted sequence as $r_{(1)_k}(\hat{x}_k) \leq \cdots \leq r_{(M)_k}(\hat{x}_k)$. Here, $(\cdot)_k : \{1, \ldots, M\} \to \{1, \ldots, M\}$ is a one-to-one mapping related to $\hat{x}_k$. Let $J := \lceil (1-\alpha)M \rceil$ and

$$w_{k+1}(x) := \left(1 - \frac{M-J}{\alpha M}\right) r_{(J)_k}(x) + \frac{1}{\alpha M}\sum_{m=J+1}^{M} r_{(m)_k}(x). \tag{7}$$

If $\hat{\eta}_k \geq w_{k+1}(\hat{x}_k)$, stop; $\hat{x}_k$ is an optimal solution. Otherwise, add w_{k+1} to the constraint set (5b), set $k = k + 1$, and return to Step 2.

4.1 Cut Generation Method

We now propose a cut generation solution approach to BP-APUB-M. The details of our method are provided in Algorithm 1. Denote the objective value of BP-APUB-M at a given $x \in \mathcal{X}$ as

$$\widetilde{\vartheta}(x) := \min_{t\in\mathbb{R}} \ t + \frac{1}{\alpha M}\sum_{m=1}^{M}\left[\frac{1}{N}\sum_{n=1}^{N} V_{m,n} F(x, \xi_n) - t\right]_+ = \min_{t\in\mathbb{R}} \ t + \frac{1}{\alpha M}\sum_{m=1}^{M}\left[r_m(x) - t\right]_+, \tag{8}$$

where the function r_m is defined in (6). Let $J = \lceil (1-\alpha)M \rceil$, where the ceiling function $\lceil z \rceil$ gives the smallest integer greater than or equal to z. Definition 2.2 indicates that APUB is derived by taking a weighted average of the "extreme" losses in the tail of the distribution of possible outcomes, beyond its $(1-\alpha)$-th percentile which is the cutoff point of Efron's percentile-based upper bound

($\mathbb{U}^{\alpha}_{\text{Efron}}[\mu|\widehat{\mathbb{P}}_N]$). In our case, the tail of the sequence $r_1(x), \dots, r_M(x)$ beyond its $(1-\alpha)$-th percentile has $M-J+1$ items.

Consider a combination without replacement which firstly selects an element in $\{1, \dots, M\}$ and next pick out other $M-J$ different elements from the remains. We depict this combination as a set $\{(i_1), (i_2, \dots, i_{M-J+1})\}$, where each i_j is unique and belongs to $\{1, \dots, M\}$. In this way, we can form $\binom{M}{1}\binom{M-1}{M-J}$ different sets, denoted as Γ_ℓ for $\ell = 1, \dots, \binom{M}{1}\binom{M-1}{M-J}$. For $\Gamma_\ell = \{(i_1), (i_2, \dots, i_{M-J+1})\}$, we write $\overline{\Gamma}_\ell = i_1$ and $\underline{\Gamma}_\ell = \{i_2, \dots, i_{M-J+1}\}$. The weighted average for Γ_ℓ is represented as

$$\tau_\ell(x) := \left(1 - \frac{M-J}{\alpha M}\right) r_{\overline{\Gamma}_\ell}(x) + \frac{1}{\alpha M} \sum_{j \in \underline{\Gamma}_\ell} r_j(x).$$

Then, as the weighted average of the extreme losses, $\widetilde{\vartheta}(x)$ equals to the maximum of $\tau_\ell(x)$ for all $\ell \in \left\{1, \dots, \binom{M}{1}\binom{M-1}{M-J}\right\}$. In other words, problem (8) is equivalent to

$$\widetilde{\vartheta}(x) = \max_{\ell \in \left\{1, \dots, \binom{M}{1}\binom{M-1}{M-J}\right\}} \tau_\ell(x). \tag{9}$$

This observation motivates us to develop the cut generation algorithm to solve BP-APUB-M.

In Algorithm 1, we first solve a sequence of master programs denoted as (5), where constraint (5b) represents optimality cuts generated throughout the algorithm. In the k-th iteration, the master program (5) finds an optimal solution, $(\hat{x}_k, \hat{\eta}_k)$. To calculate $\vartheta(\hat{x}_k)$, Algorithm 1 presents a sorting method to solve problem (9). This method sorts the sequence $r_1(\hat{x}_k), \dots, r_M(\hat{x}_k)$. Notably, in the sorted sequence $r_{(1)_k}(\hat{x}_k), \dots, r_{(M)_k}(\hat{x}_k)$, $r_{(\lceil(1-\alpha)M\rceil)_k}(\hat{x}_k)$ serves as the cutoff point of Efron's percentile-based upper bound. Accordingly, $\widetilde{\vartheta}(\hat{x}_k) = w_{k+1}(\hat{x}_k)$, where the function w_{k+1} is given in (7). As a concequence, if $\hat{\eta}_k \geq w_{k+1}(\hat{x}_k)$, $\hat{x}_k$ is an optimal solution of BP-APUB-M. Otherwise, we generate an additional optimality cut, as $\eta \geq w_{k+1}(x)$, in the master program. It is worth mentioning that Algorithm 1 adds a mere few bootstrap samples into the cut set instead of all M samples, which effectively addresses issues when M is very large.

4.2 Practical Solution Method of a Two-Stage APUB Model

As an illustrative application, we now demonstrate the utility of APUB within the context of a specific class of optimization problems known as two-stage linear stochastic optimization with random recourse. In this case, APUB-M is adapted to formulate the first stage as

$$\min_x \quad c^\intercal x + \mathbb{U}^{\alpha}_{\text{APUB}}\left[\mathbb{E}[Q(x,\xi)] \,\middle|\, \widehat{\mathbb{P}}_N\right] \tag{10a}$$

$$\text{s.t.} \quad Ax = b, \tag{10b}$$

$$x \geq 0. \tag{10c}$$

By letting $\xi = (q, h, T, W)$, the second stage is represented as

$$Q(x,\xi) := \min_y \quad q^\intercal y \tag{11a}$$

$$\text{s.t.} \quad Wy = h - Tx, \tag{11b}$$

$$y \geq 0. \tag{11c}$$

Let $\xi_n = (q_n, h_n, T_n, W_n)$ be the nth point of the original random sample associated with the empirical distribution $\widehat{\mathbb{P}}_N$. The bootstrap sampling approximation of the two-stage APUB-M (10)-(11) can be written as

$$\min_{x,t} \quad c^\intercal x + t + \frac{1}{\alpha M}\sum_{m=1}^{M}\left[\frac{1}{N}\sum_{n=1}^{N} V_{m,n} Q(x,\xi_n) - t\right]_+ \tag{12a}$$

$$\text{s.t.} \quad Ax = b, \tag{12b}$$

$$x \geq 0, \tag{12c}$$

which can be further reformulated as the following linear program,

$$\min_{x,y,s,t} \quad c^\intercal x + t + \frac{1}{\alpha M}\sum_{m=1}^{M} s_m \tag{13a}$$

$$\text{s.t.} \quad Ax = b, \tag{13b}$$

$$s_m + t \geq \frac{1}{N}\sum_{n=1}^{N} V_{m,n} q_n^\intercal y_n, \qquad m = 1,\ldots,M, \tag{13c}$$

$$W_n y_n = h_n - T_n x, \qquad n = 1,\ldots,N, \tag{13d}$$

$$x \geq 0,\ y \geq 0,\ s \geq 0. \tag{13e}$$

We now propose adapting the well-established L-shaped method (Van Slyke and Wets, 1969; Birge and Louveaux, 2011) to achieve an efficient solution derivation described in Algorithm 2. Firstly, we solve a sequence of master programs denoted as in (14), where constraints (14c) and (14d) describe feasibility and optimality cuts, respectively, throughout the algorithm. The master program (14) finds an optimal solution, $(\hat{x}, \hat{\eta})$. We represent the objective value of the bootstrap sampling approximation problem (12) at $\hat{x}$ as

$$\tilde{\tilde{\vartheta}}(\hat{x}) := \min_{t\in\mathbb{R}} \quad t + \frac{1}{\alpha M}\sum_{m=1}^{M}\left[\frac{1}{N}\sum_{n=1}^{N} V_{m,n} Q(\hat{x},\xi_n) - t\right]_+ . \tag{22}$$

Step 3 in the algorithm considers the feasibility of the second-stage problem (11) at $\hat{x}$ and some ξ_n. If constraint (11b) is unsatisfied with $\hat{x}$ and ξ_n, problem (11) is infeasible and consequently $\tilde{\tilde{\vartheta}}(\hat{x}) = Q(\hat{x}, \xi_n) = \infty$. This necessitates the generation of one or more feasibility cuts of the type (14c). In the algorithm, the feasibility is checked solving the linear program (15). At step 4, we generate the optimality cut of type (14d) when $\hat{\eta} < \tilde{\tilde{\vartheta}}(\hat{x}) < \infty$. Analogous to Algorithm 1, we also employ the sorting method to calculate problem (22) and thus obtain $\tilde{\tilde{\vartheta}}(\hat{x}) = w(\hat{x})$, where the function w is described in (19). Notably, the sorting method used in Algorithm 2 facilitates the rapid computation of problem (22), which involves solving the second-stage optimization problems, effectively addressing performance issues when M is very large.

Algorithm 2 L-Shaped Algorithm for the Two-Stage APUB-M (10)-(11).

Step 1: Initialization

Let $\ell = k = 0$.

Step 2: Solve Master Problem.

Solve the linear program as follows:

$$\min_{x,\eta} \quad c^\intercal x + \eta \tag{14a}$$
$$\text{s.t.} \quad Ax = b, \tag{14b}$$
$$D_j x \geq d_j, \qquad j = 0, \dots, \ell, \tag{14c}$$
$$E_j x + \eta \geq e_j, \qquad j = 0, \dots, k, \tag{14d}$$
$$x \geq 0. \tag{14e}$$

Let $(\hat{x}, \hat{\eta})$ be an optimal solution. If no constraint (14d) is present, let $\hat{\eta} = -\infty$ and be not considered in the computation of $\hat{x}$.

Step 3: Generate feasibility Cuts.

For $n = 1, \dots, N$, solve the linear program

$$\begin{aligned} u(\hat{x}) := \min_{y, v^+, v^-} \quad & \mathbf{1}^\intercal v^+ + \mathbf{1}^\intercal v^- \\ \text{s.t.} \quad & W_n y + I v^+ - I v^- = h_n - T_n \hat{x}, \\ & y \geq 0, \ v^+ \geq 0, \ v^- \geq 0, \end{aligned} \tag{15}$$

where $\mathbf{1}$ is the all-ones vector and I is the identity matrix, until for some n, the optimal value $u(\hat{x}) > 0$. In this case, let ϕ be the associated simplex multipliers and define

$$D_{\ell+1} := \phi^\intercal T_n, \quad d_{\ell+1} := \phi^\intercal h_n, \tag{16}$$

to generate a feasibility cut of type (14c). Add $D_{\ell+1}$, $d_{\ell+1}$ to the constraint set (14c), set $\ell = \ell + 1$, and return to Step 2. If for all n, $u(\hat{x}) = 0$, go to Step 4.

Step 4: Generate Optimality Cuts.

For $n = 1, \dots, N$, solve the linear program

$$\begin{aligned} Q(\hat{x}, \xi_n) = \min_{y} \quad & q_n^\intercal y \\ \text{s.t.} \quad & W_n y = h_n - T_n \hat{x}, \\ & y \geq 0. \end{aligned} \tag{17}$$

Let ψ_n be the simplex multipliers associated with the optimal solution of Problem n of type (17).

Sort the following sequence

$$r_m(\hat{x}) := \frac{1}{N} \sum_{n=1}^{N} V_{m,n} Q(\hat{x}, \xi_n), \quad m = 1, \dots, M, \tag{18}$$

in ascending order and label the sorted sequence as $r_{(1)}(\hat{x}) \leq \cdots \leq r_{(M)}(\hat{x})$. Here, $(\cdot) : \{1, \dots, M\} \to \{1, \dots, M\}$ is a one-to-one mapping related to $\hat{x}$. Let $J := \lceil (1 - \alpha) M \rceil$ and

$$w(x) := \left(1 - \frac{M - J}{\alpha M}\right) r_{(J)}(x) + \frac{1}{\alpha M} \sum_{m=J+1}^{M} r_{(m)}(x). \tag{19}$$

Algorithm 2 L-Shaped Algorithm for the Two-Stage APUB-M (10)-(11) (continued)

If $\hat{\eta} \geq w(\hat{x})$, stop; $\hat{x}$ is an optimal solution. Otherwise, define

$$E_{k+1} := \left(1 - \frac{M-J}{\alpha M}\right)\left(\frac{1}{N}\sum_{n=1}^{N} V_{(J),n}\psi_n^{\intercal}T_n\right) + \frac{1}{\alpha M}\sum_{m=J+1}^{M}\left(\frac{1}{N}\sum_{n=1}^{N} V_{(m),n}\psi_n^{\intercal}T_n\right), \tag{20}$$

$$e_{k+1} := \left(1 - \frac{M-J}{\alpha M}\right)\left(\frac{1}{N}\sum_{n=1}^{N} V_{(J),n}\psi_n^{\intercal}h_n\right) + \frac{1}{\alpha M}\sum_{m=J+1}^{M}\left(\frac{1}{N}\sum_{n=1}^{N} V_{(m),n}\psi_n^{\intercal}h_n\right), \tag{21}$$

add $E_{k+1},\ e_{k+1}$ to the constraint set (14d), set $k = k+1$, and return to Step 2.

5 Numerical Analyses

We assess the efficacy of APUB-M through an extensive examination of classic problems in stochastic optimization, spanning both single-stage and two-stage scenarios. Section 5.1 provides a comparative analysis between APUB-M and traditional DRO utilizing Wasserstein distance. This comparison involves evaluating their respective out-of-sample performances and coverage probabilities in addressing a two-stage benchmark problem with fixed recourse (Dantzig, 2016). The comparative analysis reveals that, although APUB-M does not guarantee 100% coverage probability in situations characterized by a severe lack of data, it demonstrates a potential advantage by mitigating over-conservatism and achieving better average out-of-sample performance than the DRO approach. Furthermore, Section 5.2 extends the application scope of APUB-M to encompass problems featuring random recourse. It shows that the robustness and favorable performance of APUB-M are maintained. Traditional DRO methodologies encounter inherent limitations when confronted with the computational complexity resulting from random recourse. In Section 5.3, we subject APUB-M to rigorous testing using a multi-product newsvendor problem (Hanasusanto et al., 2015). APUB-M provides stable and high-quality solutions even when the sample size is small. Notably, Mohajerin Esfahani and Kuhn (2018) highlight that the Wasserstein distance based DRO fails to perform effectively in situations, like the newsvendor problem, where the random loss function exhibits a Lipschitz modulus concerning random scenarios, independent of decision variables. Across all investigated scenarios, the close correspondence between nominal levels and actual coverage probabilities serves as a testament to the methodological reliability of APUB-M.

5.1 A Two-Stage Product Mix Problem with Fixed Recourse

We adapt the benchmark two-stage product mix problem presented by King (1988) to our test case, which seeks to optimize the product mix of a furniture shop amid uncertain labor conditions. During the 'here-and-now' stage, the company commits to a long-term contract, promising to deliver a set quantity of furniture in each time period. This quantity can be adjusted due to strong market demand. Labor hours, which are crucial to production, are uncertain and variable, partly because of factors such as the COVID-19 pandemic. The production involves four distinct products and two workstations each constrained by the availability of labor hours. Each product requires varying amounts of labor across these workstations and contributes specific profit margins upon sale. Labor availability dictates production time, with more hours leading to reduced production time, a phenomenon attributed to skill diversity and improved efficiency. At this stage, the company's objective is to determine the most profitable product mix that meets contractual requirements while contending with the unpredictability of labor availability.

In the subsequent 'wait-and-see' stage, the company must confront the actual labor hours available, which may deviate from earlier estimates. In instances where there is a shortfall in the labor hours anticipated by the production plan conceived in the 'here-and-now' stage, the option exists to outsource additional labor hours for workstation. However, this supplemental workforce is not as efficient as the in-house labor. Thus, at this juncture, the firm's focus pivots to minimizing the expenses linked to acquiring these supplemental, less efficient labor hours, while still fulfilling the contractual furniture delivery commitments. The decision-making process in this stage is heavily dependent on actual labor availability and is geared toward economical adjustments to labor shortages.

In practice, the company determines the production quantities outlined in the contract by analyzing historical labor hours. However, the unpredictability of absenteeism, exacerbated by the COVID-19 pandemic, has led to a significant lack of reliable data. In response to this uncertainty, the company seeks to define a profit threshold that the expected profit from this contract is likely to meet or exceed, maintaining a confidence level of approximately $100(1-\alpha)\%$. We describe the profit threshold in form of APUB, which ensures the statistical reliability of the company's objective. On this basis, we represent this product mix problem as the two-stage APUB-M (10)-(11). In the first stage, x is the decision vector for the product mix and the negative value of c represents the per-unit profits of products. In the second stage, the decision vector y signifies the outsourced labor hours assigned to workstations and associated with unit cost q (the last two components with a cost of zero correspond to two slack variables), h stands for the random labor hours available at workstations, T includes the production times required for products, and the negative value of W represents the efficiency rate of outsourced labor. The numerical parameters in our study are specified as follows:

$$\begin{aligned}
&A = 0,\ b = 0,\ c = [-12, -20, -18, -40]^\intercal, \\
&q = [6, 12, 0, 0]^\intercal,\ h = [500\gamma_1, 500\gamma_2]^\intercal, \\
&T = \begin{bmatrix} 4 - \frac{\gamma_1}{4} & 9 - \frac{\gamma_1}{4} & 7 - \frac{\gamma_1}{4} & 10 - \frac{\gamma_1}{4} \\ 3 - \frac{\gamma_2}{4} & 1 - \frac{\gamma_2}{4} & 3 - \frac{\gamma_2}{4} & 6 - \frac{\gamma_2}{4} \end{bmatrix},\ W = \begin{bmatrix} -0.9 & 0 & 1 & 0 \\ 0 & -0.9 & 0 & 1 \end{bmatrix},
\end{aligned}$$

where

$$[\gamma_1, \gamma_2]^\intercal \sim \frac{7}{10}\mathcal{N}\left(\begin{bmatrix} 12 \\ 8 \end{bmatrix}, \begin{bmatrix} 5.76 & 1.92 \\ 1.92 & 2.56 \end{bmatrix}\right) + \frac{3}{10}\mathcal{N}\left(\begin{bmatrix} 2 \\ 1 \end{bmatrix}, \begin{bmatrix} 0.16 & 0.04 \\ 0.04 & 0.04 \end{bmatrix}\right)$$

has a mixed 2-dimensional normal distribution. In the two-stage APUB-M (10)-(11), we generate a random sample with a size N from the mixed normal distribution.

5.1.1 Convergence of Bootstrap Sampling Approximation.

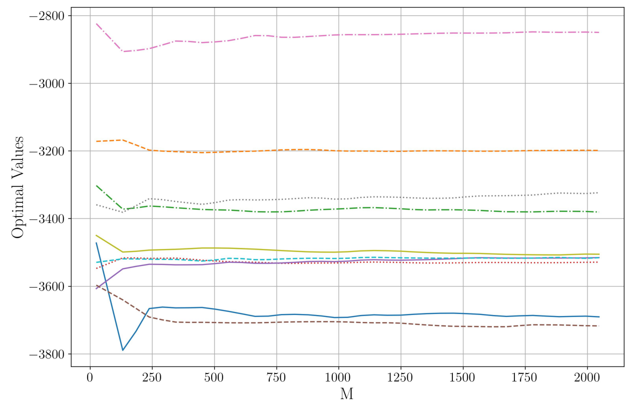

Figure 2: Convergence of the bootstrap sampling approximation.

We now assess the convergence of the bootstrap sampling method applied to our model. Our evaluation encompasses 10 independent simulations, each producing $N = 30$ sample data points and resolving the subsequent approximation as defined in (13). Throughout these tests, we maintain a consistent nominal level of $(1 - \alpha) = 0.8$. Figure 2 illustrates the relationship between the number M of bootstrap samples and the optimal values of our approximation problem, with M reaching up to 2000. A discernible stabilization trend is evident in the data: as M increases, variability in the optimal values conspicuously decreases. Notably, for $M \geq 1000$, the convergence of the approximation becomes evident as the fluctuation in the optimal values significantly lessens. This consistency bolsters our decision to adopt $M = 2000$ for all subsequent experiments in this section.

5.1.2 Comparative Analysis between APUB-M and DRO.

We carry out a comparative analysis between APUB-M and the Wasserstein distance based DRO approach when applied to this two-stage product mix problem. Recall that, used in the two-stage problem, APUB-M is formed as (10)-(11). Our aim is to evaluate their performance under various conditions and identify their respective strengths and limitations.

All the tests are conducted using a Monte Carlo simulation of 1000 replications. In each replication, we generate a training sample set of size N, using which APUB-M provides an optimal solution $\widehat{x}_N^\alpha$. The out-of-sample performance, $\mu(\widehat{x}_N^\alpha)$, is evaluated using a unique large-sized test sample set. The all replications provide the approximate probability distribution of $\mu(\widehat{x}_N^\alpha)$. For $N = 30$, 120, and 480, respectively, Figures 3a-3c display the curve of the mean of $\mu(\widehat{x}_N^\alpha)$ and the range from the 10th to the 90th percentile when varying $(1 - \alpha)$ from 0 to 1. Note that the leftmost case represents the out-of-sample performance of SAA-M, which is equivalently represented as APUB-M with $(1 - \alpha) = 0$. We also estimate the coverage probability of APUB-M concerning $\widehat{x}_N^\alpha$ as

$$\beta(\widehat{\vartheta}_N^\alpha, \{\widehat{x}_N^\alpha\}) = \Pr\left(\widehat{\vartheta}_N^\alpha \geq \mu(\widehat{x}_N^\alpha)\right).$$

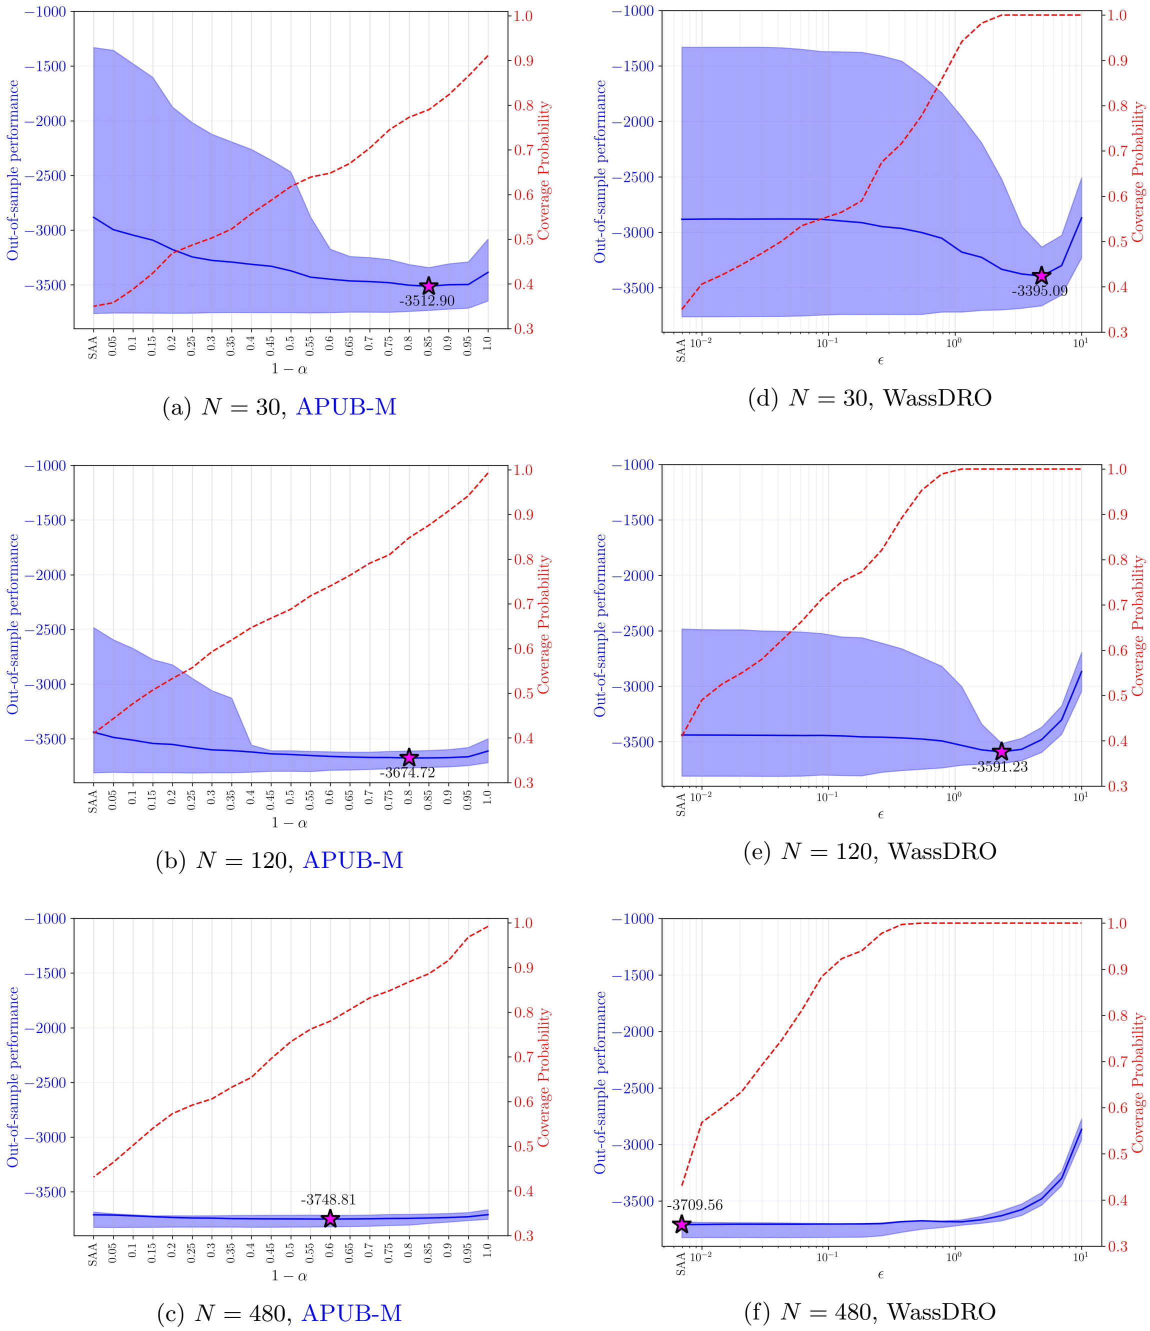

(a) $N = 30$, APUB-M

(d) $N = 30$, WassDRO

(b) $N = 120$, APUB-M

(e) $N = 120$, WassDRO

(c) $N = 480$, APUB-M

(f) $N = 480$, WassDRO

Figure 3: Out-of-sample performance (left axis, solid line, and shaded area) and the coverage probability (right axis, dashed line) as a function of the nominal level $(1 - \alpha)$ in APUB-M and a function of ϵ in WassDRO. The star symbol indicates the point where the mean of the out-of-sample performance attains its minimum. The minimum value of the mean is written next to the star symbol.

Recall that the function β is defined in (1). By Theorem 3.4, we know the asymptotic correctness of APUB-M concerning $\widehat{x}_N^\alpha$, i.e.,

$$\lim_{N\to\infty} \beta(\widehat{\vartheta}_N^\alpha, \{\widehat{x}_N^\alpha\}) \geq (1-\alpha).$$

Figures 3a-3c draw the curve of $\beta(\widehat{\vartheta}_N^\alpha, \{\widehat{x}_N^\alpha\})$ with respect to $(1-\alpha)$. Similarly, Figures 3d-3f report the out-of-sample performance and the coverage probability of the Wasserstein distance based DRO (labelled WassDRO in the figures), when altering the radius ϵ of the ball in 1-Wasserstein distance metric.

Comparing the out-of-sample performances and coverage probabilities of the two approaches in Figure 3, we have the following observations:

- Analysis of the minimum average costs (where negative values represent profits), as indicated by stars, and their associated 90th percentiles across varying sample size N, consistently showed that APUB-M achieves lower costs in comparison to WassDRO. Despite this, WassDRO demonstrates higher coverage probabilities when achieving its minimum average costs. This observation indicates a potential strategic compromise between minimizing the cost objective and improving the reliability of the solution.

- For small and medium sample sizes ($N = 30$ an $N = 120$), both APUB-M and WassDRO substantially outpace SAA-M. This advantage is evident from the notably lower and more focused cost distributions over a broad spectrum of nominal levels and radii. Additionally, both APUB-M and WassDRO reveal improvements in coverage probability, lending further support to their robustness and operational efficacy.

- For medium and large sample sizes ($N = 120$ an $N = 480$), we observe an accurate alignment of nominal level and actual coverage probabilities attests to the methodological soundness of APUB-M. This alignment not only acts as a validation of the fidelity of the method in asymptotic correctness but also highlights a significant methodological stride in the interpretability of intuitively chosen reliability levels.

- In the large sample scenarios ($N = 480$), while an increase in ϵ within the WassDRO framework improves the coverage probability, the associated cost is concurrently magnified, especially when ϵ is selected to be excessively large. This phenomenon suggests that an inappropriately chosen ϵ can lead to significant over-conservatism. In contrast, APUB-M sustains the out-of-sample performance at a level analogous to SAA-M, while exhibiting a consistent increment in coverage probability. Such findings validate the stability of APUB-M, affirming that the asymptotic consistency of APUB-M is maintained irrespective of the chosen nominal level.

- Lastly, it is noteworthy that, for both APUB-M and WassDRO, the critical $(1-\alpha)$ and ϵ points, corresponding to the minimum average costs, approach the leftmost point representing

SAA-M as N increases. This observation accentuates the benefits of SAA-M in a large sample scenario, where distributional ambiguity is reduced.

Overall, APUB-M not only exemplifies robustness but also avoids the over-conservatism often seen in traditional DRO approaches. Crucially, the nominal level $(1-\alpha)$ functions as a faithful reflection of its statistical meaning, aligning with the company's preference for a specific confidence level amidst distributional ambiguity. This functionality bolsters the statistical interpretability of our method, providing clear, relevant insights even before model training commences. However, it is important to recognize that, in contrast to DRO, APUB-M does not invariably guarantee a 100% coverage probability. This limitation is particularly evident in cases with too few data points, where even the worst-case scenarios in the sample may fail to encompass the most extreme eventualities, an issue that becomes pronounced in the face of severe sample scarcity.

5.2 A Two-Stage Product Mix Problem with Random Recourse

We now extend our analysis to a two-stage stochastic optimization problem incorporating random recourse. To facilitate direct comparison, we modify the test problem outlined in Section 5.1 to incorporate this element of randomness in the recourse measures. Accordingly, we define the random course as

$$W = \begin{bmatrix} -\mathcal{U}(0.6, 1.2) & 0 & 1 & 0 \\ 0 & -\mathcal{U}(0.6, 1.2) & 0 & 1 \end{bmatrix},$$

where $\mathcal{U}$ represents a uniform distribution. In this product mix optimization, the random recourse refers to the treatment of externally sourced labor hours as stochastic variables in the second stage, reflecting the real-world variability in the labor market.

We first observe the out-of-sample performance of SAA-M in both fixed and random recourse scenarios. The results, depicted in Figures 4a through 4b, reveal that the random recourse scenario exhibits a wider 10-90th percentile range and more than a 300-unit increase in the 90th percentile for both $N = 30$ and $N = 120$, as well as a higher mean, compared to the fixed recourse scenario (shown in Figures 3a through 3b). This indicates a heightened volatility in the random recourse case.

Nevertheless, APUB-M maintains a consistent performance profile in the random recourse scenario, as demonstrated in Figure 4, similar to its behavior in fixed recourse settings. This supports the model's methodological flexibility. Specifically, Figure 4a shows that APUB-M is exceptionally resilient when dealing with limited data, effectively reducing the mean of the cost and enhancing solution stability. This underscores the robustness of APUB-M when a suitable nominal level is chosen. Additionally, Figure 4b confirms the asymptotic correctness of APUB-M. In other words, APUB-M consistently meets the actual coverage probability when varying nominal levels, paralleling its fixed recourse performance. Moreover, Figure 4c suggests that with a large sample size (N = 480), APUB-M can avoid excessive conservatism irrespective of the nominal level.

On the other hand, in the random recourse scenario, using the same sample size, achieving

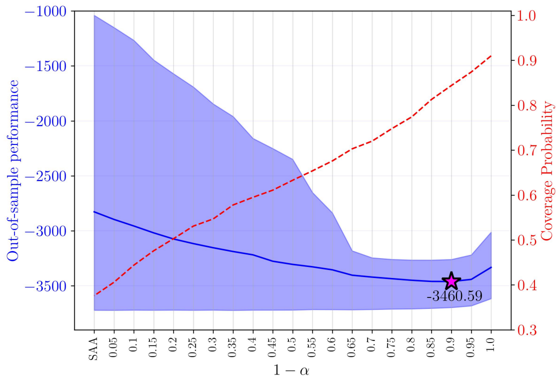

(a) $N = 30$

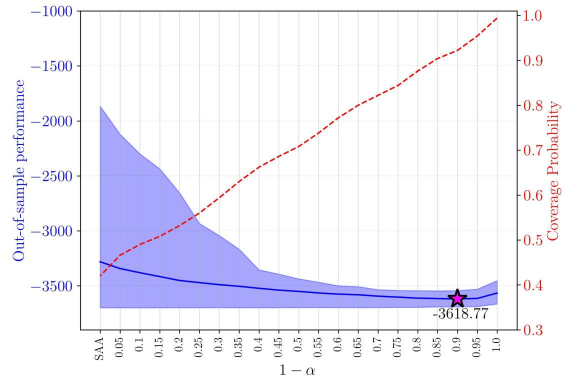

(b) $N = 120$

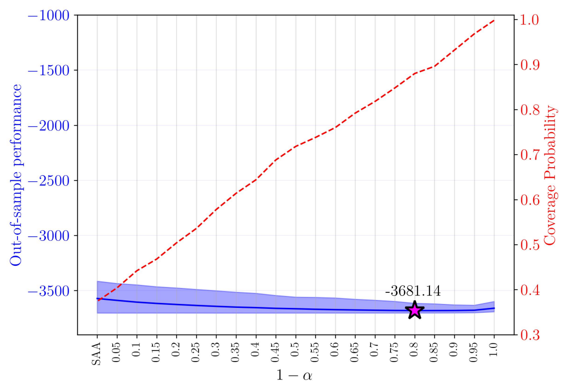

(c) $N = 480$

Figure 4: Out-of-sample performance (left axis, solid line, and shaded area) and the coverage probability (right axis, dashed line) as a function of the nominal level $(1 - \alpha)$ in APUB-M. The star symbol indicates the point where the mean of the out-of-sample performance attains its minimum. The minimum value of the mean is written next to the star symbol.

the minimum average cost (signified by a star in Figure 4) requires a higher nominal level to ensure greater robustness. Consequently, this minimum average cost is higher than that of the fixed recourse scenario. Additionally, the 10-90th percentile range is marginally broader across all nominal levels, compared to the fixed recourse case. These findings indicate a necessity for setting a larger nominal level in APUB-M to effectively manage the increased uncertainty introduced by random recourse.

5.3 A Multi-Product Newsvendor Problem

Consider a multi-product newsvendor problem, described in Hanasusanto et al. (2015), with the following random cost function,

$$F(x, \xi) = p^\intercal x + h^\intercal (x - \xi)_+ + b^\intercal (\xi - x)_+,$$

where x is the vector of order quantities for ten products, ξ represents random demand, p is the unit profit ($p < 0$ in the cost function), h and b are overage and underage costs.

5.3.1 Out-of-Sample Performance Analysis.

This test compares two cases: Case I assumes that ξ follows a mixed normal distribution as

$$\text{Case I:} \qquad \xi \sim \frac{1}{2}\mathcal{N}(\mu_1, \Sigma_1) + \frac{1}{2}\mathcal{N}(\mu_2, \Sigma_2),$$

and on this basis, Case II considers a biased noise added in the data generation as

$$\text{Case II:} \qquad \xi' = \xi + \varepsilon.$$

With the noise ε included, Case II has a large variation and, as a result, its distributional ambiguity is more serious. The numerical values of all parameters are provided in Appendix C.

Figures 5a to 5c depict the out-of-sample performances and coverage probabilities of APUB-M in Case I as the sample size N varies from 30 to 120, while Figures 5d to 5f illustrate Case II. In both cases, APUB-M outperforms SAA-M in terms of lower average cost and a narrower range from the 10th to the 90th percentile. Additionally, Figures 5 showcase the asymptotic correctness and consistency of APUB-M, as observed in Sections 5.1 and 5.2.

On the other hand, SAA-M exhibits more stable performance in Case I than in Case II. This suggests that Case I experiences less distributional ambiguity. We observe that APUB-M is adaptable to different levels of ambiguity. However, the model's performance is sensitive to the level of ambiguity. When $N = 30$, Case II demonstrates a much wider 10-90th percentile range and requires a larger nominal level $(1 - \alpha)$ to achieve the minimum average cost indicated by a star. Increasing N weakens the impact of noise on distributional ambiguity. APUB-M with a higher $(1 - \alpha)$ exhibits similar performance in both cases, particularly when $N = 120$. This observation underscores the capability of APUB-M to adjust to different uncertainty levels and the positive effect of increased sample sizes in lessening disparities caused by system noise.

5.3.2 Optimal Solutions Analysis.

We now compare the optimal solutions of Models SAA-M and APUB-M in Case I, as illustrated in Figures 6a through 6c, with N varying from 30 to 120. These solutions dictate the recommended order quantities for the ten products.

When $N = 30$, Figure 6a depicts significant fluctuations in the curves of the order quantities as $(1 - \alpha)$ increases, notably for product P2. The order quantity of product P2 decreases by 10.14% from SAA-M to APUB-M with $(1 - \alpha) = 0.5$. The increase in N noticeably stabilizes the order quantities. Upon reaching $N = 60$, the relative difference in the order quantities of product P2 reduces to 9.95% between SAA-M and APUB-M with $(1 - \alpha) = 0.5$. With a larger $N = 120$, all curves become flattened. In this scenario, as depicted in Figure 5c, APUB-M and SAA-M appear to achieve comparable performance. It can be seen in Figure 6c that their recommended optimal solutions are also very close.

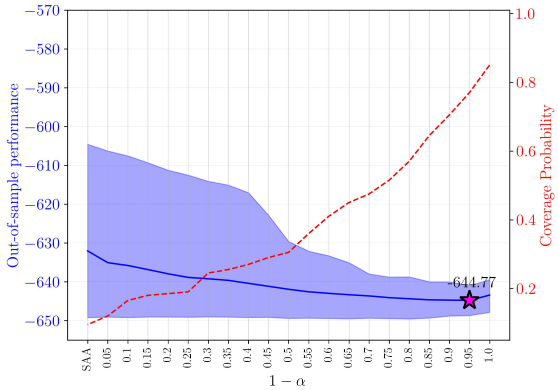

(a) Case I, $N = 30$

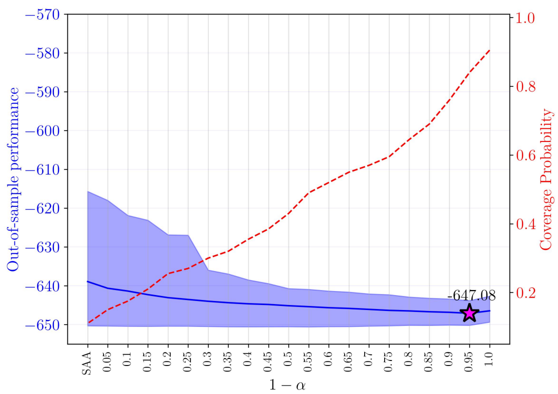

(b) Case I, $N = 60$

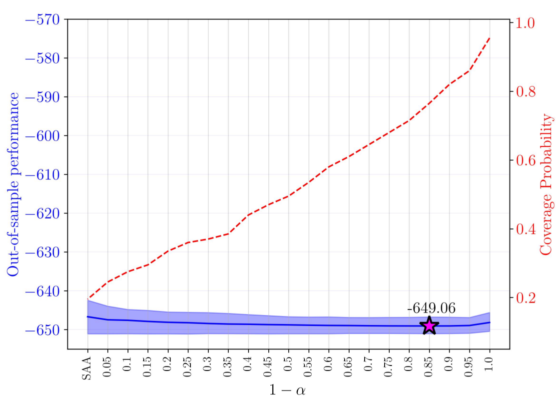

(c) Case I, $N = 120$

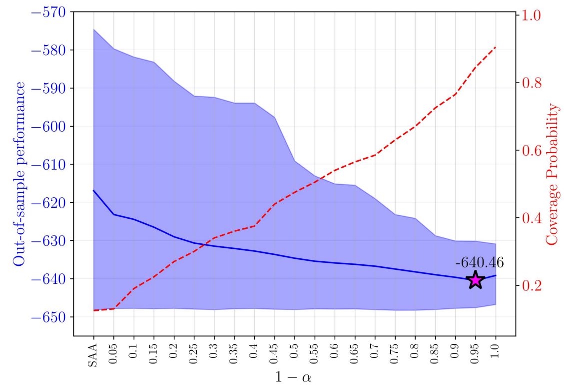

(d) Case II, $N = 30$

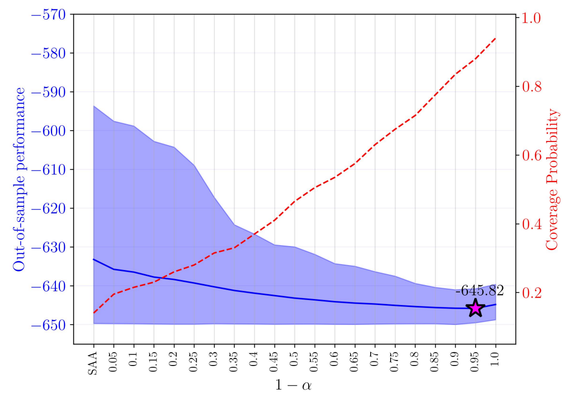

(e) Case II, $N = 60$

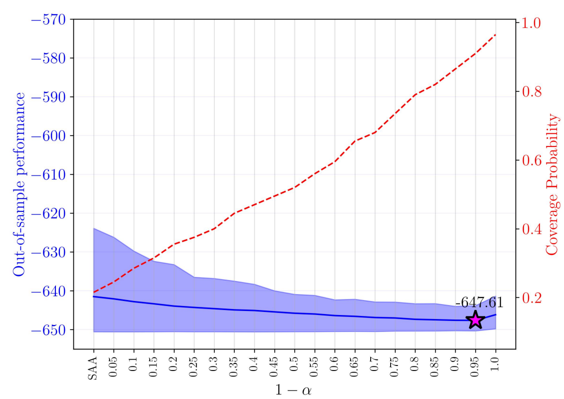

(f) Case II, $N = 120$

Figure 5: Out-of-sample performance (left axis, solid line, and shaded area) and the coverage probability (right axis, dashed line) as a function of the nominal level $(1 - \alpha)$ in APUB-M. The star symbol indicates the point where the mean of the out-of-sample performance attains its minimum. The minimum value of the mean is written next to the star symbol.

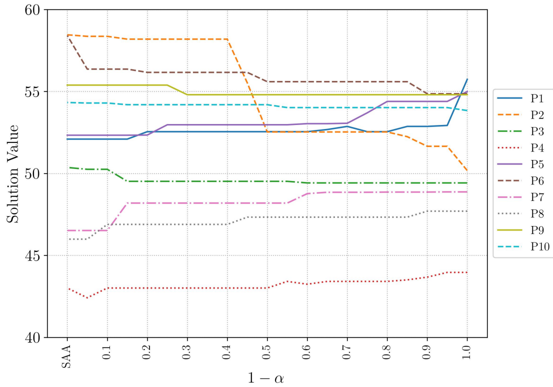

(a) $N = 30$

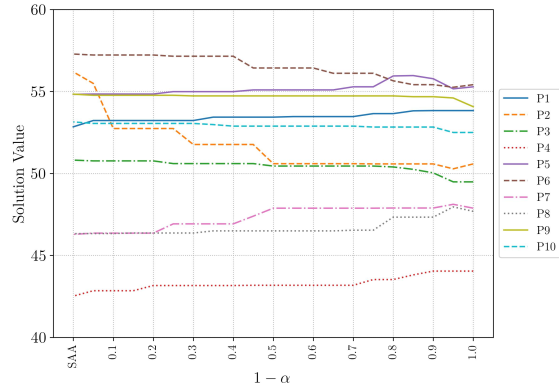

(b) $N = 60$

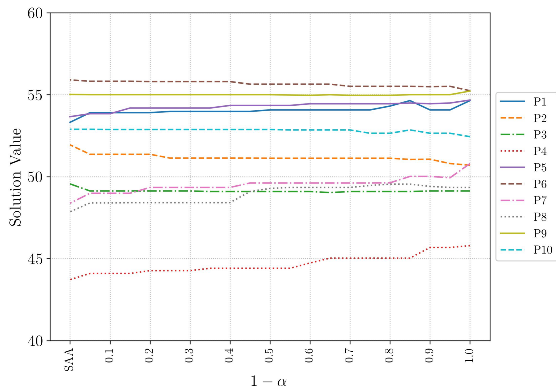

(c) $N = 120$

Figure 6: Optimal order quantities of the ten products.

SAA-M appears to be much more sensitive to N than APUB-M. Let us quantify the difference between the two solutions using the 2-norm. When N changes from 30 to 120, the difference is 7.89 for SAA-M, 3.99 for APUB-M with $(1 - \alpha) = 0.5$, and 3.36 for APUB-M with $(1 - \alpha) = 0.95$. This observation suggests that APUB-M can provide a high-quality optimal solution even with a small sample size. This capability underscores the ability of APUB-M to simulate scenarios typically requiring a larger volume of data.

6 Conclusions

In this work, we introduce APUB, a novel statistical upper bound that acts as a critical bridge between the realms of statistical inference and stochastic optimization. APUB enriches the theoretical landscape and highlights practical implications for the interpretability and application of stochastic optimization models. APUB serves as both an upper bound for the population mean, enhancing statistical analysis, and a coherent risk measure for the sample mean, focusing particularly on tail distribution errors due to insufficient sample sizes. By rigorously proving the statistical soundness of our approach, including its asymptotic correctness and consistency, we lay a solid foundation for integrating statistical methods into decision-making frameworks under distributional ambiguity.

Furthermore, APUB-M that we innovatively develop integrates APUB into stochastic optimization. This integration makes the reliability of APUB-M transparent, by ensuring that the attribute of asymptotic correctness inherent in APUB is seamlessly transferred to APUB-M in the optimization context. Indeed, the coverage probability of APUB-M aligns with asymptotic correctness at this predefined nominal level. Also, we meticulously show the asymptotic consistency of APUB-M, ensuring that our approach remains the nature of data-driven statistical methods, thereby promising stability and unbiasedness of APUB-M, as the sample size increases, but avoiding over-conservatism.

We employ a bootstrap sampling approximation method, BP-APUB-M, to manage the computational complexity, demonstrating that a significantly smaller number of bootstrap samples effectively maintains model integrity and reliability. Furthermore, we develop a cut generation algorithm for solving large-scaled cases and an L-shaped method for two-stage linear stochastic optimization with random recourse. These solution approaches, underpinning the practical viability of APUB-M, confirm their applicability in real-world scenarios, particularly in two-stage problems. Moreover, our empirical studies across various stochastic optimization problems, including single-stage and two-stage models, underscore the robustness and practicality of APUB-M. The comparative analysis with traditional DRO methods, particularly in settings of fixed and random recourse, highlights the enhanced interpretability and reduced conservatism of APUB-M. These results not only validate our theoretical findings but also showcase the broad applicability and effectiveness of our approach in real-world scenarios.

We demonstrate that APUB is asymptotically correct to the first order. In the context of optimization, APUB-M also shows asymptotic correctness. Future research will investigate the order of asymptotic correctness for APUB-M, specifically examining the rate at which its coverage probability converges to the predefined nominal level as the sample size increases.

Acknowledgments

The authors wish to extend their sincere gratitude to Prof. Jim Luedtke for his invaluable contributions to this work. His innovative insights have significantly contributed to the efficient reformulation of the bootstrap sampling approximation model. We immensely appreciate his willingness to share his expert knowledge, which has greatly enhanced the quality of our research.

References

Y. Abbasi-yadkori, D. Pál, and C. Szepesvári. Improved Algorithms for Linear Stochastic Bandits. In *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL https://proceedings.neurips.cc/paper_files/paper/2011/hash/e1d5be1c7f2f456670de3d53c7b54f4a-Abstract.html.

P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999. ISSN 1467-9965. doi: 10.1111/1467-9965.00068. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9965.00068. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-9965.00068.

K. B. Athreya. Strong law for the bootstrap. *Statistics & Probability Letters*, 1(3):147–150, Mar. 1983. ISSN 0167-7152. doi: 10.1016/0167-7152(83)90063-9. URL https://www.sciencedirect.com/science/article/pii/0167715283900639.

P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2):235–256, May 2002. ISSN 1573-0565. doi: 10.1023/A:1013689704352. URL https://doi.org/10.1023/A:1013689704352.

G. Bayraksan and D. K. Love. Data-Driven Stochastic Programming Using Phi-Divergences. In *The Operations Research Revolution*, INFORMS TutORials in Operations Research, pages 1–19. INFORMS, Sept. 2015. ISBN 978-0-9843378-8-0. doi: 10.1287/educ.2015.0134. URL https://pubsonline.informs.org/doi/abs/10.1287/educ.2015.0134. Section: 1.

A. Ben-Tal, D. den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen. Robust Solutions of Optimization Problems Affected by Uncertain Probabilities. *Management Science*, 59(2):341–357, Feb. 2013. ISSN 0025-1909. doi: 10.1287/mnsc.1120.1641. URL https://pubsonline.informs.org/doi/abs/10.1287/mnsc.1120.1641. Publisher: INFORMS.

A. C. Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the american mathematical society*, 49(1):122–136, 1941.

J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer Science & Business Media, June 2011. ISBN 978-1-4614-0237-4. Google-Books-ID: Vp0Bp8kjPxUC.

J. Blanchet and K. Murthy. Quantifying Distributional Model Risk via Optimal Transport. *Mathematics of Operations Research*, 44(2):565–600, May 2019. ISSN 0364-765X. doi: 10.1287/moor.2018.0936. URL https://pubsonline.informs.org/doi/abs/10.1287/moor.2018.0936. Publisher: INFORMS.

G. C. Calafiore and L. E. Ghaoui. On Distributionally Robust Chance-Constrained Linear Programs. *Journal of Optimization Theory and Applications*, 130(1):1–22, July 2006. ISSN 1573-2878. doi: 10.1007/s10957-006-9084-x. URL https://doi.org/10.1007/s10957-006-9084-x.

G. B. Dantzig. Linear Programming and Extensions. In *Linear Programming and Extensions*. Princeton University Press, Aug. 2016. ISBN 978-1-4008-8417-9. doi: 10.1515/9781400884179. URL https://www.degruyter.com/document/doi/10.1515/9781400884179/html.

E. Delage and Y. Ye. Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems. *Operations Research*, 58(3):595–612, June 2010. ISSN 0030-364X. doi: 10.1287/opre.1090.0741. URL https://pubsonline.informs.org/doi/abs/10.1287/opre.1090.0741. Publisher: INFORMS.

J. Devore. Probability and statistics for engineering and the sciences, Mar. 2009. URL https://hero.epa.gov/hero/index.cfm/reference/details/reference_id/196740.

D. Duque, S. Mehrotra, and D. P. Morton. Distributionally robust two-stage stochastic programming. *SIAM Journal on Optimization*, 32(3):1499–1522, 2022.

B. Efron. Nonparametric standard errors and confidence intervals. *Canadian Journal of Statistics*, 9(2):139–158, 1981. ISSN 1708-945X. doi: 10.2307/3314608. URL https://onlinelibrary.wiley.com/doi/abs/10.2307/3314608. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.2307/3314608.

B. Efron. Better Bootstrap Confidence Intervals. *Journal of the American Statistical Association*, 82(397):171–185, 1987. ISSN 0162-1459. doi: 10.2307/2289144. URL https://www.jstor.org/stable/2289144. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

R. Gao and A. Kleywegt. Distributionally Robust Stochastic Optimization with Wasserstein Distance. *Mathematics of Operations Research*, 48(2):603–655, May 2023. ISSN 0364-765X. doi: 10.1287/moor.2022.1275. URL https://pubsonline.informs.org/doi/full/10.1287/moor.2022.1275. Publisher: INFORMS.

P. Hall. On the Bootstrap and Confidence Intervals. *The Annals of Statistics*, 14(4):1431–1452, 1986. ISSN 0090-5364. URL https://www.jstor.org/stable/2241480. Publisher: Institute of Mathematical Statistics.

G. A. Hanasusanto, D. Kuhn, S. W. Wallace, and S. Zymler. Distributionally robust multi-item newsvendor problems with multimodal demand distributions. *Mathematical Programming*, 152 (1):1–32, Aug. 2015. ISSN 1436-4646. doi: 10.1007/s10107-014-0776-y. URL https://doi.org/10.1007/s10107-014-0776-y.

A. Hazra. Using the confidence interval confidently. *Journal of Thoracic Disease*, 9(10):4125–4130, Oct. 2017. ISSN 2072-1439. doi: 10.21037/jtd.2017.09.14. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5723800/.

A. J. King. Stochastic programming problems: Examples from the literature. *Numerical techniques for stochastic optimization*, pages 543–567, 1988.

F. Lin, X. Fang, and Z. Gao. Distributionally robust optimization: A review on theory and applications. *Numerical Algebra, Control and Optimization*, 12(1):159–212, 2022.

C. L. Mallows and D. Richter. Inequalities of Chebyshev Type Involving Conditional Expectations. *The Annals of Mathematical Statistics*, 40(6):1922–1932, 1969. ISSN 0003-4851. URL https://www.jstor.org/stable/2239511. Publisher: Institute of Mathematical Statistics.

V. Mnih, C. Szepesvári, and J.-Y. Audibert. Empirical Bernstein stopping. In *Proceedings of the 25th international conference on Machine learning*, ICML '08, pages 672–679, New York, NY, USA, July 2008. Association for Computing Machinery. ISBN 978-1-60558-205-4. doi: 10.1145/1390156.1390241. URL https://dl.acm.org/doi/10.1145/1390156.1390241.

P. Mohajerin Esfahani and D. Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, Sept. 2018. ISSN 1436-4646. doi: 10.1007/s10107-017-1172-1. URL https://doi.org/10.1007/s10107-017-1172-1.

H. Rahimian and S. Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.

T. R. C. Read and N. Cressie. Goodness-Of-Fit Statistics for Discrete Multivariate Data | Semantic Scholar, 2012. URL https://www.semanticscholar.org/paper/Goodness-Of-Fit-Statistics-for-Discrete-Data-Read-Cressie/563e27a8b8c4b1e412f306be93cea9467789b560.

R. T. Rockafellar. Convex Analysis: (PMS-28). In *Convex Analysis*. Princeton University Press, Apr. 2015. ISBN 978-1-4008-7317-3. doi: 10.1515/9781400873173. URL https://www.degruyter.com/document/doi/10.1515/9781400873173/html?lang=en.

R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41, 2000. ISSN 14651211. doi: 10.21314/JOR.2000.038. URL http://www.risk.net/journal-of-risk/technical-paper/2161159/optimization-conditional-value-risk.

S. Sarykalin, G. Serraino, and S. Uryasev. Value-at-Risk vs. Conditional Value-at-Risk in Risk Management and Optimization. In *State-of-the-Art Decision-Making Tools in the Information-Intensive Age*, INFORMS TutORials in Operations Research, pages 270–294. INFORMS, Sept. 2008. ISBN 978-1-877640-23-0. doi: 10.1287/educ.1080.0052. URL https://pubsonline.informs.org/doi/abs/10.1287/educ.1080.0052. Section: 13.

J. Shao and D. Tu. *The Jackknife and Bootstrap*. Springer Science & Business Media, Dec. 2012. ISBN 978-1-4612-0795-5. Google-Books-ID: VO3SBwAAQBAJ.

A. Shapiro, D. Dentcheva, and A. Ruszczynski. *Lectures on Stochastic Programming: Modeling and Theory, Third Edition*. MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, July 2021. ISBN 978-1-61197-658-8. doi: 10.1137/1.9781611976595. URL https://epubs.siam.org/doi/book/10.1137/1.9781611976595.

J. E. Smith and R. L. Winkler. The Optimizer's Curse: Skepticism and Postdecision Surprise in Decision Analysis. *Management Science*, 52(3):311–322, 2006. ISSN 0025-1909. URL https://www.jstor.org/stable/20110511. Publisher: INFORMS.

A. W. v. d. Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1998. ISBN 978-0-521-78450-4. doi: 10.1017/CBO9780511802256. URL https://www.cambridge.org/core/books/asymptotic-statistics/A3C7DAD3F7E66A1FA60E9C8FE132EE1D.

R. M. Van Slyke and R. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17(4):638–663, 1969. ISSN 0036-1399. URL https://www.jstor.org/stable/2099310. Publisher: Society for Industrial and Applied Mathematics.

W. Wiesemann, D. Kuhn, and M. Sim. Distributionally Robust Convex Optimization. *Operations Research*, 62(6):1358–1376, Dec. 2014. ISSN 0030-364X. doi: 10.1287/opre.2014.1314. URL https://pubsonline.informs.org/doi/10.1287/opre.2014.1314. Publisher: INFORMS.

W. Xie. Tractable reformulations of two-stage distributionally robust linear programs over the type-∞ Wasserstein ball. *Operations Research Letters*, 48, June 2020. doi: 10.1016/j.orl.2020.06.003.

A Proofs of Theorems and Propositions

A.1 Proof of Proposition 2.1

For a given $\widehat{\mathbb{P}}_N$, we briefly denote by $\boldsymbol{\zeta} := (\zeta_1, \dots, \zeta_N)$ a $(d_\xi \times N)$-dimensional random matrix. Let $G(\boldsymbol{\zeta}) := \frac{1}{N}\sum_{n=1}^{N} F(\zeta_n)$. It means that $\Psi(t) := \Pr\left(G(\boldsymbol{\zeta}) \leq t \;\middle|\; \widehat{\mathbb{P}}_N\right)$ is the cumulative probability distribution (cdf) of $G(\boldsymbol{\zeta})$ and $\Psi_M(t) := \frac{1}{M}\sum_{m=1}^{M} \mathbb{1}\{G(\boldsymbol{\zeta}_m) \leq t\}$ is an empirical cdf associated with a random sample $G(\boldsymbol{\zeta}_1), \dots, G(\boldsymbol{\zeta}_M)$. Thus, $\Psi_M(t)$ uniformly converges to $\Psi(t)$ as M goes to ∞ w.p.1 with respect to the bootstrap sample. Note that the fixed $\widehat{\mathbb{P}}_N$ is deterministic. Therefore, any percentile of $\Psi_M(t)$ converges to the counterpart of $\Psi(t)$ w.p.1. It follows that U_M^α converges to $\mathbb{U}^\alpha_{\text{Efron}}[\mu|\widehat{\mathbb{P}}_N]$.

A.2 Proof of Proposition 2.6

Efron's percentile-based upper bound ($\mathbb{U}^\alpha_{\text{Efron}}[\mu|\widehat{\mathbb{P}}_N]$) is asymptotically accurate to the first-order, i.e.,

$$\mathbb{U}^\alpha_{\text{Efron}}[\mu|\widehat{\mathbb{P}}_N] = (1-\alpha) + O(N^{-1/2}).$$

Its proof refers to Section 4.2 in Shao and Tu (2012). By Definition 2.2, we know that $\mathbb{U}^\alpha_{\text{APUB}}[\mu|\widehat{\mathbb{P}}_N] \geq \mathbb{U}^\alpha_{\text{Efron}}[\mu|\widehat{\mathbb{P}}_N]$.

A.3 Proof of Theorem 2.7

To prove Theorem 2.7, we need the following lemma about the bootstrap law of large numbers.

Lemma A.1. *Let $(\zeta_1, \dots, \zeta_N) \sim \widehat{\mathbb{P}}_N$. Then, as $N \to \infty$,*

$$\frac{1}{N}\sum_{n=1}^{N} F(\zeta_n) \to \mu \quad \textit{w.p.1.}$$

Proof. According to Theorem 2 in Athreya (1983) (Theorem B.2), if $\liminf_{M,N\to\infty} MN^{-\phi} > 0$ for some $\phi > 0$, and $\mathbb{E}|F(\xi) - \mu|^\theta < \infty$ for some $\theta \geq 1$ such that $\theta\phi > 1$, we have that, as $M, N \to \infty$,

$$\frac{1}{M}\sum_{m=1}^{M} F(\zeta_m) \to 1 \quad \text{w.p.1},$$

where $(\zeta_1, \dots, \zeta_M) \sim \widehat{\mathbb{P}}_N$.

In our case, choose $\phi = 1$, $\theta = 2$, and $M = N$. This ensures that $\liminf_{M,N\to\infty} MN^{-\phi} = 1 > 0$. The condition $\mathbb{E}|Z(\xi) - \mu|^\theta < \infty$ is satisfied due to finite variance. This completes the proof.

□

Proof of Theorem 2.7 Let $(\bar{\xi}_1, \bar{\xi}_2, \dots)$ be a realization of the sample path and $\bar{\mathbb{P}}_N$ be the empirical distribution associated with the first N sample points. Denote by $(\zeta_1(\bar{\mathbb{P}}_N), \dots, \zeta_N(\bar{\mathbb{P}}_N)) \sim$

$\bar{\mathbb{P}}_N$ a random sample under $\bar{\mathbb{P}}_N$. Denote a collection of realizations as

$$\mathfrak{S} := \left\{ (\bar{\xi}_1, \bar{\xi}_2, \dots) \; : \; \begin{array}{l} \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} F(\bar{\xi}_n) = \mu, \\ \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} F\left(\zeta_n(\bar{\mathbb{P}}_N)\right) = \mu \quad \text{w.p.1 (for } \zeta) \end{array} \right\}. \tag{23}$$

It follows by the Strong Law of Large Number and Lemma A.1 that $\Pr\{(\xi_1, \xi_2, \dots) \in \mathfrak{S}\} = 1$.

We now fix $(\bar{\xi}_1, \bar{\xi}_2, \dots) \in \mathfrak{S}$ along with its corresponding $(\bar{\mathbb{P}}_1, \bar{\mathbb{P}}_2, \dots)$. Here, $(\bar{\xi}_1, \bar{\xi}_2, \dots)$ and $\bar{\mathbb{P}}_N$ are deterministic, while $(\zeta_1(\bar{\mathbb{P}}_N), \dots, \zeta_N(\bar{\mathbb{P}}_N))$ remains randomness. For clarity, denote

$$\widehat{\mu}_N(\bar{\mathbb{P}}_N) := \frac{1}{N} \sum_{n=1}^{N} F\left(\zeta_n(\bar{\mathbb{P}}_N)\right).$$

By Proposition 2.3, we have

$$\mathbb{U}^{\alpha}_{\text{APUB}}[\mu | \bar{\mathbb{P}}_N] = \text{CVaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right).$$

Then, to prove Theorem 2.7, it suffices to show

$$\lim_{N\to\infty} \text{CVaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) = \mu. \tag{24}$$

Sarykalin et al. (2008) provide an expression of CVaR as

$$\text{CVaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) = \eta_{\alpha} \text{VaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) + (1 - \eta_{\alpha}) \text{CVaR}^{+}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right), \tag{25}$$

where

$$\text{CVaR}^{+}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) = \mathbb{E}\left[\widehat{\mu}_N(\bar{\mathbb{P}}_N) \;\middle|\; \widehat{\mu}_N(\bar{\mathbb{P}}_N) > \text{VaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right)\right],$$

$$\eta_{\alpha} = \frac{\Pr\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N) \le \text{VaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) \;\middle|\; \bar{\mathbb{P}}_N\right) - \alpha}{1 - \alpha}.$$

Describe $\widehat{\mu}_N(\bar{\mathbb{P}}_N) > \text{VaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right)$ as event $\mathcal{A}_N$. By the definition of VaR, it is clear to see $\Pr(\mathcal{A}_N \mid \bar{\mathbb{P}}_N) \ge \alpha$ for all N. The expression (25) implies that

$$\text{VaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) \le \text{CVaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) \le \text{CVaR}^{+}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right).$$

According to Lemma 21.2 in (Vaart, 1998) (Theorem B.3), we understand that convergence in distribution implies the convergence of the quantile function. Thus, we have

$$\lim_{N\to\infty} \text{VaR}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) = \mu,$$

On the other hand, applying Theorem 2.1 in (Mallows and Richter, 1969) (Theorem B.4), we have

$$\begin{aligned} \left|\text{CVaR}^{+}_{\alpha}\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) - \mathbb{E}\left[\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right]\right| &= \left|\mathbb{E}\left[\widehat{\mu}_N(\bar{\mathbb{P}}_N) \;\middle|\; \mathcal{A}_N\right] - \mathbb{E}\left[\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right]\right| \\ &\le \widetilde{\sigma}_N \left(\frac{1 - \Pr(\mathcal{A}_N \mid \bar{\mathbb{P}}_N)}{\Pr(\mathcal{A}_N \mid \bar{\mathbb{P}}_N)}\right)^{1/2} \\ &\le \widetilde{\sigma}_N \left(\frac{1-\alpha}{\alpha}\right)^{1/2}, \end{aligned}$$

where $\widetilde{\sigma}_N$ represents the standard deviation of $\widehat{\mu}_N(\bar{\mathbb{P}}_N)$. Since $\widehat{\mu}_N(\bar{\mathbb{P}}_N)$ converges to μ w.p.1, we know that $\widetilde{\sigma}_N$ converges to 0. Therefore, we establish:

$$\begin{aligned}\lim_{N\to\infty} \mathrm{CVaR}_\alpha^+\left(\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right) &= \lim_{N\to\infty} \mathbb{E}\left[\widehat{\mu}_N(\bar{\mathbb{P}}_N)\right] \\ &\overset{(a)}{=} \lim_{N\to\infty} \mathbb{E}\big[F(\zeta_1(\bar{\mathbb{P}}_N))\big] \\ &= \lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} F(\bar{\xi}_n) \\ &\overset{(b)}{=} \mu,\end{aligned}$$

where (a) holds since $(\zeta_1(\bar{\mathbb{P}}_N), \dots, \zeta_N(\bar{\mathbb{P}}_N))$ are drawn i.i.d. from $\bar{\mathbb{P}}_N$, and (b) holds due to the definition of $\mathfrak{S}$ (see (23)). Hence, we obtain (24) and complete the proof.

A.4 Proof of Theorem 2.7

To prove Theorem 3.1, we first prove the following two lemmas: the first lemma show the continuity of $\mu(x)$ and $\mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]$; the second lemma shows the uniform consistency of $\mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]$ on $\mathcal{K}$.

Lemma A.2. *Suppose Assumption B holds. Then $\mu(x)$ is continuous on $\mathcal{N}$, and $\mathbb{U}^\alpha_{APUB}[\mu(x)|\widehat{\mathbb{P}}_N]$ is a continuous convex function on $\mathcal{N}$ w.p.1.*

Proof. Since $F(x,\xi)$ is convex on $\mathcal{N}$, $\mu(x)$ is also convex on $\mathcal{N}$. Hence, $\mu(x)$ is continuous. Under Assumption B, $F(\cdot,\xi)$ is continuous and convex on $\mathcal{N}$. It is easy to see that, by Proposition 2.3, $\mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]$ is continuous convex on $\mathcal{N}$. □

Lemma A.3. *Suppose Assumption A and B holds. Then, we have as $N \to \infty$,*

$$\sup_{x\in\mathcal{K}} \left|\mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N] - \mu(x)\right| \to 0, \text{ w.p.1.}$$

Proof. Note that the open convex set $\mathcal{N} \subseteq \mathbb{R}^{d_x}$. We first construct a countable dense subset of $\mathcal{N}$ as $\mathcal{D} := \mathbb{Q}^{d_x} \cap \mathcal{N}$, where $\mathbb{Q}^{d_x}$ represents the set of d_x-dimensional rational numbers. Choose a sample path $(\xi_1, \xi_2, \dots)$ and hence $\widehat{\mathbb{P}}_N$ is the empirical distribution associated to the first N sample points. For $x \in \mathcal{D}$, we denote an event as

$$\Upsilon_x := \left\{(\xi_1,\xi_2,\dots) \;:\; \lim_{N\to\infty} \mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N] = \mu(x)\right\}.$$

Since $\mu(x) < \infty$ and $\sigma(x) < \infty$ under Assumption **(B2)**, it follows by Theorem 2.7 that $\Pr\big((\xi_1,\xi_2,\dots) \in \Upsilon_x\big) = 1$, which implies that $\Pr\big((\xi_1,\xi_2,\dots) \in \bigcap_{x\in\mathcal{D}} \Upsilon_x\big) = 1$. In other words, $\mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]$ converges pointwisely to $\mu(x)$ on $\mathcal{D}$ w.p.1. Furthermore, by Proposition A.2 and Theorem 10.8 in Rockafellar (2015) (Theorem B.5), we can conclude that $\mathbb{U}^\alpha_{\mathrm{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]$ converges uniformly a certain continuous function ν on $\mathcal{K}$ w.p.1. Since $\nu(x)$ and $\mu(x)$ coincidence on a dense subset of $\mathcal{K}$ and they are both continuous on $\mathcal{K}$, we know that $\nu(x) = \mu(x)$ for all $x \in \mathcal{K}$. This completes the proof. □

Proof of Theorem 2.7

i) Proof of the consistency of $\widehat{\vartheta}_N^{\alpha}$. Choose $x^* \in \mathcal{S}$ and $\widehat{x}_N \in \widehat{\mathcal{S}}_N$. It is easy to see that

$$\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_N)|\widehat{\mathbb{P}}_N] \leq \mathbb{U}_{\text{APUB}}^{\alpha}[\mu(x^*)|\widehat{\mathbb{P}}_N]$$

and

$$\mu(x^*) \leq \mu(\widehat{x}_N).$$

Thus, we have

$$\begin{aligned}
|\widehat{\vartheta}_N^{\alpha} - \vartheta^*| &= \left|\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_N)|\widehat{\mathbb{P}}_N] - \mu(x^*)\right| \\
&= \max\left\{\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_N)|\widehat{\mathbb{P}}_N] - \mu(x^*),\ \mu(x^*) - \mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_N)|\widehat{\mathbb{P}}_N]\right\}, \\
&\leq \max\left\{\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(x^*)|\widehat{\mathbb{P}}_N] - \mu(x^*),\ \mu(\widehat{x}_N) - \mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_N)|\widehat{\mathbb{P}}_N]\right\} \\
&\leq \sup_{x\in\mathcal{K}}\left|\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(x)|\widehat{\mathbb{P}}_N] - \mu(x)\right|,
\end{aligned}$$

which converges to 0 w.p.1 by Theorem A.3. This completes the proof.

ii) Proof of the consistency of $\widehat{\mathcal{S}}_N^{\alpha}$. Let $\mathcal{O}$ as a collection of sample paths along which $\widehat{\mathcal{S}}_N^{\alpha} \subseteq \mathcal{K}$ for a sufficiently large N and $\widehat{\vartheta}_N^{\alpha} \to \vartheta^*$. By the above proof and Assumption (**A2**), we have $\Pr\big((\xi_1, \xi_2, \dots) \in \mathcal{O}\big) = 1$. We now choose $(\xi_1, \xi_2, \dots) \in \mathcal{O}$. Thus $\widehat{\mathcal{S}}_N^{\alpha}$ is the optimal solution set of APUB-M using the first N sample points.

Suppose by contradiction that $\mathbb{D}(\widehat{\mathcal{S}}_N^{\alpha}, \mathcal{S}) \not\to 0$ along the sample path $(\xi_1, \xi_2, \dots)$. Then, there exists $\varepsilon > 0$ such that for all $M \in \mathbb{N}$, there exists some $N > M$ for which $\mathbb{D}(\widehat{\mathcal{S}}_N^{\alpha}, \mathcal{S}) > \varepsilon$. Specifically, there exists $\widehat{x}_N \in \widehat{\mathcal{S}}_N^{\alpha}$ such that $\inf_{y\in\mathcal{S}} \|\widehat{x}_N, y\| > \varepsilon$. Because of the compactness of $\mathcal{K}$, we can find a subsequence $\widehat{x}_{N_k} \in \widehat{\mathcal{S}}_{N_k}^{\alpha}$ such that $\widehat{x}_{N_k} \subseteq \mathcal{K}$ for all $k \in \mathbb{N}$, and

$$\lim_{k\to\infty} \widehat{x}_{N_k} = \widehat{x} \in \mathcal{K}, \qquad \inf_{y\in\mathcal{S}} \|\widehat{x}_N, y\| > \varepsilon, \quad \text{for all } k.$$

It follows that $\widehat{x} \notin \mathcal{S}$ and hence $\mu(\widehat{x}) > \vartheta^*$. On the other hand, we have

$$\left|\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_{N_k})|\widehat{\mathbb{P}}_N] - \mu(\widehat{x})\right| \leq \left|\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_{N_k})|\widehat{\mathbb{P}}_N] - \mu(\widehat{x}_{N_k})\right| + \left|\mu(\widehat{x}_{N_k}) - \mu(\widehat{x})\right|.$$

On the right hand of the above inequality, the first term converges to zero by Theorem A.3, and the second term converges to zero because of the continuity of $\mu(x)$. Thus,

$$\lim_{k\to\infty} \mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_{N_k})|\widehat{\mathbb{P}}_N] = \mu(\widehat{x}).$$

The definition of $\mathcal{O}$ ensures that $\mathbb{U}_{\text{APUB}}^{\alpha}[\mu(\widehat{x}_{N_k})|\widehat{\mathbb{P}}_N] = \widehat{\vartheta}_N^{\alpha} \to \vartheta^*$. It implies that $\mu(\widehat{x}) = \vartheta^*$. This is contradictory to the assertion that $\mathbb{D}(\widehat{\mathcal{S}}_N^{\alpha}, \mathcal{S}) \not\to 0$.

A.5 Proof of Theorem 3.4

For the convenience of reading, we restate some notations here. Recall $(\xi_1, \dots, \xi_N)$ is an i.i.d. random sample generated from $(\Omega, \mathcal{F}, \mathbb{P})$. Consider a function $F : \mathcal{X} \times \Xi \to \mathbb{R}$, where Ξ is the support of ξ_1 and $\mathcal{X}$ is a decision region. We define the population mean $\mu(x) = \mathbb{E}[F(x, \xi)]$ and variance $\sigma^2(x) = \mathbb{E}[(F(x,\xi) - \mu(x))^2]$. Let $\widehat{\mathbb{P}}_N$ be the empirical distribution associated with the random sample, and $(\zeta_1, \ldots, \zeta_N)$ is a bootstrap sample generated from $\widehat{\mathbb{P}}_N$. Accordingly, denote $\widehat{\mu}_N(x) = \mathbb{E}[F(x,\zeta)|\widehat{\mathbb{P}}_N]$, $\widehat{\sigma}^2_N(x) = \mathbb{E}\left[(F(x,\zeta) - \widehat{\mu}_N(x))^2 \,\middle|\, \widehat{\mathbb{P}}_N\right]$, and the standardized sample mean

$$S_N(x) := \frac{\sqrt{N}\,(\widehat{\mu}_N(x) - \mu(x))}{\sigma(x)}.$$

Denote the bootstrap counterparts of $\widehat{\mu}_N(x)$ and $S_N(x)$ by

$$\widehat{\mu}^*_N(x) := \frac{1}{N}\sum_{n=1}^{N} F(x,\zeta_n) \quad \text{and} \quad S^*_N(x) := \frac{\sqrt{N}\,(\widehat{\mu}^*_N(x) - \widehat{\mu}_N(x))}{\widehat{\sigma}_N(x)}.$$

We now outline our proof as follows. We know by Proposition A.2 that, the objective function of APUB-M, $\mathbb{U}^{\alpha}_{\text{APUB}}[\mu(x)|\widehat{\mathbb{P}}_N]$, is continuous. Also, Assumption A states that the set $\widehat{\mathcal{S}}^{\alpha}_N$ of the optimal solutions of APUB-M is contained in the compact set $\mathcal{K}$ for a sufficiently large N w.p.1. This implies that $\widehat{\mathcal{S}}^{\alpha}_N$ is compact and then there exists $\tilde{x}_N \in \widehat{\mathcal{S}}^{\alpha}_N$ such that

$$\mu(\tilde{x}_N) = \max_{x \in \widehat{\mathcal{S}}^{\alpha}_N} \mu(x).$$

Subsequently, we have

$$\begin{aligned} \beta(\widehat{\vartheta}^{\alpha}_N, \widehat{\mathcal{S}}^{\alpha}_N) =& \Pr\left(\mathbb{U}^{\alpha}_{\text{APUB}}[\mu(\tilde{x}_N)|\widehat{\mathbb{P}}_N] \geq \mu(\tilde{x}_N)\right) \\ \geq& \Pr\left(\mathbb{U}^{\alpha}_{\text{Efron}}[\mu(\tilde{x}_N)|\widehat{\mathbb{P}}_N] \geq \mu(\tilde{x}_N)\right), \end{aligned}$$

which means that, to obtain

$$\lim_{N\to\infty} \beta(\widehat{\vartheta}^{\alpha}_N, \widehat{\mathcal{S}}^{\alpha}_N) \geq (1-\alpha),$$

it suffices to prove the following uniform convergence as

$$\lim_{N\to\infty} \sup_{x\in\mathcal{K}} \left|\Pr\left(\mathbb{U}^{\alpha}_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N] \geq \mu(x)\right) - (1-\alpha)\right| = 0. \tag{26}$$

In this section we describe the proof of equation (26) in five steps as follows:

(i) Section A.5.1 shows the first three moments of $F(x,\xi)$ are finite.

(ii) Section A.5.2 proves that as N increases to ∞, the cdf of $S_N(x)$ uniformly (in $x \in \mathcal{K}$) approximates to the standard normal distribution function.

(iii) Section A.5.3 proves a bootstrap version of the result of step (iii).

(iv) Section A.5.4 links $\mathbb{U}^{\alpha}_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N]$ to $S_N(x)$, and utilizes the results from the step (ii) and (iii) to prove the bootstrap percentile uniformly converges to normal percentile.

(v) Section A.5.5 rigorously proves (26).

A.5.1 Finite Moments and Domination Property.

The following lemma shows the first three moments of $F(x,\xi)$ are finite.

Lemma A.4. *Under the assumptions of Theorem 3.4, for $k \leq 3$, there exist K_k such that*

$$\sup_{x \in \mathcal{K}} \mathbb{E}\left[|F(x,\xi)|^k\right] \leq K_k.$$

Proof. We first prove the boundedness for $\mathbb{E}\left[|F(x,\xi)|^k\right]$ over $x \in \mathcal{K}$. The Lipschitz continuity of $F(x,\xi)$ leads to

$$|F(x,\xi)| \leq |F(x_0,\xi)| + |F(x,\xi) - F(x_0,\xi)| \leq |F(x_0,\xi)| + L(\xi)\|x - x_0\| \leq |F(x_0,\xi)| + D \cdot L(\xi),$$

where $D := \sup_{x \in \mathcal{K}} \|x - x_0\|$ is the diameter of $\mathcal{K}$. Raising both sides to the power k and applying the inequality $(a+b)^k \leq 2^{k-1}(a^k + b^k)$ for $a, b \geq 0$, we obtain

$$|F(x,\xi)|^k \leq 2^{k-1}\left(|F(x_0,\xi)|^k + D^k L(\xi)^k\right).$$

Taking expectations on both sides,

$$\mathbb{E}\left[|F(x,\xi)|^k\right] \leq 2^{k-1}\left(\mathbb{E}\left[|F(x_0,\xi)|^k\right] + D^k \mathbb{E}\left[L(\xi)^k\right]\right).$$

The assumptions of Theorem 3.4 require that $\mathbb{E}\left[|F(x_0,\xi)|^k\right] < \infty$ and $\mathbb{E}\left[L(\xi)^k\right] < \infty$ for $k \leq 3$. Let

$$K_k := 2^{k-1}\left(\mathbb{E}\left[|F(x_0,\xi)|^k\right] + D^k \mathbb{E}\left[L(\xi)^k\right]\right).$$

Then,

$$\sup_{x \in \mathcal{K}} \mathbb{E}\left[|F(x,\xi)|^k\right] \leq K_k.$$

□

Lemma A.5. *Under the assumptions of Theorem 3.4, there exists an integrable function $G(x,\xi)$ with $\mathbb{E}[G(x,\xi)] < \infty$ for each $x \in \mathcal{K}$ such that $|F(x,\xi)| \leq G(x,\xi)$.*

Proof. By Assumption (ii) of Theorem 3.4,

$$|F(x,\xi)| \leq |F(x,\xi) - F(x_0,\xi)| + |F(x_0,\xi)| < |F(x_0,\xi)| + L(\xi)\|x - x_0\|.$$

Define

$$G(x,\xi) := |F(x_0,\xi)| + L(\xi)\|x - x_0\|.$$

Then, we have

$$\mathbb{E}[G(x,\xi)] = \mathbb{E}|F(x_0,\xi)| + \mathbb{E}[L(\xi)]\|x - x_0\|$$

Assumption (i) and (iii) of Theorem 3.4 implies $\mathbb{E}|F(x_0,\xi)| < \infty$ and $\mathbb{E}[L(\xi)] < \infty$. Thus, $\mathbb{E}[G(x,\xi)] < \infty$ for all $x \in \mathcal{K}$. □

A.5.2 Uniform Berry-Esseen Inequality.

The following theorem states the Berry-Esseen bound for $S_N(x)$, which is fundamental in our proof.

Theorem A.6 (Berry-Esseen Inequality (Berry, 1941))**.** *For a fixed $x \in \mathcal{K}$, if $\mathbb{E}|F(x,\xi)|^3 < \infty$, then*

$$\sup_{z\in\mathbb{R}} \left| \Pr(S_N(x) \le z) - \Phi(z) \right| \le C \frac{\mathbb{E}|F(x,\xi)|^3}{\sigma^3(x)\sqrt{N}}$$

where C is a constant independent of x.

Lemma A.7 implies the pdf of $S_N(x)$ converges to the standard normal uniformly in $x \in \mathcal{K}$.

Lemma A.7. *Under the assumptions of Theorem 3.4, we have $\Pr(S_N(x) \le z) = \Phi(z) + o(1)$ uniformly in $x \in \mathcal{K}$ and $z \in \mathbb{R}$.*

Proof. It follows by the assumptions of Theorem 3.4 that there exist a constant $\sigma_{\min}$ such that $\sigma(x) \ge \sigma_{\min}$ for all $x \in \mathcal{K}$. In addition, by Lemma A.4, $\sup_{x\in\mathcal{K}} \mathbb{E}|F(x,\xi)|^3 \le K_3$. Then, by Theorem A.6, we have

$$\sup_{x\in\mathcal{K}} \sup_{z\in\mathbb{R}} \left| \Pr(S_N(x) \le z) - \Phi(z) \right| \le C \frac{\sup_{x\in\mathcal{K}} \mathbb{E}|F(x,\xi)|^3}{\inf_{x\in\mathcal{K}} \sigma^3(x)\sqrt{N}} \le C \frac{K_3}{\sigma_{min}^3 \sqrt{N}} = o(1).$$

□

A.5.3 Uniform Berry-Esseen Inequality for Percentile Bootstrap Sampling.

We now show the uniform asymptotic behavior of the moments of $\widehat{\mathbb{P}}_N$ in the following lemma.

Lemma A.8. *Under the assumptions of Theorem 3.4, the following two conditions hold as $N \to \infty$ w.p.1.,*

(**C1***) $\sup_{x\in\mathcal{K}} \left| \mathbb{E}\left[|F(x,\zeta)|^k \;\middle|\; \widehat{\mathbb{P}}_N \right] - \mathbb{E}\left[|F(x,\xi)|^k \right] \right| \to 0$ *for $k \le 3$.*

(**C2***) $\sup_{x\in\mathcal{X}} |\widehat{\sigma}_N(x) - \sigma(x)| \to 0.$

Proof. We now prove (**C1***). Assumption B states that $F(x,\xi)$ is continuous in $x \in \mathcal{K}$ for any ξ. In addition, Lemma A.5 implies $F(x,\xi)$ dominated by an integrable function for each $x \in \mathcal{K}$. According to Theorem 7.53 in Shapiro et al. (2021) (Theorem B.6), we know that, as $N \to \infty$,

$$\sup_{x\in\mathcal{K}} \left| \mathbb{E}\left[|F(x,\zeta)|^k \;\middle|\; \widehat{\mathbb{P}}_N \right] - \mathbb{E}\left[|F(x,\xi)|^k \right] \right| \to 0, \quad \text{w.p.1.} \tag{27}$$

We next prove (**C2***). Since

$$\begin{aligned} \widehat{\sigma}_N^2(x) &= \mathbb{E}\left[(F(x,\zeta) - \widehat{\mu}_N(x))^2 \;\middle|\; \widehat{\mathbb{P}}_N \right] \\ &= \mathbb{E}\left[((F(x,\zeta) - \mu(x)) - (\widehat{\mu}_N(x) - \mu(x)))^2 \;\middle|\; \widehat{\mathbb{P}}_N \right] \\ &= \mathbb{E}\left[(F(x,\zeta) - \mu(x))^2 \;\middle|\; \widehat{\mathbb{P}}_N \right] - (\widehat{\mu}_N(x) - \mu(x))^2, \end{aligned}$$

we obtain

$$\sup_{x\in\mathcal{K}} \left| \widehat{\sigma}_N^2(x) - \sigma^2(x) \right| \le \sup_{x\in\mathcal{K}} \left| \mathbb{E}\left[(F(x,\zeta) - \mu(x))^2 \;\middle|\; \widehat{\mathbb{P}}_N \right] - \sigma^2(x) \right| + \sup_{x\in\mathcal{K}} (\widehat{\mu}_N(x) - \mu(x))^2 .$$

Equation (27) has shown that the two terms at the right-hand-side of the above inequality converge to 0 as $N \to \infty$ w.p.1, which implies (**C2***) holds. □

The following lemma shows the Berry-Esseen Inequality for the percentile bootstrap approach.

Lemma A.9. *Under the assumptions of Theorem 3.4, we have* $\Pr\left(S_N^*(x) \le z \;\middle|\; \widehat{\mathbb{P}}_N\right) = \Phi(z)+o(1)$ *uniformly in* $x \in \mathcal{K}$ *and* $z \in \mathbb{R}$ *w.p.1.*

Proof. By the assumptions of Theorem 3.4, we know there exist σ_{min} such that $\sigma(x) > \sigma_{min} > 0$. In addition, by Lemma A.4, $\sup_{x\in\mathcal{K}} \mathbb{E}\left[|F(x,\xi)|^3\right] \le K_3$. By **(C1*)** and **(C2*)** in Lemma A.8, the following two inequalities hold w.p.1,

$$\lim_{N\to\infty} \sup_{x\in\mathcal{K}} \mathbb{E}\left[|F(x,\zeta)|^3 \;\middle|\; \widehat{\mathbb{P}}_N\right] \le \lim_{N\to\infty} \sup_{x\in\mathcal{K}} \left|\mathbb{E}\left[|F(x,\zeta)|^3 \;\middle|\; \widehat{\mathbb{P}}_N\right] - \mathbb{E}\,|F(x,\xi)|^3\right| + \sup_{x\in\mathcal{K}} \mathbb{E}\,|F(x,\xi)|^3 \le K_3,$$

$$\lim_{N\to\infty} \inf_{x\in\mathcal{K}} |\widehat{\sigma}_N(x)| \ge \inf_{x\in\mathcal{K}} |\sigma(x)| - \lim_{N\to\infty} \sup_{x\in\mathcal{K}} |\widehat{\sigma}_N(x) - \sigma(x)| \ge \sigma_{min}.$$

Then, by Lemma A.6, we have

$$\begin{aligned}\lim_{N\to\infty} \sup_{x\in\mathcal{K}} \sup_{z\in\mathbb{R}} \left|\Pr\left(S_N^*(x) \le z \;\middle|\; \widehat{\mathbb{P}}_N\right) - \Phi(z)\right| &\le \lim_{N\to\infty} C \frac{\sup_{x\in\mathcal{K}} \mathbb{E}\left[|F(x,\xi)|^3 \;\middle|\; \widehat{\mathbb{P}}_N\right]}{\inf_{x\in\mathcal{K}} \widehat{\sigma}_N^3(x)\sqrt{N}} \\ &\le \lim_{N\to\infty} C \frac{K_3}{\sigma_{min}^3 \sqrt{N}} = 0.\end{aligned}$$

□

A.5.4 Uniform Convergence on Bootstrap Percentile.

Define

$$z_N^\alpha(x) := \inf_z \left\{ z \in \mathbb{R} : \Pr\left(S_N^*(x) \le z \;\middle|\; \widehat{\mathbb{P}}_N\right) \ge 1-\alpha \right\}.$$

Recall

$$\mathbb{U}^\alpha_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N] = \inf\left\{ z \in \mathbb{R} : \Pr\left(\widehat{\mu}_N^*(x) \le z \;\middle|\; \widehat{\mathbb{P}}_N\right) \ge 1-\alpha \right\}.$$

We derive

$$\mathbb{U}^\alpha_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N] = \widehat{\mu}_N(x) + \frac{\widehat{\sigma}_N(x)}{\sqrt{N}} z_N^\alpha(x).$$

Letting

$$T_N^\alpha(x) := -\frac{\widehat{\sigma}_N(x)}{\sigma(x)} \cdot z_N^\alpha(x),$$

we further express the coverage probability as

$$\begin{aligned}\Pr\left(\mathbb{U}^\alpha_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N] \ge \mu(x)\right) &= \Pr\left(\widehat{\mu}_N(x) + \frac{\widehat{\sigma}_N(x)}{\sqrt{N}} \cdot z_N^\alpha(x) \ge \mu(x)\right) \\ &= \Pr\left(\frac{\sqrt{N}(\widehat{\mu}_N(x) - \mu(x))}{\sigma(x)} \ge -\frac{\widehat{\sigma}_N(x)}{\sigma(x)} \cdot z_N^\alpha(x)\right) \\ &= \Pr\left(S_N(x) \ge T_N^\alpha(x)\right). \qquad (28)\end{aligned}$$

Lemma A.10. $\lim_{N\to\infty} \sup_{x\in\mathcal{K}} |T_N^\alpha(x) - \Phi^{-1}(\alpha)| = 0,$ *w.p.1.*

Proof. Consider a sample path $\boldsymbol{\xi} = (\xi_1, \xi_2, \dots)$ and denote its sample space by Ξ^∞. Define

$$\Upsilon := \left\{ \boldsymbol{\xi} \in \Xi^\infty : \lim_{N\to\infty} \sup_{x\in\mathcal{K}} |T_N^\alpha(x) - \Phi^{-1}(\alpha)| = 0 \right\},$$

$$\Upsilon_0 := \left\{ \boldsymbol{\xi} \in \Xi^\infty : \lim_{N\to\infty} \sup_{x\in\mathcal{K}} \left| \frac{\widehat{\sigma}_N(x)}{\sigma(x)} - 1 \right| = 0 \right\},$$

$$\Upsilon_1 := \left\{ \boldsymbol{\xi} \in \Xi^\infty : \lim_{N\to\infty} \sup_{x\in\mathcal{K}} \left|(-z_N^\alpha(x)) - \Phi^{-1}(\alpha)\right| = 0 \right\},$$

and

$$\Upsilon_2 := \left\{ \boldsymbol{\xi} \in \Xi^\infty : \lim_{N\to\infty} \sup_{x\in\mathcal{K}} \sup_{z\in\mathbb{R}} \left| \Pr\left(S_N^*(x) \le z \;\middle|\; \widehat{\mathbb{P}}_N \right) - \Phi(z) \right| = 0 \right\}.$$

Our goal is to prove $\Pr(\boldsymbol{\xi} \in \Upsilon) = 1$. We now give our proof in the following three steps.

Step 1: we prove that $(\Upsilon_0 \cap \Upsilon_1) \subseteq \Upsilon$.

Consider $\boldsymbol{\xi} \in \Upsilon_0 \cap \Upsilon_1$. Observe that

$$\begin{aligned}
\left|T_N^\alpha(x) - \Phi^{-1}(\alpha)\right| &= \left| \left(\frac{\widehat{\sigma}_N(x)}{\sigma(x)} - 1 \right) (-z_N^\alpha(x)) + \left(-z_N^\alpha(x) - \Phi^{-1}(\alpha) \right) \right| \\
&\le \left| \frac{\widehat{\sigma}_N(x)}{\sigma(x)} - 1 \right| |-z_N^\alpha(x)| + \left| -z_N^\alpha(x) - \Phi^{-1}(\alpha) \right|.
\end{aligned}$$

By definition of Υ_1, we know $-z_N^\alpha(x) \to \Phi^{-1}(\alpha)$ uniformly in $x \in \mathcal{K}$. There exists N_1 such that for all $N \ge N_1$,

$$\sup_{x\in\mathcal{K}} |-z_N^\alpha(x)| \le \left|\Phi^{-1}(\alpha)\right| + 1.$$

Similarly, by definition of Υ_0 and Υ_1, for any $\varepsilon > 0$, there exists N_2 such that for all $N \ge N_2$,

$$\sup_{x\in\mathcal{K}} \left| \frac{\widehat{\sigma}_N(x)}{\sigma(x)} - 1 \right| < \varepsilon, \quad \text{and} \quad \sup_{x\in\mathcal{K}} \left| -z_N^\alpha(x) - \Phi^{-1}(\alpha) \right| < \varepsilon.$$

Therefore, for $N \ge N_0 := \max\{N_1, N_2\}$, we have

$$\left|T_N^\alpha(x) - \Phi^{-1}(\alpha)\right| \le \varepsilon \left(\left|\Phi^{-1}(\alpha)\right| + 1 \right) + \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, it follows that

$$\lim_{N\to\infty} \sup_{x\in\mathcal{K}} \left|T_N^\alpha(x) - \Phi^{-1}(\alpha)\right| = 0.$$

Thus, $\boldsymbol{\xi} \in \Upsilon$, and therefore $(\Upsilon_0 \cap \Upsilon_1) \subseteq \Upsilon$, which completes the proof for **Step 1**.

Step 2: we prove that $\Upsilon_2 \subseteq \Upsilon_1$.

Consider $\boldsymbol{\xi} \in \Upsilon_2$ and a given $\delta > 0$. Since $\Phi(z)$ is continuous and strictly increasing, its inverse $\Phi^{-1}(\cdot)$ is also continuous and strictly increasing on its domain $(0, 1)$. Let

$$\varepsilon(\delta) := \min\left\{ \Phi\left(\Phi^{-1}(1-\alpha) + \delta\right) - (1-\alpha),\ (1-\alpha) - \Phi\left(\Phi^{-1}(1-\alpha) - \delta\right) \right\}.$$

By definition of Υ_2, there exists N_0 such that for all $N \ge N_3$, we have

$$\sup_{x\in\mathcal{K}} \sup_{z\in\mathbb{R}} \left| \Pr\left(S_N^*(x) \le z \;\middle|\; \widehat{\mathbb{P}}_N \right) - \Phi(z) \right| < \varepsilon(\delta).$$

In other words, for all $N \geq N_3$, we have:

$$\left| \Pr\left(S_N^*(x) \leq z \;\middle|\; \widehat{\mathbb{P}}_N \right) - \Phi(z) \right| < \varepsilon(\delta), \quad \text{for all } x \in \mathcal{K},\ z \in \mathbb{R}.$$

Thus, at $z = \Phi^{-1}(1-\alpha) - \delta$:

$$\Pr\left(S_N^*(x) \leq \Phi^{-1}(1-\alpha) - \delta \;\middle|\; \widehat{\mathbb{P}}_N \right) < \Phi\left(\Phi^{-1}(1-\alpha) - \delta\right) + \varepsilon(\delta) \leq (1-\alpha) - \varepsilon(\delta) + \varepsilon(\delta) = 1-\alpha.$$

Similarly, at $z = \Phi^{-1}(1-\alpha) + \delta$:

$$\Pr\left(S_N^*(x) \leq \Phi^{-1}(1-\alpha) + \delta \;\middle|\; \widehat{\mathbb{P}}_N \right) > \Phi\left(\Phi^{-1}(1-\alpha) + \delta\right) - \varepsilon(\delta) \geq (1-\alpha) + \varepsilon(\delta) - \varepsilon(\delta) = 1-\alpha.$$

Therefore, it follows by the definition of $z_N^\alpha(x)$ that, for all $N \geq N_3$ and $x \in \mathcal{K}$, $z_N^\alpha(x) \in \left[\Phi^{-1}(1-\alpha) - \delta,\ \Phi^{-1}(1-\alpha) + \delta\right]$. Thus,

$$\sup_{x \in \mathcal{K}} \left| z_N^\alpha(x) - \Phi^{-1}(1-\alpha) \right| \leq \delta.$$

Note that $\Phi^{-1}(\alpha) = -\Phi^{-1}(1-\alpha)$. When choosing an arbitrarily small δ, we obtain

$$\lim_{N \to \infty} \sup_{x \in \mathcal{K}} \left| (-z_N^\alpha(x)) - \Phi^{-1}(\alpha) \right| = 0,$$

which completes the proof for **Step 2**.

Step 3: we prove that $\Pr(\boldsymbol{\xi} \in \Upsilon) = 1$.

By **(C2*)** in Lemma A.8, it is clear to see $\Pr(\boldsymbol{\xi} \in \Upsilon_0) = 1$. In addition, by Lemma A.7, $\Pr(\xi \in \Upsilon_2) = 1$. By previous two steps, we establish

$$(\Upsilon_0 \cap \Upsilon_2) \subseteq (\Upsilon_0 \cap \Upsilon_1) \subseteq \Upsilon.$$

Therefore, $\Pr(\boldsymbol{\xi} \in \Upsilon) = 1$. It completes the proof of Lemma A.10. □

A.5.5 Proof of Equation (26).

Now we finalize the proof of (26). It follows from (28) that

$$\begin{aligned}
&\Pr\left(\mathbb{U}^\alpha_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N] \geq \mu(x) \right) - (1-\alpha) \\
=& \Pr\left(S_N(x) \geq T_N^\alpha(x) \right) - (1-\alpha) \\
=& \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}\{S_N(x) \geq T_N^\alpha(x)\} \;\middle|\; T_N^\alpha(x) \right] \right] - 1 + \Phi\left(\Phi^{-1}(\alpha)\right) \\
=& \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}\{S_N(x) \geq T_N^\alpha(x)\} - [1 - \Phi(T_N^\alpha(x))] \;\middle|\; T_N^\alpha(x) \right] \right] + \mathbb{E}\left[[1 - \Phi(T_N^\alpha(x))]\right] - 1 + \Phi\left(\Phi^{-1}(\alpha)\right) \\
=& \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}\{S_N(x) \geq T_N^\alpha(x)\} - [1 - \Phi(T_N^\alpha(x))] \;\middle|\; T_N^\alpha(x) \right] \right] + \mathbb{E}[\Phi\left(\Phi^{-1}(\alpha)\right) - \Phi(T_N^\alpha(x))]
\end{aligned}$$

We can further derive

$$
\begin{aligned}
&\sup_{x\in\mathcal{K}}\left|\Pr\left(\mathbb{U}^{\alpha}_{\text{Efron}}[\mu(x)|\widehat{\mathbb{P}}_N]\geq\mu(x)\right)-(1-\alpha)\right|\\
\leq&\sup_{x\in\mathcal{K}}\left|\mathbb{E}\left[\mathbb{E}\left[\mathbb{1}\{S_N(x)\geq T^{\alpha}_N(x)\}-[1-\Phi\left(T^{\alpha}_N(x)\right)]\;\middle|\;T^{\alpha}_N(x)\right]\right]\right|+\sup_{x\in\mathcal{K}}\left|\mathbb{E}[\Phi\left(\Phi^{-1}(\alpha)\right)-\Phi\left(T^{\alpha}_N(x)\right)]\right|\\
\leq&\mathbb{E}\left[\sup_{x\in\mathcal{K}}\left|\mathbb{E}\left[\mathbb{1}\{S_N(x)\geq T^{\alpha}_N(x)\}-[1-\Phi\left(T^{\alpha}_N(x)\right)]\;\middle|\;T^{\alpha}_N(x)\right]\right|\right]+\mathbb{E}\left[\sup_{x\in\mathcal{K}}\left|\Phi\left(\Phi^{-1}(\alpha)\right)-\Phi\left(T^{\alpha}_N(x)\right)\right|\right]
\end{aligned}
$$

Therefore, to prove (26), it suffices to prove

$$
\lim_{N\to\infty}\mathbb{E}\left[\sup_{x\in\mathcal{K}}\left|\mathbb{E}\left[\mathbb{1}\{S_N(x)\geq T^{\alpha}_N(x)\}-[1-\Phi\left(T^{\alpha}_N(x)\right)]\;\middle|\;T^{\alpha}_N(x)\right]\right|\right]=0, \tag{29}
$$

and

$$
\lim_{N\to\infty}\mathbb{E}\left[\sup_{x\in\mathcal{K}}\left|\Phi\left(\Phi^{-1}(\alpha)\right)-\Phi\left(T^{\alpha}_N(x)\right)\right|\right]=0. \tag{30}
$$

Let us first prove (29). By Lemma A.7, for any $\epsilon>0$, there exists N_4 such that for $N>N_4$,

$$
\begin{aligned}
&\sup_{x\in\mathcal{K}}\left|\mathbb{E}\left[\mathbb{1}\{S_N(x)\geq T^{\alpha}_N(x)\}-[1-\Phi\left(T^{\alpha}_N(x)\right)]\;\middle|\;T^{\alpha}_N(x)\right]\right|\\
\leq&\sup_{x\in\mathcal{K}}\sup_{z\in\mathbb{R}}\left|\Pr(S_N(x)\geq z)-[1-\Phi(z)]\right|<\epsilon,\quad\text{w.p.1},
\end{aligned}
$$

and then

$$
\mathbb{E}\left[\sup_{x\in\mathcal{K}}\left|\mathbb{E}\left[\mathbb{1}\{S_N(x)\geq T^{\alpha}_N(x)\}-[1-\Phi\left(T^{\alpha}_N(x)\right)]\;\middle|\;T^{\alpha}_N(x)\right]\right|\right]<\epsilon.
$$

This proves (29).

Next, we prove (30). Since $\Phi(\cdot)$ is continuous on $\mathbb{R}$, for any $\epsilon>0$, there exists $\delta>0$ such that whenever

$$
|T^{\alpha}_N(x)-\Phi^{-1}(\alpha)|<\delta,
$$

we have

$$
\left|\Phi(T^{\alpha}_N(x))-\Phi(\Phi^{-1}(\alpha))\right|<\epsilon.
$$

By Lemma A.10, there exists N_5 such that for $N>N_5$,

$$
\sup_{x\in\mathcal{K}}|T^{\alpha}_N(x)-\Phi^{-1}(\alpha)|<\delta,\quad\text{w.p.1},
$$

and thus,

$$
\sup_{x\in\mathcal{K}}\left|\Phi(T^{\alpha}_N(x))-\Phi(\Phi^{-1}(\alpha))\right|<\epsilon,\quad\text{w.p.1}.
$$

This implies (30), that is, for $N>N_5$,

$$
\mathbb{E}\left[\sup_{x\in\mathcal{K}}\left|\Phi\left(\Phi^{-1}(\alpha)\right)-\Phi\left(T^{\alpha}_N(x)\right)\right|\right]<\epsilon.
$$

B Theorems Used in the Paper

Theorem B.1 (Theorem 1, Rockafellar and Uryasev (2000)). *Let $h(x, \omega)$ be a random function where $x \in \mathcal{X}$ and ω belong to an arbitrary probability space with distribution $\mathbb{Q}$. Let $q_\alpha(x)$ denote the $100(1-\alpha)$-percentile of $h(x, \omega)$ and*

$$H_\alpha(x, t) = t + \frac{1}{\alpha} \int [h(x, \omega) - t]_+ \mathbb{Q}(d\omega),$$

where $t \in \mathbb{R}$. Then, for all $x \in \mathcal{X}$, we have

$$\frac{1}{\alpha} \int_0^\alpha q_\tau(x) d\tau = \min_{t \in \mathbb{R}} H_\alpha(x, t).$$

Theorem B.2 (Theorem 2, Athreya (1983)). *Suppose $\liminf MN^{-\phi} > 0$ for some $\phi > 0$ as $M, N \to \infty$, and $\mathbb{E}_{\mathbb{P}} |F(\xi) - \mu|^\theta < \infty$ for some $\theta \geq 1$ such that $\theta\phi > 1$. Then, as $M, N \to \infty$, we have*

$$\frac{1}{M} \sum_{m=1}^{M} F\left(\zeta_m(\widehat{\mathbb{P}}_N)\right) \to 1 \quad \textit{w.p.1.}$$

Theorem B.3 (Lemma 21.2, Vaart (1998)). *The quantile function of a cumulative distribution function $\mathcal{F}$ is the generalized inverse $\mathcal{F}^{-1} : (0, 1) \to \mathbb{R}$ given by*

$$\mathcal{F}^{-1}(p) = \inf\{x : \mathcal{F}(x) \leq p\}.$$

For any any sequence of cumulative distribution functions, $\mathcal{F}_N$ converges to $\mathcal{F}$ in distribution if and only if $\mathcal{F}_N^{-1}$ converges to $\mathcal{F}^{-1}$ in distribution.

Theorem B.4 (Theorem 2.1, Mallows and Richter (1969)). *Let η be a random variable, $\mathcal{A}$ be an event, and σ_η be its standard deviation. Then we have:*

$$\left| \mathbb{E}[\eta | \mathcal{A}] - \mathbb{E}[\eta] \right| \leq \sigma_\eta \left(\frac{1 - \Pr(\mathcal{A})}{\Pr(\mathcal{A})} \right)^{1/2}.$$

Theorem B.5 (Theorem 10.8 (Rockafellar, 2015)). *Let $\mathcal{C}$ be an open convex set. Let $(g_1, g_2, \cdots)$ be a sequence of finite convex functions on $\mathcal{C}$. Suppose that the sequence converges pointwise on a dense subset $\mathcal{D} \subseteq \mathcal{C}$ and the limit is finite. Then, the sequence $(g_1, g_2, \cdots)$ converges uniformly to a continuous function on any compact subset inside $\mathcal{C}$.*

Theorem B.6 (Theorem 7.53, Shapiro et al. (2021)). *Let $\mathcal{K}$ be a nonempty compact subset of $\mathbb{R}^n$ and suppose that (i) for any $x \in \mathcal{K}$ the function $F(\cdot, \xi)$ is continuous at x for almost every $\xi \in \Xi$, (ii) $F(x, \xi)$, $x \in \mathcal{K}$, is dominated by an integrable function, and (iii) the sample is iid. Then, the expected function $f(x)$ is finite valued and continuous on $\mathcal{K}$, and the sample mean $\hat{f}_N(x)$ converges to $f(x)$ w.p.1 uniformly on $\mathcal{K}$.*

C Data for the Newsvendor problem in Section 5.3

$$p = -2, h = 9, b = 5$$

$$\mu_1 = [60.89, 48.58, 46.81, 56.54, 61.58, 52.69, 69.42, 60.54, 54.43, 51.76]^\intercal$$

$$\mu_2 = [50.30, 61.87, 53.16, 41.79, 51.94, 62.14, 45.47, 45.26, 55.95, 55.95]^\intercal$$

$$\Sigma_1 = \begin{bmatrix} 9.27 & 2.84 & -0.07 & 1.19 & -0.48 & 1.40 & 2.87 & 4.06 & -1.40 & -1.96 \\ 2.84 & 5.90 & -2.83 & 0.21 & 2.27 & -2.40 & -0.89 & 4.22 & 3.43 & 2.78 \\ -0.07 & -2.83 & 5.48 & -0.30 & 0.90 & 3.54 & -4.51 & -2.45 & -2.91 & -4.95 \\ 1.19 & 0.21 & -0.30 & 7.99 & -1.02 & -1.27 & -0.15 & -1.55 & -1.69 & -0.36 \\ -0.48 & 2.27 & 0.90 & -1.02 & 9.48 & -0.08 & -3.69 & 2.71 & -0.69 & -0.34 \\ 1.40 & -2.40 & 3.54 & -1.27 & -0.08 & 6.94 & -1.26 & -2.73 & 0.01 & -5.19 \\ 2.87 & -0.89 & -4.51 & -0.15 & -3.69 & -1.26 & 12.05 & -0.16 & -0.16 & 2.44 \\ 4.06 & 4.22 & -2.45 & -1.55 & 2.71 & -2.73 & -0.16 & 9.16 & -0.77 & 1.94 \\ -1.40 & 3.43 & -2.91 & -1.69 & -0.69 & 0.01 & -0.16 & -0.77 & 7.41 & 2.24 \\ -1.96 & 2.78 & -4.95 & -0.36 & -0.34 & -5.19 & 2.44 & 1.94 & 2.24 & 6.70 \end{bmatrix}$$

$$\Sigma_2 = \begin{bmatrix} 6.32 & 2.99 & -0.06 & 0.73 & -0.33 & 1.36 & 1.55 & 2.51 & -1.19 & -1.75 \\ 2.99 & 9.57 & -4.09 & 0.19 & 2.44 & -3.60 & -0.74 & 4.02 & 4.49 & 3.83 \\ -0.06 & -4.09 & 7.06 & -0.25 & 0.86 & 4.74 & -3.35 & -2.08 & -3.40 & -6.08 \\ 0.73 & 0.19 & -0.25 & 4.37 & -0.64 & -1.11 & -0.07 & -0.86 & -1.29 & -0.29 \\ -0.33 & 2.44 & 0.86 & -0.64 & 6.74 & -0.08 & -2.04 & 1.71 & -0.60 & -0.31 \\ 1.36 & -3.60 & 4.74 & -1.11 & -0.08 & 9.65 & -0.98 & -2.41 & 0.01 & -6.62 \\ 1.55 & -0.74 & -3.35 & -0.07 & -2.04 & -0.98 & 5.17 & -0.08 & -0.10 & 1.72 \\ 2.51 & 4.02 & -2.08 & -0.86 & 1.71 & -2.41 & -0.08 & 5.12 & -0.59 & 1.57 \\ -1.19 & 4.49 & -3.40 & -1.29 & -0.60 & 0.01 & -0.10 & -0.59 & 7.83 & 2.49 \\ -1.75 & 3.83 & -6.08 & -0.29 & -0.31 & -6.62 & 1.72 & 1.57 & 2.49 & 7.83 \end{bmatrix}$$

$$\varepsilon \sim \begin{bmatrix} \mathcal{U}(-5.37, 26.27) \\ \mathcal{U}(6.74, 14.16) \\ \mathcal{U}(3.22, 17.68) \\ \mathcal{U}(-7.48, 28.38) \\ \mathcal{U}(-4.89, 25.79) \\ \mathcal{U}(-0.21, 16.11) \\ \mathcal{U}(-12.14, 32.99) \\ \mathcal{U}(-7.74, 28.64) \\ \mathcal{U}(0.77, 20.13) \\ \mathcal{U}(2.13, 18.77) \end{bmatrix}.$$