PARAMETRIC APUB FOR DECISION MAKING UNDER PARAMETER UNCERTAINTY

Anonymous authors
Paper under double-blind review

ABSTRACT

Parametric optimization can exploit a known distribution family but usually treats its fitted parameters as certain. We develop parametric Average Percentile Upper Bound (APUB) optimization to adjust a decision and certify its expected cost against this epistemic uncertainty. APUB refits the parametric model on bootstrap datasets and minimizes an average of upper percentiles of the resulting expected-cost functions. We study generative parametric bootstrap and weighted-likelihood bootstrap as two prior-free constructions in one framework. The theory establishes decision consistency under construction-specific conditions, regular fixed-decision coverage, and finite-ensemble error control; an adaptive algorithm grows and reuses the fitted ensemble. In paired, fixed-truth mixed-normal and mixed-log-normal newsvendor experiments, positive-level PB- and WLB-APUB have lower mean regret than SAA and fitted-model optimization at small samples, while the adjustment contracts as the sample size increases. Only APUB displays approximately level-tracking coverage curves, giving its nominal level an interpretable one-sided confidence meaning. APUB avoids this prior-specification risk through a data-driven resampling adjustment.

1 INTRODUCTION

Decisions about new products, new services, and changing markets often precede the accumulation of reliable data. A retailer planning an initial inventory may have only a small pilot sample, while historical demand concerns products or customer groups that are no longer representative. Expert knowledge may suggest a distribution family—for example, a mixture representing heterogeneous customer groups—without identifying its parameters or supporting a credible parameter prior. Optimization must then exploit available structure while acknowledging uncertainty about the fitted model.

The uncertainty of interest concerns a decision's expected performance. Outcomes remain random even when their distribution is known; uncertainty about expected cost arises because its parameters must be learned. Fitted-model optimization (FMO) exploits the assumed family but treats the estimate as known. Sample-average approximation (SAA) avoids specifying the family but does not exploit this information. Bayesian risk optimization and distributionally robust optimization (DRO) offer alternative adjustments through a posterior or an ambiguity set. When prior knowledge is limited, their specification and calibration also require judgment.

We develop parametric Average Percentile Upper Bound (APUB) optimization, extending the bootstrap framework of Liu et al. (2026) to fitted expected-cost models. Each bootstrap fit supplies an expected-cost function, and APUB averages their upper percentiles at the candidate decision. Two established statistical procedures generate these fits: the weighted-likelihood bootstrap (WLB), which perturbs observation weights, and the parametric bootstrap (PB), which simulates and refits datasets from the original fitted model. Both belong to parametric APUB because the expected-cost functions use the assumed parametric family. The distinction is how the ensemble is generated, not whether the decision model is parametric.

A statistical upper-confidence construction. The APUB ensemble represents estimation variability conditional on the observed data. Its target is the unknown expected cost, rather than the tail of individual realized losses. Consequently, APUB is a decision-and-certification procedure: in addition to selecting a decision, it reports an upper estimate of how costly that decision can be under

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

parameter-estimation uncertainty. Although its upper-average calculation has the mathematical representation of CVaR, the inferential role differs from conventional CVaR applied to outcome losses. This is analogous to the distinction between a bootstrap percentile confidence bound and an outcome quantile.

Advantages and scope. Parametric APUB has three related advantages for parameter uncertainty. First, it makes risk aversion data-driven: at a chosen nominal level γ, it aggregates unfavorable *bootstrap-fitted expected-cost functions*, not realized-loss tails. The adjustment therefore responds to the observed sampling variability of the fitted model and requires neither a subjective parameter prior nor a separately calibrated ambiguity radius. It is not a worst-case criterion, and under the stated contraction conditions the adjustment vanishes with information at a fixed γ; hence protection need not impose a persistent conservative penalty.

Second, γ has a statistical interpretation. Under the bootstrap and selected-decision conditions in Section 4, APUB's reported objective is intended as a nominal one-sided upper confidence assessment of the true expected cost at its recommendation. The empirical coverage curves in Section 5.1 show the corresponding approximately linear, level-tracking behavior: the chosen nominal level is close to the observed coverage, rather than an abstract tuning parameter.

Third, APUB is independent of parameter-prior specification. It retains the assumed distribution family but uses fitted-model resampling rather than a prior. It is consequently not exposed to the small-sample loss that can result when a Bayesian prior is misspecified. These are advantages for epistemic-risk adjustment and certification, not claims that APUB uniformly minimizes every repeated-sampling regret summary or protects against model-family misspecification.

2 Related work

Optimization under parameter and distribution uncertainty. Data-driven stochastic optimization must distinguish outcome randomness from epistemic uncertainty about the data-generating law or its parameters. Distributionally robust optimization (DRO) addresses the latter by minimizing worst-case expected cost over an ambiguity set. Classical moment-based formulations protect against uncertainty in estimated means, covariances, and higher-order distributional features (Delage & Ye, 2010; Wiesemann et al., 2014). More generally, data-driven robust optimization constructs uncertainty sets from data, while divergence- and Wasserstein-based DRO use neighborhoods of an empirical or fitted reference distribution (Bertsimas et al., 2018; Mohajerin Esfahani & Kuhn, 2018; Shapiro, 2017). These approaches can yield finite-sample or asymptotic guarantees when their ambiguity sets and radii are appropriately calibrated, but their decisions depend materially on the choice of statistical distance, support restrictions, and calibration target. In our additive newsvendor, first-order Wasserstein DRO with a fixed ground norm and unrestricted demand support adds an inventory-independent penalty and leaves the SAA minimizers unchanged. We therefore use full-covariance moment DRO as a comparator; its radius grid is a sensitivity analysis, not a confidence calibration.

Bayesian risk and Bayesian distributional robustness. Bayesian risk optimization propagates parameter uncertainty through a posterior and optimizes a posterior functional of model-conditional expected cost (Wu et al., 2018). Posterior-CVaR emphasizes unfavorable posterior draws of that expected cost, and nested stochastic-gradient methods have been developed for this objective (Cakmak et al., 2020). Bayesian-DRO combines posterior learning of an unknown parametric distribution with a further ambiguity set around a reference model (Shapiro et al., 2023); related work develops Bayesian criteria for constructing or comparing ambiguity sets (Gupta, 2019). Thus Bayesian risk optimization, Bayesian-DRO, and frequentist DRO all offer principled but distinct ways to respond to parameter-estimation error: they require, respectively, a prior and likelihood, a prior plus an ambiguity-set specification, or an ambiguity-set calibration rule. When the model is well specified and the likelihood becomes informative, posterior concentration can diminish the influence of either an aligned or a misspecified prior; this is the mechanism reflected by the convergence of the two Bayesian rows at the largest sample size in our study. At small samples, however, prior influence can be material. APUB takes a different route: it uses resampling of fitted models and therefore avoids parameter-prior misspecification, while retaining the assumed distributional family.

Bootstrap confidence bounds and parametric APUB. Bootstrap percentile limits use quantiles of an estimator's resampling distribution (Efron, 1987; Efron & Tibshirani, 1993; Davison & Hinkley, 1997). APUB averages upper percentiles of bootstrap estimates of expected cost (Liu et al., 2026). Our method differs from DRO because it reweights *fitted expected-cost functions*, rather than optimizing over an ambiguity set of outcome distributions; it differs from Bayesian methods because it does not posit a parameter prior. The upper average is a candidate statistical bound for a fixed decision only under the relevant bootstrap approximation conditions, and it is not a worst-case realized-loss criterion or a simultaneous guarantee for a data-selected optimizer.

3 Parametric APUB with two bootstrap constructions

Let $D_N = (\xi_1, \ldots, \xi_N)$ be iid from P_{θ_0} in a correctly specified family $\{P_\theta : \theta \in \Theta\}$ with density p_θ. For feasible decisions $x \in \mathcal{X}$, define

$$H(x,\theta) = \mathbb{E}_{P_\theta}\ell(x,\xi), \qquad v^\star = \min_{x\in\mathcal{X}} H(x,\theta_0), \qquad \mathcal{X}^\star = \arg\min_{x\in\mathcal{X}} H(x,\theta_0). \tag{1}$$

The original fitting rule returns $\widehat{\theta}_N$ by approximately maximizing $M_N(\theta) = N^{-1}\sum_i \log p_\theta(\xi_i)$ over its declared parameter constraints. The point-estimate baselines solve

$$x_N^{\mathrm{FMO}} \in \arg\min_{x\in\mathcal{X}} H(x,\widehat{\theta}_N), \qquad x_N^{\mathrm{SAA}} \in \arg\min_{x\in\mathcal{X}} \frac{1}{N}\sum_i \ell(x,\xi_i). \tag{2}$$

3.1 Bootstrap fits and percentile bounds

For a fixed decision, bootstrap fitting propagates parameter-estimation uncertainty to its expected cost. Standard PB simulates from the original fitted family and refits (Efron & Tibshirani, 1993; Davison & Hinkley, 1997; Efron, 2012). WLB assigns continuous random likelihood weights and refits the observed sample (Newton & Raftery, 1994). Exchangeably weighted bootstrap theory provides a route to sampling-distribution approximation (Præstgaard & Wellner, 1993), and loss-likelihood bootstrap methods extend weighted fitting beyond likelihoods (Lyddon et al., 2019). The present extension passes both constructions through nonlinear parametric fitting before applying upper-percentile aggregation. We compare PB and WLB as alternative inputs to the same APUB objective and adaptive solver.

Generative parametric bootstrap (PB). Conditionally on D_N, independently draw and refit

$$\xi^*_{1b}, \ldots, \xi^*_{Nb} \overset{\text{iid}}{\sim} P_{\widehat{\theta}_N}, \qquad M^{\mathrm{P}}_{N,b}(\theta) = \frac{1}{N}\sum_i \log p_\theta(\xi^*_{ib}), \tag{3}$$

obtaining $\widetilde{\theta}^{\mathrm{P}}_{N,b}$. These are draws from a fitted parametric law, not the empirical distribution of the original observations. We use the refitted parameters directly; reflecting estimates around the original fit is not part of this procedure.

Weighted-likelihood bootstrap (WLB). Independently across b, draw $e_{ib} \overset{\text{iid}}{\sim} \mathrm{Exp}(1)$ and set

$$\lambda_{ib} = \frac{e_{ib}}{\sum_r e_{rb}}, \qquad \lambda_b \sim \mathrm{Dirichlet}(1,\ldots,1), \qquad M^{\mathrm{W}}_{N,b}(\theta) = \sum_i \lambda_{ib}\log p_\theta(\xi_i). \tag{4}$$

The fit $\widetilde{\theta}^{\mathrm{W}}_{N,b}$ approximately maximizes this criterion. Each weight has conditional mean $1/N$; weights within a vector are dependent and sum to one. All observations retain positive weight almost surely. This nonparametric weighting mechanism is applied to a parametric fitting problem.

For $m \in \{\mathrm{W}, \mathrm{P}\}$, let $Q^m_N = \mathcal{L}(\widetilde{\theta}^m_N \mid D_N)$ be the conditional law of the implemented fit, including its initialization and fitting randomness, and define $h^m_b(x) = H(x, \widetilde{\theta}^m_{N,b})$. Draws are conditionally independent when both bootstrap randomization and fitting seeds are independent across b. Let $F^m_{N,x}$ be the conditional CDF of $H(x, \widetilde{\theta}^m_N)$. Its upper percentile $F^{m,-1}_{N,x}(\gamma)$, at a *nominal* one-sided confidence level of γ, is a bootstrap percentile candidate for an upper confidence bound. Its frequentist validity is an additional property, not a consequence of defining this quantile.

3.2 UPPER-PERCENTILE AGGREGATION AND OPTIMIZATION

Assume conditional fitted costs are integrable at each decision. For $\gamma \in (0,1)$, define the ideal parametric APUB problem by

$$U^m_{N,\gamma}(x) = \frac{1}{1-\gamma}\int_\gamma^1 F^{m,-1}_{N,x}(u)\,du = \inf_t \left\{ t + \frac{\mathbb{E}_{Q^m_N}[H(x,\widetilde{\theta}^m_N) - t]_+}{1-\gamma} \right\}, \tag{5}$$
$$x^m_{N,\gamma} \in \arg\min_{x\in\mathcal{X}} U^m_{N,\gamma}(x), \qquad [z]_+ = \max(z,0).$$

The upper average is at least the corresponding percentile, so any valid percentile coverage lower bound transfers to it. The outcome expectation in (1) averages market randomness for a fixed parameter; the outer bootstrap operation addresses uncertainty about that expectation. With B fits, the computable problem and its averaging endpoint are

$$U^m_{N,B,\gamma}(x) = \min_t \left\{ t + \frac{1}{(1-\gamma)B}\sum_b [h^m_b(x) - t]_+ \right\}, \qquad x^m_{N,B,\gamma} \in \arg\min_{x\in\mathcal{X}} U^m_{N,B,\gamma}(x), \tag{6}$$

$$\overline{H}^m_{N,B}(x) = \frac{1}{B}\sum_b h^m_b(x), \qquad x^{m,\mathrm{BAG}}_{N,B} \in \arg\min_{x\in\mathcal{X}} \overline{H}^m_{N,B}(x). \tag{7}$$

At $\gamma = 0$ the finite objective equals BAG, which need not equal SAA or FMO after nonlinear fitting. Comparing BAG and APUB on the same ensemble isolates upper-tail aggregation. A model module supplies resampling, fitting, and H and its (sub)gradient. The outer optimizer is unchanged between WLB and PB.

Algorithm 1 presents an adaptive framework for solving (5) by incrementally appending ΔB new bootstrap fits to the sample objective (6). Appendix A details the algorithmic updates, stopping rules, and resumption protocols. The practical implementation monitors decision stability, objective value variations at prior solutions, estimated Monte Carlo simulation error, tail ensemble size, and subproblem optimization accuracy. These operational criteria govern numerical convergence; a simultaneous statistical certificate can be established by allocating error probabilities across stages under Theorem 3.

Algorithm 1 Adaptive APUB with a selectable bootstrap generator

Require: Data D_N, generator m, level γ, increment ΔB, tolerances, resource budget
1: Initialize an empty ensemble and a feasible decision
2: **loop**
3: Append ΔB new fits generated by WLB or PB
4: Warm-start optimization using all accumulated fits; refine within the optimization budget
5: Save fits, solution, optimization accuracy, stability diagnostics, and random state
6: **if** optimization and consecutive-stage stability criteria are satisfied **then**
7: **return** the current solution
8: **end if**
9: **if** the resource budget is exhausted **then**
10: **return** the solution and achieved diagnostics
11: **end if**
12: **end loop**

4 THEORY: FROM FITTED COSTS TO CERTIFIED DECISIONS

APUB has two logically distinct ingredients. A resampling construction must first approximate the unknown expected-cost function; upper aggregation then transfers that approximation to a decision and a reported upper cost. This separation is useful because WLB and PB require different fitting arguments but have the same decision consequence. Proofs appear in Appendix B.

The following two propositions establish the fitted-function contraction needed for Theorem 1. For WLB, the key step is controlling the weighted likelihood uniformly over parameters.

Proposition 1 (WLB contraction). *For the WLB construction, suppose:*

(A1) $\Theta \subset \mathbb{R}^q$ *and* $\mathcal{X} \subset \mathbb{R}^d$ *are nonempty compact sets, and* $\theta_0 \in \Theta$.

(A2) $\ell_\theta(\xi) := \log p_\theta(\xi)$ *is measurable, finite and continuous in* θ *for* P_{θ_0}*-almost every* ξ*, with* $\mathbb{E}_{\theta_0} F(\xi) < \infty$*, where* $F(\xi) = \sup_{\theta\in\Theta} |\ell_\theta(\xi)|$.

(A3) $M(\theta) = \mathbb{E}_{\theta_0} \ell_\theta(\xi)$ *has maximizer set* $\Theta_0 = \{\theta_0\}$ *or the finite set of its equivalent component-label permutations.*

(A4) The implemented fit is measurable and its global criterion gap satisfies $R_N = \sup_{\theta\in\Theta} M_N^{\mathrm{W}}(\theta) - M_N^{\mathrm{W}}(\widetilde{\theta}_N^{\mathrm{W}}) \to 0$ *in joint probability over data, weights, and fitting randomness.*

(A5) H *is continuous on* $\mathcal{X} \times \Theta$ *and* $H(x, \theta) = H(x, \theta_0)$ *for* $\theta \in \Theta_0$.

Then $a_N^{\mathrm{W}} = \mathbb{E}[\sup_x |H(x, \widetilde{\theta}_N^{\mathrm{W}}) - H(x, \theta_0)| \mid D_N] \to 0$ *in probability.*

The conditions separate uniform likelihood approximation, identification, numerical fitting, and the expected-cost interface. The compactness assumption also supplies a bound for dominated convergence; replacing it requires suitable tightness and uniform-integrability arguments. Label invariance avoids requiring one particular mixture labeling.

PB requires a separate argument because the fitted distribution generates the synthetic training samples. For a generating parameter ϑ, write $M_\vartheta(\theta) = \mathbb{E}_{P_\vartheta} \ell_\theta(\xi)$ and

$$\omega_h(\xi) = \sup_{\theta,\theta'\in\Theta:\|\theta-\theta'\|\le h} |\ell_\theta(\xi) - \ell_{\theta'}(\xi)|.$$

Proposition 2 (PB contraction). *Assume (A1), (A3), and (A5), and the following PB-specific conditions:*

(P1) The measurable original fit lies in Θ *and* $\mathrm{dist}(\widehat{\theta}_N, \Theta_0) \to 0$ *in probability.*

(P2) For some neighborhood $V \subset \Theta$ *of* Θ_0*, likelihoods are measurable, finite and continuous in* θ*,* P_ϑ*-almost surely for each* $\vartheta \in V$*, and*

$$C_F := \sup_{\vartheta\in V} \mathbb{E}_{P_\vartheta} F^2 < \infty, \qquad \lim_{h\downarrow 0} \sup_{\vartheta\in V} \mathbb{E}_{P_\vartheta} \omega_h = 0.$$

Moreover $\sup_{\theta\in\Theta} |M_\vartheta(\theta) - M_{\theta_0}(\theta)| \to 0$ *as* $\mathrm{dist}(\vartheta, \Theta_0) \to 0$.

(P3) Synthetic observations are conditionally iid from $P_{\widehat{\theta}_N}$*, and the measurable refit's global gap satisfies* $R_N^{\mathrm{P}} := \sup_{\theta\in\Theta} M_N^{\mathrm{P}}(\theta) - M_N^{\mathrm{P}}(\widetilde{\theta}_N^{\mathrm{P}}) \to 0$ *in joint probability.*

Then $a_N^{\mathrm{P}} = \mathbb{E}[\sup_x |H(x, \widetilde{\theta}_N^{\mathrm{P}}) - H(x, \theta_0)| \mid D_N] \to 0$ *in probability.*

Condition (P1) concerns the distribution that generates the synthetic data; (P2) controls its changing sampling law and a uniform law of large numbers for the refit. The second-moment envelope is stronger than WLB's first-moment condition because it permits a short conditional Chebyshev argument. Original-fit consistency in (P1) follows, for example, from the ordinary uniform likelihood law, identification, compactness, and an asymptotically vanishing original global fitting gap.

Let Q_N denote the conditional law of the fitted parameter generated by either construction and define

$$a_N := \mathbb{E}_{Q_N} \sup_{x\in\mathcal{X}} |H(x, \widetilde{\theta}_N) - H(x, \theta_0)|.$$

Theorem 1 (Decision transfer under bootstrap contraction). *Suppose* $\mathcal{X}$ *and* Θ *are nonempty compact sets and* H *is continuous on their product, with the same values for equivalent parameter labels. For WLB assume the likelihood, identification, and fitting conditions of Proposition 1; for PB assume those of Proposition 2. Then*

$$a_N = \mathbb{E}_{Q_N} \sup_{x\in\mathcal{X}} |H(x, \widetilde{\theta}_N) - H(x, \theta_0)| \xrightarrow{\mathbb{P}} 0, \sup_{x\in\mathcal{X}} |U_{N,\gamma}(x) - H(x, \theta_0)| \le \frac{a_N}{1-\gamma}.$$

For fixed $\gamma \in (0,1)$, let $\varepsilon_{N,B} \geq \sup_x |U_{N,B,\gamma}(x) - U_{N,\gamma}(x)|$. If a measurable $\widehat{x} \in \mathcal{X}$ is δ-optimal for the finite problem, meaning $U_{N,B,\gamma}(\widehat{x}) \leq \inf_{x\in\mathcal{X}} U_{N,B,\gamma}(x) + \delta$, then

$$H(\widehat{x}, \theta_0) - v^\star \leq \frac{2a_N}{1-\gamma} + 2\varepsilon_{N,B} + \delta. \tag{8}$$

If $\varepsilon_{N,B}, \delta \to 0$ in probability, $\operatorname{dist}(\widehat{x}, \mathcal{X}^\star) \to 0$ in probability.

The bound separates statistical fitting error, finite-ensemble error, and optimization error. It neither requires convexity nor uniqueness of the optimal decision. Lemma 2 in the appendix gives the related fixed-decision normal expansion. It explains why the APUB adjustment is an estimation-uncertainty adjustment, but a fixed-decision statement alone cannot justify coverage after optimizing over the data.

Theorem 2 (Coverage at the ideal data-selected APUB optimizer)**.** *Fix $\gamma \in (0,1)$ and let $x_{N,\gamma}$ be a measurable optimizer of the ideal APUB objective. Suppose the assumptions of Theorem 1 hold, the population minimizer is unique, $\mathcal{X}^\star = \{x^\star\}$, and $x_{N,\gamma} \xrightarrow{\mathbb{P}} x^\star$. Let $\mathcal{V}$ be a neighborhood of $x^\star$. Write*

$$G_N(x) = \sqrt{N}\{H(x, \widehat{\theta}_N) - H(x, \theta_0)\}.$$

Assume that, for a continuous function $s : \mathcal{V} \to (0,\infty)$,

$$G_N(x^\star)/s(x^\star) \Rightarrow \mathcal{N}(0,1), \tag{9}$$

$$\lim_{\delta\downarrow 0} \limsup_N \mathbb{P}\left\{ \sup_{x\in\mathcal{V}:\|x-x^\star\|\leq\delta} |G_N(x) - G_N(x^\star)| > \varepsilon \right\} = 0 \quad \text{for every } \varepsilon > 0, \tag{10}$$

$$\sup_{x\in\mathcal{V}} W_1\Big(\mathcal{L}\Big(\sqrt{N}[H(x,\widetilde{\theta}_N) - H(x,\widehat{\theta}_N)] \mid D_N \Big), \mathcal{N}(0, s(x)^2) \Big) \xrightarrow{\mathbb{P}} 0. \tag{11}$$

Then, with $\kappa_\gamma = \phi(\Phi^{-1}(\gamma))/(1-\gamma)$,

$$\mathbb{P}\{H(x_{N,\gamma}, \theta_0) \leq U_{N,\gamma}(x_{N,\gamma})\} \longrightarrow \Phi(\kappa_\gamma) > \gamma. \tag{12}$$

This is a selected-decision result for the ideal bootstrap objective at a fixed population. The local uniform bootstrap condition supplies the first-order upper adjustment throughout a neighborhood of $x^\star$; stochastic equicontinuity controls the effect of moving to the random optimizer.

Finally, the following theorem bounds the simulation error from replacing the ideal conditional bootstrap law by a finite fitted ensemble.

Theorem 3 (Conditional finite-ensemble error)**.** *Condition on D_N. Suppose the fitted parameters are iid from Q_N, $H(x,\theta) \in [a,b]$ simultaneously for all $x \in \mathcal{X}$ and Q_N-almost every θ, and H is uniformly L-Lipschitz in x. For a predetermined finite h-net of size K_h and $\eta \in (0,1)$, with conditional probability at least $1-\eta$,*

$$\sup_x |U_{N,B,\gamma}(x) - U_{N,\gamma}(x)| \leq \frac{b-a}{1-\gamma}\sqrt{\frac{\log(2K_h/\eta)}{2B}} + 2Lh. \tag{13}$$

The guarantee controls simulation error, not population coverage, and does not make a practical stopping diagnostic a confidence guarantee. A growing ensemble reduces simulation error; it does not increase the number of original observations.

5 Experiments

We use a multi-product newsvendor to separate structural estimation from uncertainty aggregation. With demand vector D, inventory $x \geq 0$, selling price p_j, purchase cost c_j, and salvage s_j, where $p_j > c_j > s_j$, the cost is $\ell(x,D) = \sum_{j=1}^d \{(p_j - s_j)(x_j - D_j)_+ - (p_j - c_j)x_j\}$. Its expectation is additive in marginal distributions, so demand correlation does not by itself change the oracle solution in this unconstrained problem. Joint fitting and whole-vector resampling preserve dependence in estimation, and APUB aggregates fitted total costs. The experiment fixes one submitted Gaussian-mixture truth and generates thirty independent master training datasets with nested sample sizes.

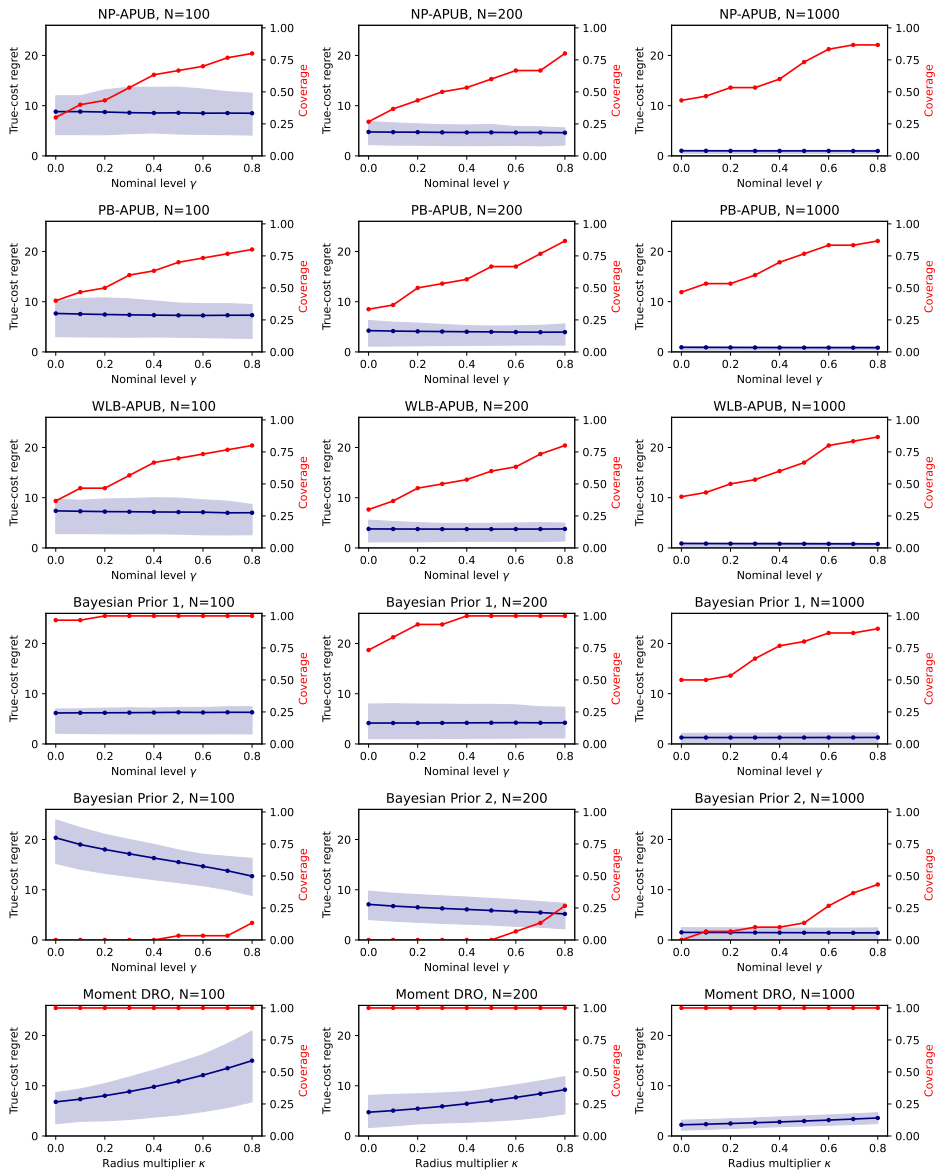

Figure 1: Fixed mixed-normal experiment, 30 paired training datasets. Rows show NP-APUB, PB-APUB, WLB-APUB, Bayesian Prior 1, Bayesian Prior 2, and moment DRO; columns show sample sizes. Bayesian Prior 1 is selected to align with the underlying distribution when good prior knowledge is assumed, whereas Bayesian Prior 2 deliberately represents a misspecified prior. Blue curves are mean true-cost regret and blue regions are the 20th–80th percentile interval across replications. Red curves use the right axis for empirical coverage.

All methods share the dataset within a replication. True parameters are used only to generate data and evaluate final decisions. Expected costs and the population optimum are analytic, avoiding test-sample noise.

SAA and FMO isolate point estimation with and without the assumed family. NP-APUB applies upper aggregation to nonparametric bootstrap sample-average cost functions; PB-APUB and WLB-APUB refit the same parametric family with matched fitting budgets. BAG ($\gamma = 0$) isolates averaging from upper aggregation. Bayesian posterior-CVaR is reported for two prior choices as a sensitivity analysis: Prior 1 aligns with the underlying distribution when good prior knowledge is assumed, whereas Prior 2 deliberately represents misspecified prior knowledge. Bayesian posterior-

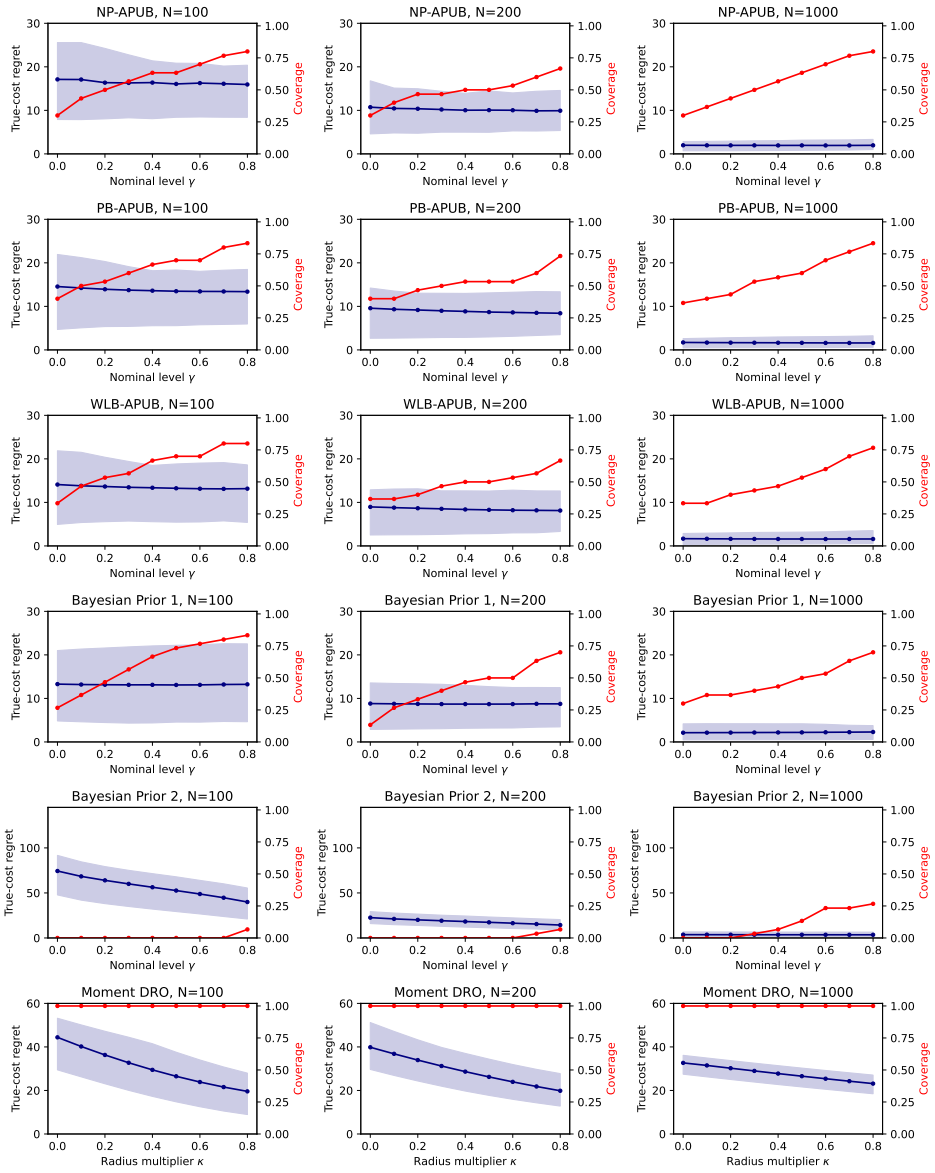

Figure 2: Fixed mixed-log-normal results. Blue curves and bands show mean true-cost regret and its 20th–80th percentile interval; red curves show empirical coverage.

CVaR is a prior-sensitivity comparator, not a complete Bayesian-method benchmark. Moment DRO uses full covariance and a radius-sensitivity grid. The DRO parameter κ is not a confidence level.

For each recommended decision we report true-cost regret $H(\widehat{x}, \theta_0) - v^\star$ and coverage, the indicator that the reported objective can serve as an upper bound for $H(\widehat{x}, \theta_0)$. These are complementary outcomes: regret measures decision quality, while coverage measures the reliability of the reported one-sided cost certificate at its own decision. Figures 1 and 2 display the mean true-cost regret, the 20th–80th percentile interval, and empirical coverage for each method.

The comparison evaluates APUB's epistemic-risk adjustment against SAA, which ignores the parametric family, and FMO, which treats one fitted model as certain. In both fixed-truth cases, PB-APUB and WLB-APUB at positive γ have lower mean regret than both SAA and FMO, while their coverage rises with γ. This advantage narrows at $N = 1000$ as parameter uncertainty contracts,

avoiding an unnecessary large adjustment when information is abundant. Full grids, numerical diagnostics, and ensemble sizes appear in the appendices.

5.1 Mixed-normal results

Figure 1 demonstrates the data-driven risk-aversion adjustment. Parametric PB-APUB and WLB-APUB improve on NP-APUB throughout the displayed grid. Increasing γ protects against unfavorable fitted models, yet WLB mean regret changes only from 7.30 to 7.01 at $N = 100$, 3.77 to 3.78 at $N = 200$, and .87 to .79 at $N = 1000$ when γ rises from .1 to .8. Thus the additional epistemic protection is not accompanied by a systematic increase in mean regret in this experiment. This is distinct from moment DRO, whose radius κ is not a confidence level and whose conservative cost can be large.

Only the APUB rows have approximately linear, level-tracking coverage curves. For WLB, empirical coverage increases from 46.7% to 80.0% at $N = 100$, from 36.7% to 80.0% at $N = 200$, and from 43.3% to 86.7% at $N = 1000$ as γ increases from .1 to .8; at the target $\gamma = .8$, the observed coverage is therefore close to the nominal level. Subject to the stated bootstrap and selected-decision conditions, this makes γ an interpretable intended one-sided confidence level for APUB's reported upper cost, not merely a risk-aversion knob.

Prior choice highlights the advantage of independence from parameter priors. At $N = 100$ and $\gamma = .5$, aligned Bayesian Prior 1 has mean regret 6.30 and 100% coverage, whereas potentially misspecified Prior 2 has mean regret 15.50 and 3.3% coverage. Their difference contracts at $N = 1000$ as the likelihood becomes informative; APUB does not require the analyst to assume that this eventual learning will rescue an initially incorrect prior. The complete mixed-normal design, numerical records, and all-level tables are in Appendix C.

5.2 Mixed-log-normal results

The fixed mixed-log-normal experiment in Figure 2 repeats the same paired design with a log-normal mixture and confirms the main qualitative pattern: FMO improves on SAA, fitted APUB methods improve on NP-APUB, and APUB coverage increases with γ without a large mean-regret change. At $\gamma = .8$, PB/WLB coverage is $83.3\%/80.0\%$ at $N = 100$, $73.3\%/66.7\%$ at $N = 200$, and $83.3\%/76.7\%$ at $N = 1000$, compared with FMO coverage of 40.0%, 40.0%, and 36.7%. The difference from the mixed-normal case is that the 80th-percentile regret need not fall below FMO at every sample size; APUB's role is reliability under parameter uncertainty, not universal improvement of every regret summary.

Prior sensitivity is sharper in the small log-normal samples. At $N = 100$ and $\gamma = .8$, Prior 1 has mean regret 13.23 and 83.3% coverage, whereas Prior 2 has mean regret 39.93 and 6.7% coverage. By $N = 1000$, the corresponding mean regrets are 2.29 and 3.52, consistent with waning prior influence. Moment DRO remains highly conservative. Appendix D gives the complete curves, numerical grid, and retained diagnostic qualifications.

6 Conclusion

Parametric APUB is a decision-and-certification method for parameter uncertainty: it combines a known distribution family with bootstrap upper-percentile aggregation to select a decision and report a nominal one-sided upper bound on that decision's true expected cost. WLB and generative parametric bootstrap are alternative constructions with shared optimization and transfer results but distinct estimation assumptions. APUB's central practical advantage is that it uses the structural distributional model without requiring a parameter prior: its uncertainty adjustment is driven by observed data and fitted-model resampling, so it is not exposed to the small-sample loss caused by choosing a misspecified prior. Under the stated conditions, that bootstrap adjustment contracts with data at a fixed nominal level, yielding consistent decisions without a prior or an ambiguity-radius schedule. The completed studies document level-dependent coverage and regret, not uniform superiority over point estimation, ensemble averaging, or either bootstrap construction. Extending selected-decision coverage to finite ensembles and understanding nonregular fitting remain important theoretical tasks.

REPRODUCIBILITY STATEMENT

The manuscript package includes audited fixed-truth result snapshots, scripts regenerating every displayed table and figure, and the exact mixed-normal and mixed-log-normal configurations. Appendices C and D specify data generation, fitting, priors, solver tolerances, evaluation, and numerical qualifications. The experiment runners and saved fit ensembles are maintained separately. Resumption uses persisted inputs, draw-indexed seeds and configuration identifiers.

AI USE STATEMENT

OpenAI Codex assisted with debugging, result auditing, literature lookup and manuscript editing. Numerical results originate from saved experiment outputs, not generated examples.

REFERENCES

Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Data-driven robust optimization. *Mathematical Programming*, 167(2):235–292, 2018. doi: 10.1007/s10107-017-1125-8.

Sait Cakmak, Di Wu, and Enlu Zhou. Solving bayesian risk optimization via nested stochastic gradient estimation. *arXiv preprint arXiv:2007.05860*, 2020. URL https://arxiv.org/abs/2007.05860.

Anthony C. Davison and David V. Hinkley. *Bootstrap Methods and Their Application*. Cambridge University Press, Cambridge, 1997.

Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010. doi: 10.1287/opre.1090.0741.

Bradley Efron. Better bootstrap confidence intervals. *Journal of the American Statistical Association*, 82(397):171–185, 1987. doi: 10.1080/01621459.1987.10478410.

Bradley Efron. Bayesian inference and the parametric bootstrap. *The Annals of Applied Statistics*, 6(4):1971–1997, 2012. URL https://arxiv.org/abs/1301.2936.

Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall, New York, 1993.

Vishal Gupta. Near-optimal bayesian ambiguity sets for distributionally robust optimization. *Management Science*, 65(9):4242–4260, 2019. doi: 10.1287/mnsc.2018.3140.

Shixin Liu, Ming Gao, and Jian Hu. Minimizing upper confidence bounds: A data-driven framework for stochastic programming. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, 2026. URL https://arxiv.org/abs/2403.08966.

Simon P. Lyddon, Christopher C. Holmes, and Stephen G. Walker. General bayesian updating and the loss-likelihood bootstrap. *Biometrika*, 106(2):465–478, 2019. doi: 10.1093/biomet/asz006.

Pascal Massart. The tight constant in the dvoretzky–kiefer–wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990. doi: 10.1214/aop/1176990746.

Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171:115–166, 2018. doi: 10.1007/s10107-017-1172-1. URL https://arxiv.org/abs/1505.05116.

Michael A. Newton and Adrian E. Raftery. Approximate bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society: Series B*, 56(1):3–26, 1994. doi: 10.1111/j.2517-6161.1994.tb01956.x.

Jens Præstgaard and Jon A. Wellner. Exchangeably weighted bootstraps of the general empirical process. *The Annals of Probability*, 21(4):2053–2086, 1993. URL https://sites.stat.washington.edu/jaw/JAW-papers/jaw-praest-93AP.pdf.

Alexander Shapiro. Distributionally robust stochastic programming. *SIAM Journal on Optimization*, 27(4):2258–2275, 2017. doi: 10.1137/16M1058297.

Alexander Shapiro, Enlu Zhou, and Yifan Lin. Bayesian distributionally robust optimization. *SIAM Journal on Optimization*, 33(2):1279–1304, 2023. doi: 10.1137/21M1465548.

Wolfram Wiesemann, Daniel Kuhn, and Melvyn Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014. doi: 10.1287/opre.2014.1314.

Di Wu, Helin Zhu, and Enlu Zhou. A bayesian risk approach to data-driven stochastic optimization: Formulations and asymptotics. *SIAM Journal on Optimization*, 28(2):1588–1612, 2018. doi: 10.1137/16M1101933.

A COMPLETE ADAPTIVE ALGORITHM

Let $h_b(x) = H(x, \widetilde{\theta}_{N,b})$ and $U_B(x) = U_{N,B,\gamma}(x)$. Fitting, feasible-set projection, and any optimization certification are supplied by the application. The bootstrap increment ΔB is distinct from the gradient minibatch size n_g. The experimental appendices specify application-dependent settings.

A.1 OPTIMIZATION UPDATES AND ENSEMBLE EXPANSION

For a fixed ensemble, a sequence $j_1, \ldots, j_{n_g}$ sampled uniformly with replacement from $\{1, \ldots, B\}$ gives stochastic subgradients for the joint (x, t) objective:

$$\widehat{g}_x = \frac{1}{(1-\gamma)n_g} \sum_{r=1}^{n_g} \mathbf{1}\{h_{j_r}(x) > t\} \nabla h_{j_r}(x), \qquad \widehat{g}_t = 1 - \frac{1}{(1-\gamma)n_g} \sum_{r=1}^{n_g} \mathbf{1}\{h_{j_r}(x) > t\}. \quad (14)$$

For a nonsmooth convex h_b, replace its gradient by a measurable subgradient selection. Choosing zero for the hinge subgradient at equality is valid. If the fitted objectives are convex and their values and subgradients are available exactly, these minibatch estimates are unbiased subgradient selections for the joint finite-ensemble objective in (x, t). After each step, the decision vector is projected onto the feasible set $\mathcal{X}$, while the auxiliary threshold is updated jointly. The step sizes, minibatch size, and iteration budget are algorithm inputs. For nonconvex objectives, the same updates may be used as a heuristic, but the convex convergence guarantees do not apply.

Start with a batch of ΔB conditionally independent fits from the selected WLB or PB generator. Optimize the finite-ensemble APUB objective using the accumulated models. If the bootstrap stability checks are not satisfied after adequate optimization, generate and fit another ΔB models, append them to the ensemble, and warm-start the next solve from the current decision. Previous fits are retained. Thus the total ensemble size after k batches is $B = k\Delta B$, while the number of batches is determined by the stopping rule. A maximum resource budget may be imposed as a practical safeguard, but it does not define the ensemble-growth rule.

Two batch sizes have different roles. The bootstrap increment ΔB determines how many new parameter estimates are fitted when the ensemble grows. The gradient minibatch size n_g determines how many existing models are sampled for each stochastic-gradient update. A gradient update does not require refitting models. Conversely, adding bootstrap models improves the Monte Carlo approximation of the fitted-parameter distribution; it does not by itself solve the current optimization problem. If optimization accuracy is inadequate, continue optimizing the current ensemble before deciding whether more bootstrap fits are needed.

A.2 STABILITY CHECKS AND SAFEGUARDS

The stability checks compare adequately optimized solutions before and after an expansion. Let $B_- = B - \Delta B$, and let (x_-, u_-) be the preceding record, with $u_- = U_{N,B_-,\gamma}(x_-)$. Define $d_x = \|x_B - x_-\|_\infty$ and $d_U = |U_{N,B,\gamma}(x_-) - u_-|$. The latter comparison evaluates both ensembles at the same previous decision, separating ensemble change from decision change. An approximate bootstrap standard error e_B supplies an additional simulation diagnostic. Once the declared optimization criterion meets tolerance ε_{opt}, a comparison is acceptable if $d_x \le \varepsilon_x$, $d_U \le \varepsilon_U$, $e_B \le \varepsilon_{\text{MC}}$, and $B(1-\gamma) \ge m_{\text{tail}}$. Stop after r_{stable} consecutive acceptable comparisons; otherwise add another batch. The first batch supplies the initial record and cannot establish between-ensemble stability. If a resource limit interrupts the process, retain the best available feasible decision together with the achieved accuracy.

An optional refinement solver can improve the stochastic-gradient candidate. In convex applications, a feasible decision and a valid lower bound can certify a finite-ensemble objective gap. Such a certificate requires an application-appropriate lower-bound procedure; an approximate subproblem value is not automatically a lower bound. When no certificate is available, a declared optimization diagnostic may be used, but the result must not be described as certified. This separates the general adaptive framework from any particular numerical solver.

Increasing B reduces simulation error conditional on the data; it does not create information about θ_0. High levels retain a smaller upper tail and can require more fitted models. The optional simultaneous certificate below is distinct from the implemented stability diagnostics above. For multiple

aggregation levels, the failure-probability budget must also be allocated across levels. The guarantee controls bootstrap simulation error; it does not establish coverage of the unknown population objective. Theorem 3 and Appendix B.6 give the bounded-oracle error bound and its proof. Outside those assumptions, the practical stopping rule remains a diagnostic rather than a confidence sequence or a post-selection guarantee.

A.3 Complete pseudocode

Algorithm 2 Incremental adaptive APUB

Require: Data D_N, level γ, feasible set $\mathcal{X}$, generator $m \in \{\mathrm{W}, \mathrm{P}\}$ and fitting rule, bootstrap increment ΔB

Require: Tolerances $\varepsilon_{\mathrm{opt}}, \varepsilon_x, \varepsilon_U, \varepsilon_{\mathrm{MC}}$, minimum tail size m_{tail}, required stable comparisons r_{stable}

Initialize $B \leftarrow 0$, $s \leftarrow 0$, previous record empty, and reproducible random generator
loop
  for $b = B + 1, \ldots, B + \Delta B$ **do**
    if $m = \mathrm{W}$ **then**
      Draw and normalize exponential weights; refit using (4)
    else
      Draw N observations from $P_{\widehat{\theta}_N}$; refit using (3)
    end if
    Cache $\widetilde{\theta}^m_{N,b}$ with its independent generation and fitting seeds
  end for
  $B \leftarrow B + \Delta B$; retain all previously fitted models
  Optimize with Algorithm 3, warm-starting from the previous decision
  while the declared optimization criterion exceeds $\varepsilon_{\mathrm{opt}}$ **do**
    Continue optimization on the current ensemble, without adding bootstrap fits
    If an optimization resource limit is reached, return the best feasible decision together with the achieved optimization accuracy
  end while
  Obtain x_B, $u_B = U_B(x_B)$ and $t_B = \mathrm{quantile}_\gamma\{h_b(x_B) : b \le B\}$
  $v_b \leftarrow t_B + (h_b(x_B) - t_B)_+/(1-\gamma)$ for each b
  $e_B \leftarrow \mathrm{sd}(v_1, \ldots, v_B)/\sqrt{B}$
  if a previous record (B_-, x_-, u_-) exists **then**
    $d_x \leftarrow \|x_B - x_-\|_\infty$; $d_U \leftarrow |U_B(x_-) - u_-|$
    good $\leftarrow [d_x \le \varepsilon_x] \wedge [d_U \le \varepsilon_U]$
    good $\leftarrow$ good $\wedge [e_B \le \varepsilon_{\mathrm{MC}}] \wedge [B(1-\gamma) \ge m_{\mathrm{tail}}]$
  else
    good $\leftarrow$ false
  end if
  $s \leftarrow s + 1$ if good; otherwise $s \leftarrow 0$
  Save decision, objective, diagnostics, fits, and random-generator state
  if $s \ge r_{\mathrm{stable}}$ **then**
    return current decision with diagnostic stopping status
  end if
  If an optional resource limit prevents another batch, return current decision together with the achieved simulation diagnostics
  $(B_-, x_-, u_-) \leftarrow (B, x_B, u_B)$
end loop

The standard-error calculation is conditional on the observed data. Because the decision and threshold are estimated from the ensemble, and stopping involves repeated checks, this diagnostic is not itself a sequential confidence guarantee. Fitting failures must be recorded rather than silently removing unfavorable bootstrap draws. If the optimization diagnostic is not a valid objective-gap bound, satisfying its tolerance is not an optimality certificate.

Algorithm 3 Finite-ensemble stochastic optimization

Require: Fitted-value and subgradient oracles, feasible-set projection $\Pi_{\mathcal{X}}$, ensemble size B, level γ

Require: Initial decision, iteration budget T, minibatch size n_g, step sizes η_j^x, η_j^t, checkpoint interval J_{check}

Initialize feasible x and threshold t at the empirical γ-quantile
Initialize the running decision average and best evaluated feasible decision
for $j = 0, \ldots, T-1$ **do**
    Sample n_g model indices uniformly with replacement
    Compute $\widehat{g}_x, \widehat{g}_t$ using (14)
    $x \leftarrow \Pi_{\mathcal{X}}(x - \eta_j^x \widehat{g}_x)$; $t \leftarrow t - \eta_j^t \widehat{g}_t$
    Update the running average when it is a feasible candidate
    if checkpoint interval reached or final iteration **then**
        Evaluate exact finite-ensemble APUB at feasible candidates and retain the best
        Save iterates, average, best decision, and random-generator state
    end if
end for
Optionally refine the best decision with an application-appropriate solver
Obtain an optimization diagnostic or a valid lower bound when available
return best feasible decision, finite-ensemble objective, and diagnostic or certified gap

For a convex feasible set, averages of feasible iterates remain feasible. For a nonconvex set, candidate feasibility must be checked, projection may be difficult, and no global convergence claim follows from these updates. A lower-bound certificate is valid only when the bound is valid for the original finite-ensemble problem, not merely for a smoothed approximation. Bootstrap weights and fitted models are reused across stages; resumption validates the saved configuration and input identifiers.

A.4 Optional simultaneous Monte Carlo certificate

For the fixed-increment rule let $B_k = k\Delta B$, with ΔB chosen before inspecting bootstrap outcomes. Under Theorem 3's hypotheses, apply (13) at stage k with $\eta_k = \eta/[k(k+1)]$ and a predetermined net resolution h_k. A union bound gives simultaneous conditional validity at every stage with probability at least $1-\eta$, since $\sum_k \eta_k = \eta$. Independence between nested prefixes is unnecessary. Evaluating the bound at a data-dependent stopping stage preserves this simultaneous Monte Carlo statement. For a finite grid of levels, allocate the failure budget across levels as well as stages.

This certificate is not the implemented standard-error or stability rule. It requires valid bounds on the fitted oracle and its Lipschitz constant over the chosen domain. It controls bootstrap simulation error, not coverage of the population objective. Unbounded realized demands do not by themselves settle boundedness of an expected-cost oracle; the parameter and decision domains matter. We have not verified the required uniform bounds for the unconstrained fitted-parameter law in the experiments.

A.5 Resumption and shared ensembles

The saved run configuration, original observations, reference fit, draw-indexed bootstrap seeds and fitted prefixes identify each experiment. Each newly completed fit is reusable across all configured levels. Inner stochastic iterates and random states are checkpointed; an interrupted deterministic refinement may be repeated on resume. A completed output is reused only when its configuration and input identifiers agree. The final B may differ across levels despite their shared prefix; counts or prefix fitting times must not be summed across levels as if they represented disjoint fits. NP-APUB uses the same outer logic, with each generated cost function replacing the corresponding parametric fit.

B TECHNICAL CONDITIONS AND PROOFS

This appendix follows the dependency structure of the theory. We first record the stability property of upper aggregation. We then verify fitted-cost contraction for WLB and PB, use that contraction to prove the decision-transfer theorem, and finally give the fixed-decision, selected-decision, and finite-ensemble results. All probability statements are with respect to repeated training datasets and any stated resampling or fitting randomness.

B.1 UPPER AGGREGATION AS A STABLE OPERATOR

Lemma 1 (Upper-average stability)**.** *For $\gamma \in (0,1)$ and an integrable scalar random variable Y, define*

$$\mathrm{CVaR}_\gamma(Y) = \inf_{t\in\mathbb{R}} \left\{ t + \frac{1}{1-\gamma}\mathbb{E}[Y-t]_+ \right\}.$$

For integrable Y, Z on a common probability space,

$$|\,\mathrm{CVaR}_\gamma(Y) - \mathrm{CVaR}_\gamma(Z)| \le \frac{\mathbb{E}|Y-Z|}{1-\gamma}.$$

If $|Y-Z| \le \varepsilon$ almost surely, the sharper bound is ε. Consequently, for scalar laws P, Q with finite first moments,

$$|\,\mathrm{CVaR}_\gamma(P) - \mathrm{CVaR}_\gamma(Q)| \le \frac{W_1(P,Q)}{1-\gamma},$$

where $\mathrm{CVaR}_\gamma(P)$ means evaluation at any variable with law P. The operator is monotone, translation equivariant, and positively homogeneous; in particular, $\mathrm{CVaR}_\gamma(c) = c$ for a constant c. All statements also hold under a conditional probability law.

The upper-average functional is mathematically the upper-tail conditional value-at-risk operator, denoted CVaR_γ. We use this standard notation for its analytic properties. In APUB, the operator acts on the conditional bootstrap distribution of fitted expected costs, so its statistical role concerns parameter-estimation uncertainty. This identity does not turn APUB into a tail-risk criterion for individual realized costs, nor does it by itself establish a confidence guarantee.

In particular, an oracle that is L-Lipschitz in x uniformly over parameters induces an L-Lipschitz APUB objective. The decision-transfer theorem uses the first bound, while the calibration lemmas below use its Wasserstein consequence.

Proof. For any common threshold t, the 1-Lipschitz positive-part function gives

$$\left| \left(t + \frac{\mathbb{E}[Y-t]_+}{1-\gamma} \right) - \left(t + \frac{\mathbb{E}[Z-t]_+}{1-\gamma} \right) \right| \le \frac{\mathbb{E}|Y-Z|}{1-\gamma}.$$

Taking infima in both directions proves the first bound. Monotonicity follows directly from the positive-part expression; shifting or positively scaling the threshold proves translation equivariance and positive homogeneity. If $|Y-Z| \le \varepsilon$, then $Z-\varepsilon \le Y \le Z+\varepsilon$ almost surely, so $\mathrm{CVaR}_\gamma(Z) - \varepsilon \le \mathrm{CVaR}_\gamma(Y) \le \mathrm{CVaR}_\gamma(Z) + \varepsilon$. Finally, the first bound holds for every coupling of P, Q; taking the infimum of $\mathbb{E}|Y-Z|$ over couplings gives the W_1 bound. Conditioning simply fixes the probability law used throughout. □

B.2 CONSTRUCTION-SPECIFIC FITTED-COST CONTRACTION

The next two proofs establish the WLB and PB propositions stated in the main text. Their conditions are sufficient rather than necessary; they make transparent where the two constructions differ.

Proof of Proposition 1.

Proof. The proof suppresses the superscript W and writes M_N^w for M_N^{W}.

Put $\ell_\theta(\xi) = \log p_\theta(\xi)$ and $F(\xi) = \sup_{\theta\in\Theta} |\ell_\theta(\xi)|$. By (A2), $\mathbb{E}F(\xi) < \infty$. The normalized exponential construction gives

$$M_N^w(\theta) = \frac{A_N(\theta)}{\overline{e}_N}, \qquad A_N(\theta) = \frac{1}{N}\sum_{i=1}^N e_i \ell_\theta(\xi_i), \qquad \overline{e}_N = \frac{1}{N}\sum_{i=1}^N e_i.$$

The pairs (ξ_i, e_i) are iid, $\mathbb{E}e_i = 1$, and $\mathbb{E}[e_i F(\xi_i)] = \mathbb{E}F(\xi_i) < \infty$. Thus the weighted summands have an integrable envelope and are almost surely continuous in θ on the compact set Θ.

For completeness, let $\omega_h(\xi) = \sup_{\|\theta-\vartheta\|\le h} |\ell_\theta(\xi) - \ell_\vartheta(\xi)|$, where both parameters lie in Θ. For each observation outside a fixed null set, continuity on compact Θ is uniform over parameters, not uniformly over all observations. This gives $\omega_h(\xi) \to 0$ as $h \downarrow 0$, and $\omega_h \le 2F$. Dominated convergence therefore gives $\mathbb{E}\omega_h \to 0$. Choose a finite h-net $\{\theta_j\}$ of Θ. The ordinary law of large numbers at its finitely many points, together with

$$\sup_\theta |A_N(\theta) - M(\theta)| \le \max_j |A_N(\theta_j) - M(\theta_j)| + \frac{1}{N}\sum_i e_i \omega_h(\xi_i) + \mathbb{E}\omega_h(\xi),$$

proves $\sup_\theta |A_N(\theta) - M(\theta)| \xrightarrow{\mathbb{P}} 0$ by first taking $N \to \infty$ and then $h \downarrow 0$. Since $\overline{e}_N \xrightarrow{\mathbb{P}} 1$ and $\sup_\theta |M(\theta)| \le \mathbb{E}F$, it follows that

$$T_N := \sup_{\theta\in\Theta} |M_N^w(\theta) - M(\theta)| \le \frac{\sup_\theta |A_N(\theta) - M(\theta)|}{\overline{e}_N} + |\overline{e}_N^{-1} - 1| \sup_\theta |M(\theta)| \xrightarrow{\mathbb{P}} 0.$$

These convergences are under the joint law of data and weights; the argument does not require a particular coupling of the bootstrap weights across sample sizes.

Let $\Theta_0 = \arg\max_{\theta\in\Theta} M(\theta)$. By (A3), this is $\{\theta_0\}$ or its finite label-permutation class. Dominated convergence makes M continuous. For any $\eta > 0$ with a nonempty set of parameters at least η from Θ_0, compactness and identification give

$$c_\eta = M(\theta_0) - \sup_{\operatorname{dist}(\theta,\Theta_0)\ge\eta} M(\theta) > 0.$$

By (A4),

$$0 \le M(\theta_0) - M(\widetilde{\theta}_N) \le 2T_N + R_N \xrightarrow{\mathbb{P}} 0.$$

Consequently, $\mathbb{P}\{\operatorname{dist}(\widetilde{\theta}_N, \Theta_0) \ge \eta\} \le \mathbb{P}\{2T_N + R_N \ge c_\eta\} \to 0$. If the set defining c_η is empty, the first probability is zero.

By (A1) and (A5), H is uniformly continuous and bounded on $\mathcal{X} \times \Theta$. Also $H(x,\theta) = H(x,\theta_0)$ for every $\theta \in \Theta_0$. Hence

$$Z_N := \sup_{x\in\mathcal{X}} |H(x,\widetilde{\theta}_N) - H(x,\theta_0)| \xrightarrow{\mathbb{P}} 0,$$
$$0 \le Z_N \le 2C_H, \qquad C_H := \sup_{(x,\theta)\in\mathcal{X}\times\Theta} |H(x,\theta)| < \infty.$$

For any $u > 0$, $\mathbb{E}Z_N \le u + 2C_H\mathbb{P}(Z_N > u)$, so $\mathbb{E}Z_N \to 0$. Since Q_N is the conditional law of the actual fit, including any fitting randomization, $a_N = \mathbb{E}[Z_N \mid D_N]$. The tower property and Markov's inequality yield, for every $\varepsilon > 0$,

$$\mathbb{P}_{D_N}(a_N > \varepsilon) \le \frac{\mathbb{E}_{D_N} a_N}{\varepsilon} = \frac{\mathbb{E}Z_N}{\varepsilon} \longrightarrow 0.$$

This proves the conditional first-moment contraction, rather than assuming it.

□

Proof of Proposition 2.

Proof. **Step 1: a conditional uniform law.** Take a finite h-net $\{\theta_1, \dots, \theta_{J_h}\}$ of compact Θ. On the event $\widehat{\theta}_N \in V$, conditional Chebyshev and a union bound give, for $u > 0$,

$$\mathbb{P}\left(\max_j |M_N^{\mathrm{P}}(\theta_j) - M_{\widehat{\theta}_N}(\theta_j)| > u \mid D_N\right) \le \frac{J_h C_F}{N u^2}.$$

For every θ, choose a net point within distance h. The triangle inequality gives

$$T_N^{\mathrm{loc}} := \sup_\theta |M_N^{\mathrm{P}}(\theta) - M_{\widehat{\theta}_N}(\theta)|$$
$$\leq \max_j |M_N^{\mathrm{P}}(\theta_j) - M_{\widehat{\theta}_N}(\theta_j)| + \frac{1}{N}\sum_i \omega_h(\xi_i^*) + \mathbb{E}_{P_{\widehat{\theta}_N}}\omega_h.$$

Write $r_h = \sup_{\vartheta \in V} \mathbb{E}_{P_\vartheta}\omega_h$. The conditional expectation of the last two terms is at most $2r_h$. Markov's inequality consequently bounds their probability of exceeding u by $2r_h/u$. For fixed h, the net-point error vanishes as N grows. Then $r_h \to 0$ as $h \downarrow 0$, and (P1) implies $\mathbb{P}(\widehat{\theta}_N \notin V) \to 0$. It follows that $T_N^{\mathrm{loc}} \to 0$ in joint probability.

Step 2: convergence to the population criterion. By (P1)–(P2), $\sup_\theta |M_{\widehat{\theta}_N}(\theta) - M_{\theta_0}(\theta)| \to 0$ in probability. Hence

$$T_N^{\mathrm{P}} := \sup_\theta |M_N^{\mathrm{P}}(\theta) - M_{\theta_0}(\theta)| \to 0.$$

The envelope and continuity in (P2) imply continuity of M_{θ_0} by dominated convergence. For any $\varepsilon > 0$ for which the set is nonempty, compactness and (A3) imply

$$c_\varepsilon = M_{\theta_0}(\theta_0) - \sup_{\mathrm{dist}(\theta,\Theta_0)\geq\varepsilon} M_{\theta_0}(\theta) > 0.$$

The refitting gap gives $0 \leq M_{\theta_0}(\theta_0) - M_{\theta_0}(\widetilde{\theta}_N^{\mathrm{P}}) \leq 2T_N^{\mathrm{P}} + R_N^{\mathrm{P}} \to 0$, so $\mathrm{dist}(\widetilde{\theta}_N^{\mathrm{P}}, \Theta_0) \to 0$ in joint probability.

Step 3: conditional first-moment contraction. Uniform continuity and boundedness of H on $\mathcal{X} \times \Theta$, together with (A5), give $Z_N^{\mathrm{P}} := \sup_x |H(x, \widetilde{\theta}_N^{\mathrm{P}}) - H(x, \theta_0)| \to 0$ in probability and $0 \leq Z_N^{\mathrm{P}} \leq 2C_H$, where $C_H = \sup_{x,\theta} |H(x,\theta)| < \infty$. Thus $\mathbb{E}Z_N^{\mathrm{P}} \to 0$. Finally,

$$\mathbb{P}\{a_N^{\mathrm{P}} > \varepsilon\} \leq \varepsilon^{-1}\mathbb{E}a_N^{\mathrm{P}} = \varepsilon^{-1}\mathbb{E}Z_N^{\mathrm{P}} \to 0,$$

which is the claimed conditional mean contraction. □

B.3 Proof of Theorem 1

Proof. **Step 1: obtain contraction for the selected construction.** Proposition 1 supplies $a_N \to 0$ for WLB; Proposition 2 supplies it for PB under its own assumptions. The remaining argument depends only on this fitted-function error, so it also applies to any other conditional fitting law satisfying the same contraction and the compactness and continuity hypotheses.

Step 2: contraction implies APUB convergence and decision consistency. Condition on D_N. For each x, couple $Y = H(x, \widetilde{\theta}_N)$ to the constant $Z = H(x, \theta_0)$. The conditional stability bound in Lemma 1 and $\mathrm{CVaR}_\gamma(Z) = Z$ imply

$$|U_{N,\gamma}(x) - H(x,\theta_0)| \leq \frac{\mathbb{E}_{Q_N}|H(x,\widetilde{\theta}_N) - H(x,\theta_0)|}{1-\gamma} \leq \frac{a_N}{1-\gamma}.$$

Taking the supremum over x and applying Step 1 proves uniform convergence for every fixed $\gamma < 1$. Set $e_{N,B} = a_N/(1-\gamma) + \varepsilon_{N,B}$, and choose any $x^\star \in \arg\min_{x\in\mathcal{X}} H(x,\theta_0)$, which exists by compactness and continuity. The uniform error bound and δ-optimality give

$$H(\widehat{x},\theta_0) \leq U_{N,B,\gamma}(\widehat{x}) + e_{N,B}$$
$$\leq U_{N,B,\gamma}(x^\star) + \delta + e_{N,B}$$
$$\leq H(x^\star,\theta_0) + 2e_{N,B} + \delta.$$

This is (8). If the simulation and optimization errors vanish in probability, Step 1 implies vanishing regret. For every fixed $\eta > 0$, continuity and compactness give a positive objective gap on the set of feasible decisions at least η from the population argmin set, whenever that set is nonempty. Vanishing regret therefore implies convergence in probability of the distance to the argmin set. No uniqueness of the optimal decision or convexity is required. □

B.4 FIXED-DECISION CALIBRATION

The next lemma is intentionally a fixed-decision statement. It concerns the scalar conditional law of fitted expected cost and can be verified separately for each regular estimator. It does not claim that WLB and PB have the same resampling law, nor does it transfer coverage automatically to a data-selected decision.

Lemma 2 (Fixed-decision normal expansion and coverage)**.** *Fix a deterministic x and $\gamma \in (0,1)$. Write $\widehat{h}_N = H(x, \widehat{\theta}_N)$ and suppose, for $s(x) > 0$,*

$$\sqrt{N}\{\widehat{h}_N - H(x,\theta_0)\} \Rightarrow \mathcal{N}(0, s(x)^2),$$
$$W_1\Big(\mathcal{L}\Big(\sqrt{N}[H(x,\widetilde{\theta}_N) - \widehat{h}_N] \mid D_N\Big), \mathcal{N}(0, s(x)^2)\Big) \xrightarrow{\mathbb{P}} 0.$$

For $z_\gamma = \Phi^{-1}(\gamma)$ and $\kappa_\gamma = \phi(z_\gamma)/(1-\gamma)$,

$$U_{N,\gamma}(x) = \widehat{h}_N + \frac{\kappa_\gamma s(x)}{\sqrt{N}} + o_{\mathbb{P}}(N^{-1/2}), \qquad \mathbb{P}\{H(x,\theta_0) \le U_{N,\gamma}(x)\} \to \Phi(\kappa_\gamma) > \gamma. \tag{15}$$

Proof. Fix a deterministic decision x and a level $\gamma \in (0,1)$, neither depending on N or on the observed data. To simplify notation, write

$$h_0 = H(x,\theta_0), \qquad \widehat{h}_N = H(x,\widehat{\theta}_N), \qquad s = s(x) > 0,$$

and define the centered bootstrap fluctuation and its conditional law by

$$Z_N^* = \sqrt{N}\{H(x,\widetilde{\theta}_N) - \widehat{h}_N\}, \qquad P_N^* = \mathcal{L}(Z_N^* \mid D_N).$$

The law P_N^* is random because it depends on D_N. Let P_s be the law of sZ, where $Z \sim \mathcal{N}(0,1)$. The second assumption of the theorem states that $\delta_N := W_1(P_N^*, P_s) \xrightarrow{\mathbb{P}} 0$ under repeated sampling of the original data. All conditional expectations below integrate over the bootstrap law with the data held fixed.

Step 1: separate the fitted value from its bootstrap adjustment. In the definition of $U_{N,\gamma}(x)$, make the bijective change of threshold $t = \widehat{h}_N + v/\sqrt{N}$. Since $[H(x,\widetilde{\theta}_N) - t]_+ = N^{-1/2}[Z_N^* - v]_+$, we obtain the exact identity

$$\begin{aligned} U_{N,\gamma}(x) &= \widehat{h}_N + \frac{1}{\sqrt{N}} \inf_{v\in\mathbb{R}} \left\{ v + \frac{1}{1-\gamma}\mathbb{E}\big[[Z_N^* - v]_+ \mid D_N\big] \right\} \\ &= \widehat{h}_N + \frac{1}{\sqrt{N}} \operatorname{CVaR}_\gamma(P_N^*). \end{aligned}$$

Thus it suffices to determine the limit of the conditional upper-average term. Lemma 1 gives the pathwise bound

$$\big|\operatorname{CVaR}_\gamma(P_N^*) - \operatorname{CVaR}_\gamma(P_s)\big| \le \frac{\delta_N}{1-\gamma} = o_{\mathbb{P}}(1).$$

Here the level is fixed, so $(1-\gamma)^{-1}$ is a constant.

Step 2: evaluate the normal upper-average term. For a standard normal variable define

$$f_\gamma(v) = v + \frac{1}{1-\gamma}\mathbb{E}[Z - v]_+.$$

Differentiation under the expectation is justified by the bounded difference quotients of the positive-part function and $\mathbb{P}(Z = v) = 0$. Therefore

$$f_\gamma'(v) = 1 - \frac{1-\Phi(v)}{1-\gamma} = \frac{\Phi(v) - \gamma}{1-\gamma}.$$

This derivative is negative for $v < z_\gamma$ and positive for $v > z_\gamma$, where $z_\gamma = \Phi^{-1}(\gamma)$. Hence z_γ is the unique minimizing threshold. Substitution gives

$$\begin{aligned} \operatorname{CVaR}_\gamma(Z) &= z_\gamma + \frac{1}{1-\gamma}\int_{z_\gamma}^{\infty} (z - z_\gamma)\phi(z)\,dz \\ &= \frac{1}{1-\gamma}\int_{z_\gamma}^{\infty} z\phi(z)\,dz = \frac{\phi(z_\gamma)}{1-\gamma} = \kappa_\gamma. \end{aligned}$$

The second equality uses $\int_{z_\gamma}^{\infty} \phi(z)\,dz = 1-\gamma$; the third uses $\phi'(z) = -z\phi(z)$ and $\phi(z) \to 0$ as $z \to \infty$. Positive homogeneity then gives $\mathrm{CVaR}_\gamma(P_s) = s\kappa_\gamma$.

Step 3: obtain the first-order APUB expansion. Combining Steps 1 and 2 yields

$$U_{N,\gamma}(x) = \widehat{h}_N + \frac{s\kappa_\gamma}{\sqrt{N}} + R_N(x), \qquad |R_N(x)| \le \frac{\delta_N}{(1-\gamma)\sqrt{N}}.$$

Consequently $\sqrt{N}R_N(x) \xrightarrow{\mathbb{P}} 0$, proving the expansion in (15). In particular, the adjustment is of order $N^{-1/2}$ for a fixed level and fixed s; it measures estimation uncertainty rather than the variability of individual outcomes.

Step 4: translate the expansion into sampling coverage. By the first assumption of the lemma,

$$T_N := \frac{\sqrt{N}}{s}(\widehat{h}_N - h_0) \Rightarrow Z.$$

The expansion implies

$$A_N := \frac{\sqrt{N}}{s}\{U_{N,\gamma}(x) - h_0\} = T_N + \kappa_\gamma + o_{\mathbb{P}}(1) \Rightarrow Z + \kappa_\gamma$$

by Slutsky's theorem. Independence between the fitted value and its bootstrap adjustment is not required: the remainder converges in probability to zero. Because the limiting distribution has no atom at zero,

$$\begin{aligned}\mathbb{P}\{h_0 \le U_{N,\gamma}(x)\} &= \mathbb{P}\{A_N \ge 0\} \\ &\longrightarrow \mathbb{P}\{Z + \kappa_\gamma \ge 0\} = 1 - \Phi(-\kappa_\gamma) = \Phi(\kappa_\gamma).\end{aligned}$$

This probability is over repeated training datasets. The conditional bootstrap law constructs the bound; the sampling law determines its frequentist coverage.

Step 5: establish strict conservatism relative to the percentile level. Using the integral representation from Step 2,

$$\kappa_\gamma - z_\gamma = \frac{1}{1-\gamma}\int_{z_\gamma}^{\infty} (z - z_\gamma)\phi(z)\,dz > 0.$$

Since Φ is strictly increasing, $\Phi(\kappa_\gamma) > \Phi(z_\gamma) = \gamma$. □

B.5 Selected-decision coverage

The proof below is stronger than the fixed-decision lemma because it must control both the fitted-cost approximation and the data-dependent movement of the optimizer.

Proof. Write $x_N = x_{N,\gamma}$ and $s_\star = s(x^\star)$. By the unique-optimum assumption and Theorem 1, $x_N \xrightarrow{\mathbb{P}} x^\star$; hence $x_N \in \mathcal{V}$ with probability tending to one. For $x \in \mathcal{V}$, let

$$P^*_{N,x} = \mathcal{L}\Big(\sqrt{N}[H(x,\widetilde{\theta}_N) - H(x,\widehat{\theta}_N)] \mid D_N\Big), \qquad \Delta_N = \sup_{x\in\mathcal{V}} W_1\big(P^*_{N,x}, \mathcal{N}(0, s(x)^2)\big).$$

The exact threshold change used in the proof of Lemma 2 gives, simultaneously for every $x \in \mathcal{V}$,

$$\sqrt{N}\{U_{N,\gamma}(x) - H(x,\widehat{\theta}_N)\} = \mathrm{CVaR}_\gamma(P^*_{N,x}).$$

The W_1 stability bound for the upper-average functional and positive homogeneity of the normal upper average therefore imply

$$\sup_{x\in\mathcal{V}} \Big|\sqrt{N}\{U_{N,\gamma}(x) - H(x,\widehat{\theta}_N)\} - \kappa_\gamma s(x)\Big| \le \frac{\Delta_N}{1-\gamma} = o_{\mathbb{P}}(1).$$

Evaluating (B.5) at x_N is valid on the event $\{x_N \in \mathcal{V}\}$ and hence in probability.

Condition (10) and $x_N \xrightarrow{\mathbb{P}} x^\star$ yield

$$G_N(x_N) - G_N(x^\star) = o_{\mathbb{P}}(1).$$

Continuity of s likewise gives $s(x_N) - s_\star = o_{\mathbb{P}}(1)$. Combining these facts with (B.5) gives

$$\begin{aligned}\frac{\sqrt{N}}{s_\star}\{U_{N,\gamma}(x_N) - H(x_N, \theta_0)\} &= \frac{G_N(x_N)}{s_\star} + \kappa_\gamma \frac{s(x_N)}{s_\star} + o_{\mathbb{P}}(1) \\ &= \frac{G_N(x^\star)}{s_\star} + \kappa_\gamma + o_{\mathbb{P}}(1) \Rightarrow \mathcal{N}(\kappa_\gamma, 1),\end{aligned}$$

where the last step uses (9). The normal limit has no atom at zero, so the portmanteau theorem yields

$$\mathbb{P}\{H(x_N, \theta_0) \le U_{N,\gamma}(x_N)\} \longrightarrow \mathbb{P}\{Z + \kappa_\gamma \ge 0\} = \Phi(\kappa_\gamma).$$

Finally, $\kappa_\gamma > \Phi^{-1}(\gamma)$ by the final step of the proof of Lemma 2; thus $\Phi(\kappa_\gamma) > \gamma$. □

B.6 Proof of Theorem 3

The theorem isolates the Monte Carlo approximation of the ideal upper-average objective. It is not a population-coverage statement.

Proof. Throughout the proof, condition on the observed dataset D_N. The fitted-parameter law Q_N is therefore fixed, and probability and expectation refer only to the bootstrap draws. Fix $B \ge 1$, $\gamma \in (0,1)$, $\eta \in (0,1)$, and a finite h-net $\{x_1, \dots, x_{K_h}\} \subset \mathcal{X}$: for every $x \in \mathcal{X}$ there is an index j with $\|x - x_j\| \le h$, using the norm in the Lipschitz assumption. The net is chosen without inspecting the bootstrap draws. Write

$$F_x(z) = Q_N\{\theta : H(x,\theta) \le z\}, \qquad \widehat{F}_{B,x}(z) = \frac{1}{B}\sum_{b=1}^{B} \mathbf{1}\{H(x, \widetilde{\theta}_{N,b}) \le z\}.$$

Let P_x and $\widehat{P}_{B,x}$ denote the scalar laws with these CDFs. By definition,

$$U_{N,\gamma}(x) = \mathrm{CVaR}_\gamma(P_x), \qquad U_{N,B,\gamma}(x) = \mathrm{CVaR}_\gamma(\widehat{P}_{B,x}).$$

Step 1: simultaneous CDF control on the finite net. For each fixed x_j, the values $H(x_j, \widetilde{\theta}_{N,b})$, $b = 1, \dots, B$, are iid under the conditional law, with CDF F_{x_j}. The Dvoretzky–Kiefer–Wolfowitz inequality with the sharp constant (Massart, 1990) gives, for every $u > 0$,

$$\mathbb{P}\left\{\sup_{z\in\mathbb{R}} |\widehat{F}_{B,x_j}(z) - F_{x_j}(z)| > u \,\middle|\, D_N\right\} \le 2e^{-2Bu^2}.$$

Take a union bound over the K_h net points. The resulting inequality is

$$\mathbb{P}\left\{\max_{1\le j\le K_h} \sup_{z\in\mathbb{R}} |\widehat{F}_{B,x_j}(z) - F_{x_j}(z)| > u \,\middle|\, D_N\right\} \le 2K_h e^{-2Bu^2}.$$

The same bootstrap draws are used at all net points, so these CDF errors need not be independent; the union bound does not require independence across j. Set

$$u_B = \sqrt{\frac{\log(2K_h/\eta)}{2B}}, \qquad \mathcal{E}_B = \left\{\max_j \sup_z |\widehat{F}_{B,x_j}(z) - F_{x_j}(z)| \le u_B\right\}.$$

Then $\mathbb{P}(\mathcal{E}_B \mid D_N) \ge 1 - \eta$.

Step 2: convert CDF control to first-moment Wasserstein control. The bounded-oracle assumption implies that P_{x_j} and $\widehat{P}_{B,x_j}$ are both supported on $[a,b]$, almost surely. For scalar laws with finite first moments, Wasserstein distance equals the integral of the absolute CDF difference. Hence, on $\mathcal{E}_B$, simultaneously for all j,

$$\begin{aligned}W_1(\widehat{P}_{B,x_j}, P_{x_j}) &= \int_{\mathbb{R}} |\widehat{F}_{B,x_j}(z) - F_{x_j}(z)|\, dz \\ &= \int_a^b |\widehat{F}_{B,x_j}(z) - F_{x_j}(z)|\, dz \le (b-a)u_B.\end{aligned}$$

The integral outside $[a, b]$ vanishes because both CDFs are zero below a and one above b.

Step 3: bound APUB error at every net point. Apply the Wasserstein stability inequality in Lemma 1 to the two scalar laws. On $\mathcal{E}_B$,

$$|U_{N,B,\gamma}(x_j) - U_{N,\gamma}(x_j)| = |\operatorname{CVaR}_\gamma(\widehat{P}_{B,x_j}) - \operatorname{CVaR}_\gamma(P_{x_j})| \\ \leq \frac{W_1(\widehat{P}_{B,x_j}, P_{x_j})}{1-\gamma} \leq \frac{b-a}{1-\gamma} u_B \qquad (1 \leq j \leq K_h).$$

Step 4: extend the bound to all feasible decisions. For any $x, y \in \mathcal{X}$, the uniform Lipschitz assumption gives

$$|H(x,\theta) - H(y,\theta)| \leq L\|x-y\| \quad \text{for } Q_N\text{-almost every } \theta.$$

Couple the two fitted costs by using the same parameter draw. The almost-sure stability bound in Lemma 1 then yields

$$|U_{N,\gamma}(x) - U_{N,\gamma}(y)| \leq L\|x-y\|.$$

The same argument under the empirical parameter law, using the same bootstrap index for both decisions, gives almost surely

$$|U_{N,B,\gamma}(x) - U_{N,B,\gamma}(y)| \leq L\|x-y\|.$$

Here the sharper almost-sure stability bound is used, so these Lipschitz constants have no extra factor $(1-\gamma)^{-1}$. For each $x \in \mathcal{X}$, choose a net point x_j with $\|x - x_j\| \leq h$. On $\mathcal{E}_B$, the triangle inequality and Step 3 imply

$$\begin{aligned} |U_{N,B,\gamma}(x) - U_{N,\gamma}(x)| &\leq |U_{N,B,\gamma}(x) - U_{N,B,\gamma}(x_j)| \\ &\quad + |U_{N,B,\gamma}(x_j) - U_{N,\gamma}(x_j)| \\ &\quad + |U_{N,\gamma}(x_j) - U_{N,\gamma}(x)| \\ &\leq Lh + \frac{b-a}{1-\gamma} u_B + Lh. \end{aligned}$$

Since the right-hand side is independent of x, taking the supremum over $\mathcal{X}$ and substituting the expression for u_B proves (13) on $\mathcal{E}_B$. Step 1 gives the claimed conditional probability of at least $1-\eta$. □

B.7 Additional properties

Proposition 3 (Convexity and model reweighting)**.** *If the fitted functions h_b are finite and convex on convex $\mathcal{X}$, then $U_{N,B,\gamma}$ is convex and*

$$U_{N,B,\gamma}(x) = \max_{w \in \mathcal{Q}_{B,\gamma}} \sum_b w_b h_b(x), \quad \mathcal{Q}_{B,\gamma} = \left\{ w \geq 0 : \sum_b w_b = 1,\ w_b \leq \frac{1}{(1-\gamma)B} \right\}. \tag{16}$$

For a common ensemble and fixed decision, this value is nondecreasing in γ.

Proof. Introduce variables $u_b \geq 0$ with $u_b \geq h_b(x) - t$ in (6). For fixed x this is a feasible finite linear program, bounded below. Multipliers $w_b \geq 0$ on the second inequalities give the Lagrangian

$$t\Big(1 - \sum_b w_b\Big) + \sum_b u_b \Big(\frac{1}{(1-\gamma)B} - w_b\Big) + \sum_b w_b h_b(x).$$

Minimizing over unrestricted t requires $\sum_b w_b = 1$, while minimizing over $u_b \geq 0$ requires $w_b \leq [(1-\gamma)B]^{-1}$. Together with $w_b \geq 0$, these conditions define $\mathcal{Q}_{B,\gamma}$ in (16). Linear-programming duality proves (16). Convexity follows as a supremum of nonnegative weighted sums of convex functions. Increasing γ decreases $(1-\gamma)$ and relaxes the upper bounds on w, proving monotonicity for a common ensemble. □

C FIXED-TRUTH MIXED-NORMAL EXPERIMENT

C.1 DESIGN AND COMPUTATION

The completed mixed-normal study uses one fixed 20-product, three-component Gaussian-mixture population. It is not a randomized-environment experiment. We generate 30 independent master training datasets, each with 1000 joint observations, using the declared master seed 20260910. Within replication r, the datasets at $N = 100$ and $N = 200$ are saved prefixes of the $N = 1000$ dataset. Every method receives the same $D_{r,N}$ at a given replication and sample size. The 30 replication indices are the independent experimental units.

All products use selling price $p = 10$, purchase cost $c = 5$, and salvage $s = 1$, giving

$$\ell(x, D) = \sum_{j=1}^{20} \{9(x_j - D_j)_+ - 5x_j\}.$$

The fixed mixture and its component count are unavailable to fitting and optimization. They are used only to generate the training data and to evaluate each completed recommendation analytically. For a Gaussian marginal, $\mathbb{E}(x - D)_+ = (x - \mu)\Phi(z) + \sigma\phi(z)$, where $z = (x - \mu)/\sigma$; mixture costs and derivatives are the corresponding mixture-weighted sums. Thus the population optimum is the coordinatewise nonnegative $5/9$ mixture quantile and no finite test sample is used in evaluation.

SAA minimizes empirical average cost. FMO, PB-APUB, and WLB-APUB use the known three-component Gaussian-mixture family. Original-data EM uses ten starts; bootstrap refits use three starts, covariance-eigenvalue floor .25, relative likelihood tolerance 10^{-8}, and initial maximum 1000 iterations. The weighted-bootstrap continuation budget is 4000 iterations. PB generates an N-point joint sample from the original fitted mixture; WLB uses normalized independent exponential likelihood weights; and NP-APUB resamples complete observed demand vectors. At $\gamma = 0$, APUB is its matched bootstrap-average companion (BAG), which averages fitted cost functions rather than mixture parameters. APUB and posterior-CVaR use $\gamma \in \{0, .1, \ldots, .8\}$; moment DRO uses $\kappa \in \{0, .1, \ldots, .8\}$.

Bootstrap ensembles grow in increments of 500 to a maximum of 10000. The hybrid solver uses minibatches of 128, 1200 stochastic steps, deterministic refinement, and finite-ensemble gap tolerance .01. Operational stopping requires tail size at least 100, two stable comparisons, decision change at most .25, preceding-decision objective change at most 2, and estimated Monte Carlo standard error at most 3. These are finite-computation diagnostics, not a sequential confidence guarantee. The Bayesian comparison uses four chains, 1000 warmup sweeps, stages of 500 retained draws per chain, and at most 4000 retained draws per chain. Prior 1 aligns with the underlying distribution when good prior knowledge is available; Prior 2 is potentially misspecified. Moment DRO uses full empirical covariance and retains moment and shape ambiguity at $\kappa = 0$.

C.2 COMPLETE RESULTS AND QUALIFICATIONS

All 30 replications are complete for every sample size, method, and configured level. The audit retains 5040 saved decisions and verifies nested training prefixes, common fixed-population analytic evaluation, and the absence of discarded completed replications. Tables 1–3 report mean regret, the 20th and 80th percentiles across the 30 replication-specific regrets, and empirical coverage at every level. The percentile values describe repeated-training-sample dispersion, not confidence intervals for the mean. Figure 1 in the main body displays the full curves.

Coverage is the fraction of reported objectives that exceed analytic true expected cost at the corresponding data-selected recommendation. It is descriptive evidence about the reported objective at that recommendation, not a uniform confidence guarantee for the optimizer, regret, or realized losses. All completed bootstrap outputs satisfy the declared operational stopping criteria before the 10000-model cap. Shared ensemble prefixes are reused across levels, so their model counts must not be summed as independent computational work.

C.3 COMPLETE NUMERICAL RESULTS

Columns separately report mean true-cost regret, the 20th and 80th percentiles of regret, and empirical coverage (%). The percentiles and coverage are across the 30 paired training datasets. For APUB and Bayesian posterior-CVaR, the level is γ; for moment DRO, it is the radius multiplier κ. At $\gamma = 0$, APUB is BAG.

Table 1: Complete fixed mixed-normal results for $N = 100$.

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
NP-APUB	$\gamma = 0.0$	8.81	4.23	11.96	30.0
NP-APUB	$\gamma = 0.1$	8.81	4.23	11.96	40.0
NP-APUB	$\gamma = 0.2$	8.73	4.21	13.14	43.3
NP-APUB	$\gamma = 0.3$	8.61	4.41	13.67	53.3
NP-APUB	$\gamma = 0.4$	8.54	4.57	13.65	63.3
NP-APUB	$\gamma = 0.5$	8.58	4.35	13.69	66.7
NP-APUB	$\gamma = 0.6$	8.50	4.28	13.30	70.0
NP-APUB	$\gamma = 0.7$	8.52	4.23	12.78	76.7
NP-APUB	$\gamma = 0.8$	8.49	4.14	12.49	80.0
PB-APUB	$\gamma = 0.0$	7.66	3.04	10.13	40.0
PB-APUB	$\gamma = 0.1$	7.54	2.98	10.65	46.7
PB-APUB	$\gamma = 0.2$	7.45	2.96	10.76	50.0
PB-APUB	$\gamma = 0.3$	7.38	2.93	10.57	60.0
PB-APUB	$\gamma = 0.4$	7.32	2.98	10.16	63.3
PB-APUB	$\gamma = 0.5$	7.28	2.92	9.73	70.0
PB-APUB	$\gamma = 0.6$	7.26	2.85	9.60	73.3
PB-APUB	$\gamma = 0.7$	7.30	2.77	9.60	76.7
PB-APUB	$\gamma = 0.8$	7.30	2.70	9.44	80.0
WLB-APUB	$\gamma = 0.0$	7.38	2.85	9.70	36.7
WLB-APUB	$\gamma = 0.1$	7.30	2.85	9.49	46.7
WLB-APUB	$\gamma = 0.2$	7.23	2.82	9.74	46.7
WLB-APUB	$\gamma = 0.3$	7.20	2.78	9.83	56.7
WLB-APUB	$\gamma = 0.4$	7.16	2.81	9.99	66.7
WLB-APUB	$\gamma = 0.5$	7.14	2.75	9.93	70.0
WLB-APUB	$\gamma = 0.6$	7.12	2.61	9.58	73.3
WLB-APUB	$\gamma = 0.7$	6.99	2.60	9.31	76.7
WLB-APUB	$\gamma = 0.8$	7.01	2.64	8.65	80.0
Bayesian Prior 1	$\gamma = 0.0$	6.16	2.15	6.96	96.7
Bayesian Prior 1	$\gamma = 0.1$	6.20	2.08	6.99	96.7
Bayesian Prior 1	$\gamma = 0.2$	6.20	2.04	7.13	100.0
Bayesian Prior 1	$\gamma = 0.3$	6.23	2.02	7.19	100.0
Bayesian Prior 1	$\gamma = 0.4$	6.25	2.00	7.15	100.0
Bayesian Prior 1	$\gamma = 0.5$	6.30	2.00	7.26	100.0
Bayesian Prior 1	$\gamma = 0.6$	6.27	2.00	7.27	100.0
Bayesian Prior 1	$\gamma = 0.7$	6.30	2.03	7.40	100.0
Bayesian Prior 1	$\gamma = 0.8$	6.30	1.99	7.41	100.0
Bayesian Prior 2	$\gamma = 0.0$	20.34	15.24	23.92	0.0
Bayesian Prior 2	$\gamma = 0.1$	19.00	14.08	22.37	0.0
Bayesian Prior 2	$\gamma = 0.2$	18.02	13.27	21.03	0.0
Bayesian Prior 2	$\gamma = 0.3$	17.15	12.63	20.00	0.0
Bayesian Prior 2	$\gamma = 0.4$	16.31	12.01	19.04	0.0
Bayesian Prior 2	$\gamma = 0.5$	15.50	11.42	17.92	3.3
Bayesian Prior 2	$\gamma = 0.6$	14.67	10.81	17.05	3.3
Bayesian Prior 2	$\gamma = 0.7$	13.77	10.00	16.63	3.3
Bayesian Prior 2	$\gamma = 0.8$	12.71	8.86	16.25	13.3
Moment DRO	$\kappa = 0.0$	6.79	2.44	8.63	100.0
Moment DRO	$\kappa = 0.1$	7.33	2.93	9.35	100.0
Moment DRO	$\kappa = 0.2$	8.02	3.03	10.42	100.0

Table 1 continued

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
Moment DRO	$\kappa = 0.3$	8.83	3.29	11.73	100.0
Moment DRO	$\kappa = 0.4$	9.79	3.75	13.13	100.0
Moment DRO	$\kappa = 0.5$	10.88	4.19	14.62	100.0
Moment DRO	$\kappa = 0.6$	12.12	4.85	16.18	100.0
Moment DRO	$\kappa = 0.7$	13.49	5.67	18.37	100.0
Moment DRO	$\kappa = 0.8$	15.00	6.78	20.89	100.0
SAA	—	8.81	4.23	11.96	30.0
FMO	—	7.51	2.84	10.18	40.0

Table 2: Complete fixed mixed-normal results for $N = 200$.

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
NP-APUB	$\gamma = 0.0$	4.77	2.25	6.85	26.7
NP-APUB	$\gamma = 0.1$	4.73	2.15	6.58	36.7
NP-APUB	$\gamma = 0.2$	4.72	2.14	6.38	43.3
NP-APUB	$\gamma = 0.3$	4.68	2.06	6.20	50.0
NP-APUB	$\gamma = 0.4$	4.66	2.03	6.16	53.3
NP-APUB	$\gamma = 0.5$	4.67	2.05	6.24	60.0
NP-APUB	$\gamma = 0.6$	4.65	2.09	5.84	66.7
NP-APUB	$\gamma = 0.7$	4.66	2.00	5.78	66.7
NP-APUB	$\gamma = 0.8$	4.62	2.15	5.57	80.0
PB-APUB	$\gamma = 0.0$	4.23	1.14	6.26	33.3
PB-APUB	$\gamma = 0.1$	4.15	1.18	5.92	36.7
PB-APUB	$\gamma = 0.2$	4.09	1.22	5.74	50.0
PB-APUB	$\gamma = 0.3$	4.05	1.27	5.47	53.3
PB-APUB	$\gamma = 0.4$	4.01	1.29	5.24	56.7
PB-APUB	$\gamma = 0.5$	3.99	1.34	5.16	66.7
PB-APUB	$\gamma = 0.6$	3.94	1.36	5.19	66.7
PB-APUB	$\gamma = 0.7$	3.92	1.37	5.28	76.7
PB-APUB	$\gamma = 0.8$	3.94	1.36	5.58	86.7
WLB-APUB	$\gamma = 0.0$	3.78	1.19	5.53	30.0
WLB-APUB	$\gamma = 0.1$	3.77	1.21	5.26	36.7
WLB-APUB	$\gamma = 0.2$	3.76	1.21	5.05	46.7
WLB-APUB	$\gamma = 0.3$	3.75	1.28	4.89	50.0
WLB-APUB	$\gamma = 0.4$	3.74	1.29	4.87	53.3
WLB-APUB	$\gamma = 0.5$	3.74	1.32	4.91	60.0
WLB-APUB	$\gamma = 0.6$	3.75	1.28	4.96	63.3
WLB-APUB	$\gamma = 0.7$	3.75	1.27	5.04	73.3
WLB-APUB	$\gamma = 0.8$	3.78	1.39	4.93	80.0
Bayesian Prior 1	$\gamma = 0.0$	4.16	1.06	7.94	73.3
Bayesian Prior 1	$\gamma = 0.1$	4.17	1.04	8.00	83.3
Bayesian Prior 1	$\gamma = 0.2$	4.17	1.06	7.94	93.3
Bayesian Prior 1	$\gamma = 0.3$	4.19	1.09	7.92	93.3
Bayesian Prior 1	$\gamma = 0.4$	4.20	1.12	7.86	100.0
Bayesian Prior 1	$\gamma = 0.5$	4.22	1.14	7.89	100.0
Bayesian Prior 1	$\gamma = 0.6$	4.24	1.16	7.79	100.0
Bayesian Prior 1	$\gamma = 0.7$	4.21	1.21	7.40	100.0
Bayesian Prior 1	$\gamma = 0.8$	4.22	1.18	7.32	100.0
Bayesian Prior 2	$\gamma = 0.0$	7.09	4.07	9.72	0.0
Bayesian Prior 2	$\gamma = 0.1$	6.75	3.76	9.30	0.0
Bayesian Prior 2	$\gamma = 0.2$	6.50	3.51	9.02	0.0
Bayesian Prior 2	$\gamma = 0.3$	6.28	3.30	8.78	0.0
Bayesian Prior 2	$\gamma = 0.4$	6.08	3.13	8.54	0.0
Bayesian Prior 2	$\gamma = 0.5$	5.87	2.98	8.26	0.0

Table 2 continued

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
Bayesian Prior 2	$\gamma = 0.6$	5.66	2.72	7.99	6.7
Bayesian Prior 2	$\gamma = 0.7$	5.46	2.55	7.63	13.3
Bayesian Prior 2	$\gamma = 0.8$	5.19	2.21	7.31	26.7
Moment DRO	$\kappa = 0.0$	4.73	1.68	8.06	100.0
Moment DRO	$\kappa = 0.1$	5.05	2.01	8.28	100.0
Moment DRO	$\kappa = 0.2$	5.44	2.40	8.40	100.0
Moment DRO	$\kappa = 0.3$	5.90	2.58	8.57	100.0
Moment DRO	$\kappa = 0.4$	6.42	2.73	8.82	100.0
Moment DRO	$\kappa = 0.5$	7.02	2.95	9.46	100.0
Moment DRO	$\kappa = 0.6$	7.68	3.25	10.18	100.0
Moment DRO	$\kappa = 0.7$	8.41	3.73	10.96	100.0
Moment DRO	$\kappa = 0.8$	9.22	4.41	11.82	100.0
SAA	—	4.70	2.23	6.82	26.7
FMO	—	4.23	1.20	6.34	33.3

Table 3: Complete fixed mixed-normal results for $N = 1000$.

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
NP-APUB	$\gamma = 0.0$	1.02	0.50	1.13	43.3
NP-APUB	$\gamma = 0.1$	1.01	0.49	1.13	46.7
NP-APUB	$\gamma = 0.2$	1.00	0.49	1.13	53.3
NP-APUB	$\gamma = 0.3$	0.99	0.49	1.13	53.3
NP-APUB	$\gamma = 0.4$	0.99	0.48	1.12	60.0
NP-APUB	$\gamma = 0.5$	0.99	0.48	1.11	73.3
NP-APUB	$\gamma = 0.6$	0.99	0.50	1.10	83.3
NP-APUB	$\gamma = 0.7$	0.98	0.50	1.09	86.7
NP-APUB	$\gamma = 0.8$	0.98	0.51	1.11	86.7
PB-APUB	$\gamma = 0.0$	0.92	0.34	1.08	46.7
PB-APUB	$\gamma = 0.1$	0.90	0.36	1.05	53.3
PB-APUB	$\gamma = 0.2$	0.89	0.37	1.02	53.3
PB-APUB	$\gamma = 0.3$	0.87	0.37	1.01	60.0
PB-APUB	$\gamma = 0.4$	0.86	0.36	1.01	70.0
PB-APUB	$\gamma = 0.5$	0.86	0.36	0.99	76.7
PB-APUB	$\gamma = 0.6$	0.85	0.37	1.02	83.3
PB-APUB	$\gamma = 0.7$	0.85	0.39	1.01	83.3
PB-APUB	$\gamma = 0.8$	0.84	0.38	1.03	86.7
WLB-APUB	$\gamma = 0.0$	0.88	0.31	1.02	40.0
WLB-APUB	$\gamma = 0.1$	0.87	0.32	1.02	43.3
WLB-APUB	$\gamma = 0.2$	0.85	0.32	1.02	50.0
WLB-APUB	$\gamma = 0.3$	0.84	0.34	1.04	53.3
WLB-APUB	$\gamma = 0.4$	0.84	0.34	1.04	60.0
WLB-APUB	$\gamma = 0.5$	0.83	0.33	1.05	66.7
WLB-APUB	$\gamma = 0.6$	0.82	0.33	1.02	80.0
WLB-APUB	$\gamma = 0.7$	0.80	0.34	0.98	83.3
WLB-APUB	$\gamma = 0.8$	0.79	0.31	0.92	86.7
Bayesian Prior 1	$\gamma = 0.0$	1.27	0.26	2.11	50.0
Bayesian Prior 1	$\gamma = 0.1$	1.28	0.26	2.12	50.0
Bayesian Prior 1	$\gamma = 0.2$	1.27	0.26	2.14	53.3
Bayesian Prior 1	$\gamma = 0.3$	1.27	0.26	2.15	66.7
Bayesian Prior 1	$\gamma = 0.4$	1.28	0.26	2.17	76.7
Bayesian Prior 1	$\gamma = 0.5$	1.28	0.26	2.18	80.0
Bayesian Prior 1	$\gamma = 0.6$	1.28	0.26	2.20	86.7
Bayesian Prior 1	$\gamma = 0.7$	1.29	0.27	2.19	86.7
Bayesian Prior 1	$\gamma = 0.8$	1.29	0.27	2.18	90.0

Table 3 continued

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
Bayesian Prior 2	$\gamma = 0.0$	1.51	0.28	2.45	0.0
Bayesian Prior 2	$\gamma = 0.1$	1.50	0.26	2.43	6.7
Bayesian Prior 2	$\gamma = 0.2$	1.48	0.25	2.41	6.7
Bayesian Prior 2	$\gamma = 0.3$	1.47	0.24	2.39	10.0
Bayesian Prior 2	$\gamma = 0.4$	1.46	0.23	2.39	10.0
Bayesian Prior 2	$\gamma = 0.5$	1.45	0.23	2.35	13.3
Bayesian Prior 2	$\gamma = 0.6$	1.43	0.23	2.36	26.7
Bayesian Prior 2	$\gamma = 0.7$	1.41	0.24	2.35	36.7
Bayesian Prior 2	$\gamma = 0.8$	1.41	0.26	2.40	43.3
Moment DRO	$\kappa = 0.0$	2.23	1.16	3.13	100.0
Moment DRO	$\kappa = 0.1$	2.35	1.31	3.28	100.0
Moment DRO	$\kappa = 0.2$	2.48	1.49	3.43	100.0
Moment DRO	$\kappa = 0.3$	2.63	1.67	3.60	100.0
Moment DRO	$\kappa = 0.4$	2.80	1.86	3.80	100.0
Moment DRO	$\kappa = 0.5$	2.98	1.99	3.99	100.0
Moment DRO	$\kappa = 0.6$	3.17	2.13	4.19	100.0
Moment DRO	$\kappa = 0.7$	3.37	2.29	4.40	100.0
Moment DRO	$\kappa = 0.8$	3.59	2.46	4.63	100.0
SAA	—	1.02	0.50	1.11	43.3
FMO	—	0.92	0.34	1.07	46.7

D FIXED-TRUTH MIXED-LOG-NORMAL EXPERIMENT

D.1 DESIGN AND ANALYTIC EVALUATION

The mixed-log-normal experiment repeats the fixed-truth, paired-replication design with a 20-product, three-component log-normal mixture. Thirty independent master datasets of size 1000 are generated from the declared fixed population using seed 20260911; the $N = 100$ and $N = 200$ datasets are nested prefixes within each replication. Every method receives the same data at a given replication and sample size. The distributional truth is used only for data generation and analytic evaluation, never for fitting, prior selection, or regret-based tuning.

The newsvendor economics remain $p = 10$, $c = 5$, and $s = 1$. For a log-normal marginal $D = \exp(Y)$ with $Y \sim N(\mu, \sigma^2)$,

$$\mathbb{E}(x - D)_+ = x\Phi\left(\frac{\log x - \mu}{\sigma}\right) - \exp(\mu + \sigma^2/2)\Phi\left(\frac{\log x - \mu - \sigma^2}{\sigma}\right), \qquad x > 0.$$

The mixture expected cost and gradient are the mixture-weighted sums of this expression. The population optimum is again the coordinatewise $5/9$ marginal quantile, so reported regret uses analytic true costs rather than a finite test sample.

D.2 METHODS, PRIORS, AND NUMERICAL SETTINGS

The method set, levels, paired data protocol, and evaluation metrics are the same as in Appendix C: SAA, FMO, NP-APUB, PB-APUB, WLB-APUB, Bayesian posterior-CVaR under Prior 1 and Prior 2, and full-covariance moment DRO. Fitted methods operate on the three-component log-normal-mixture family, using EM in log-demand space. Original-data fits use ten starts and bootstrap fits use three starts, covariance-eigenvalue floor 1.5625×10^{-4}, relative likelihood tolerance 10^{-8}, and a 1000-iteration initial budget; weighted fits can continue to 4000 iterations. Bootstrap ensembles grow by 500 models up to 10000, with the same hybrid solver and stopping diagnostics used in the mixed-normal case.

The Bayesian definitions are identical in role to those in the mixed-normal case. Prior 1 is selected to align with the underlying distribution when the analyst has good prior knowledge; Prior 2 represents potentially misspecified knowledge. Their log-space hyperparameters are fixed before evaluation: Prior 1 has $\kappa_0 = .1$, log-scale prior standard deviation .5, and prior-mean multiplier 1.0; Prior 2 has $\kappa_0 = 10$, log-scale prior standard deviation .35, and prior-mean multiplier 1.2. Neither is tuned using true regret. Moment DRO uses the same full-covariance ambiguity construction and κ grid; κ is a radius multiplier, not a confidence level.

D.3 COMPLETE CURVES, TABLES, AND NUMERICAL QUALIFICATIONS

All 30 replications are complete for every method, sample size, and configured level. The audit retains 5040 saved decisions. Figure 2 in the main body gives the complete curves: blue lines and bands are mean regret and the 20th–80th percentile interval, and red lines are empirical coverage. The panels use row-specific regret scales so that the Prior 2 small-sample stress outcomes remain visible without obscuring the other methods. The full numerical grid follows in Tables 4–6.

At $N = 100$, $16/30$ Prior 2 posterior runs do not meet the declared diagnostic criteria at every displayed γ; the saved outcomes are retained in the table rather than discarded. The corresponding count is $4/30$ at $N = 200$ and zero at $N = 1000$. One Prior 1 run at $N = 1000$ does not meet the diagnostic criteria. For DRO, two $N = 1000$ outputs at $\kappa = 0$ exhaust the iteration budget; all other reported DRO outputs meet the configured gap criterion. These counts qualify the numerical comparison and do not change the reported summaries.

D.4 COMPLETE NUMERICAL RESULTS

Columns report mean true-cost regret, the 20th and 80th percentiles of regret, and empirical coverage (%). All summaries use the 30 paired training datasets. For APUB and Bayesian posterior-CVaR, the level is γ; for moment DRO, it is the radius multiplier κ. At $\gamma = 0$, APUB is BAG.

Table 4: Complete fixed mixed-log-normal results for $N = 100$.

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
NP-APUB	$\gamma = 0.0$	17.11	7.90	25.57	30.0
NP-APUB	$\gamma = 0.1$	17.08	7.87	25.57	43.3
NP-APUB	$\gamma = 0.2$	16.36	8.04	24.23	50.0
NP-APUB	$\gamma = 0.3$	16.30	8.29	22.76	56.7
NP-APUB	$\gamma = 0.4$	16.38	8.09	21.38	63.3
NP-APUB	$\gamma = 0.5$	16.07	8.41	20.84	63.3
NP-APUB	$\gamma = 0.6$	16.26	8.48	20.79	70.0
NP-APUB	$\gamma = 0.7$	16.11	8.42	20.18	76.7
NP-APUB	$\gamma = 0.8$	15.94	8.40	20.37	80.0
PB-APUB	$\gamma = 0.0$	14.55	4.70	21.92	40.0
PB-APUB	$\gamma = 0.1$	14.23	5.07	21.20	50.0
PB-APUB	$\gamma = 0.2$	13.93	5.37	20.30	53.3
PB-APUB	$\gamma = 0.3$	13.73	5.40	19.13	60.0
PB-APUB	$\gamma = 0.4$	13.61	5.56	18.19	66.7
PB-APUB	$\gamma = 0.5$	13.48	5.57	18.36	70.0
PB-APUB	$\gamma = 0.6$	13.43	5.78	18.05	70.0
PB-APUB	$\gamma = 0.7$	13.42	5.86	18.31	80.0
PB-APUB	$\gamma = 0.8$	13.38	5.95	18.48	83.3
WLB-APUB	$\gamma = 0.0$	14.10	4.94	21.86	33.3
WLB-APUB	$\gamma = 0.1$	13.81	5.38	21.50	46.7
WLB-APUB	$\gamma = 0.2$	13.65	5.57	20.38	53.3
WLB-APUB	$\gamma = 0.3$	13.47	5.71	19.40	56.7
WLB-APUB	$\gamma = 0.4$	13.36	5.56	18.49	66.7
WLB-APUB	$\gamma = 0.5$	13.24	5.46	18.83	70.0
WLB-APUB	$\gamma = 0.6$	13.14	5.51	19.03	70.0
WLB-APUB	$\gamma = 0.7$	13.11	5.75	19.14	80.0
WLB-APUB	$\gamma = 0.8$	13.16	5.40	18.61	80.0
Bayesian Prior 1	$\gamma = 0.0$	13.29	4.82	20.99	26.7
Bayesian Prior 1	$\gamma = 0.1$	13.19	4.63	21.34	36.7
Bayesian Prior 1	$\gamma = 0.2$	13.14	4.46	21.58	46.7
Bayesian Prior 1	$\gamma = 0.3$	13.12	4.33	21.84	56.7
Bayesian Prior 1	$\gamma = 0.4$	13.12	4.37	22.06	66.7
Bayesian Prior 1	$\gamma = 0.5$	13.09	4.56	22.20	73.3
Bayesian Prior 1	$\gamma = 0.6$	13.11	4.55	22.33	76.7
Bayesian Prior 1	$\gamma = 0.7$	13.20	4.71	22.57	80.0
Bayesian Prior 1	$\gamma = 0.8$	13.23	4.64	22.53	83.3
Bayesian Prior 2	$\gamma = 0.0$	74.35	47.78	91.69	0.0
Bayesian Prior 2	$\gamma = 0.1$	68.37	42.00	84.53	0.0
Bayesian Prior 2	$\gamma = 0.2$	63.98	38.12	79.27	0.0
Bayesian Prior 2	$\gamma = 0.3$	60.04	34.95	75.08	0.0
Bayesian Prior 2	$\gamma = 0.4$	56.32	32.05	71.42	0.0
Bayesian Prior 2	$\gamma = 0.5$	52.64	29.33	67.91	0.0
Bayesian Prior 2	$\gamma = 0.6$	48.81	26.67	63.95	0.0
Bayesian Prior 2	$\gamma = 0.7$	44.69	23.83	60.13	0.0
Bayesian Prior 2	$\gamma = 0.8$	39.93	21.04	55.48	6.7
Moment DRO	$\kappa = 0.0$	44.40	29.45	53.18	100.0
Moment DRO	$\kappa = 0.1$	40.20	26.25	50.11	100.0
Moment DRO	$\kappa = 0.2$	36.31	23.08	47.19	100.0
Moment DRO	$\kappa = 0.3$	32.73	20.05	44.45	100.0
Moment DRO	$\kappa = 0.4$	29.46	17.28	41.48	100.0
Moment DRO	$\kappa = 0.5$	26.51	14.77	37.57	100.0
Moment DRO	$\kappa = 0.6$	23.87	12.52	34.01	100.0
Moment DRO	$\kappa = 0.7$	21.56	10.55	30.79	100.0
Moment DRO	$\kappa = 0.8$	19.58	8.93	27.96	100.0
SAA	—	17.11	7.90	25.57	26.7

Table 4 continued

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
FMO	—	14.54	4.69	22.54	40.0

Table 5: Complete fixed mixed-log-normal results for $N = 200$.

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
NP-APUB	$\gamma = 0.0$	10.73	4.59	16.79	30.0
NP-APUB	$\gamma = 0.1$	10.44	4.81	15.12	40.0
NP-APUB	$\gamma = 0.2$	10.37	4.74	14.97	46.7
NP-APUB	$\gamma = 0.3$	10.20	4.99	14.38	46.7
NP-APUB	$\gamma = 0.4$	10.03	4.99	13.92	50.0
NP-APUB	$\gamma = 0.5$	10.05	4.96	14.46	50.0
NP-APUB	$\gamma = 0.6$	10.03	5.28	14.00	53.3
NP-APUB	$\gamma = 0.7$	9.89	5.25	14.38	60.0
NP-APUB	$\gamma = 0.8$	9.92	5.37	14.57	66.7
PB-APUB	$\gamma = 0.0$	9.58	2.67	14.22	40.0
PB-APUB	$\gamma = 0.1$	9.33	2.71	13.53	40.0
PB-APUB	$\gamma = 0.2$	9.17	2.77	13.02	46.7
PB-APUB	$\gamma = 0.3$	8.99	2.85	13.00	50.0
PB-APUB	$\gamma = 0.4$	8.85	2.87	12.98	53.3
PB-APUB	$\gamma = 0.5$	8.69	2.98	13.10	53.3
PB-APUB	$\gamma = 0.6$	8.61	3.12	13.25	53.3
PB-APUB	$\gamma = 0.7$	8.51	3.31	13.40	60.0
PB-APUB	$\gamma = 0.8$	8.42	3.50	13.35	73.3
WLB-APUB	$\gamma = 0.0$	8.96	2.51	12.89	36.7
WLB-APUB	$\gamma = 0.1$	8.80	2.55	13.07	36.7
WLB-APUB	$\gamma = 0.2$	8.66	2.58	13.12	40.0
WLB-APUB	$\gamma = 0.3$	8.52	2.65	12.66	46.7
WLB-APUB	$\gamma = 0.4$	8.38	2.77	12.65	50.0
WLB-APUB	$\gamma = 0.5$	8.28	2.83	12.61	50.0
WLB-APUB	$\gamma = 0.6$	8.21	2.99	12.81	53.3
WLB-APUB	$\gamma = 0.7$	8.17	3.02	12.64	56.7
WLB-APUB	$\gamma = 0.8$	8.13	3.32	12.63	66.7
Bayesian Prior 1	$\gamma = 0.0$	8.82	2.88	13.56	13.3
Bayesian Prior 1	$\gamma = 0.1$	8.76	2.94	13.43	26.7
Bayesian Prior 1	$\gamma = 0.2$	8.73	2.99	13.36	33.3
Bayesian Prior 1	$\gamma = 0.3$	8.70	3.04	13.24	40.0
Bayesian Prior 1	$\gamma = 0.4$	8.70	3.14	13.00	46.7
Bayesian Prior 1	$\gamma = 0.5$	8.69	3.19	12.71	50.0
Bayesian Prior 1	$\gamma = 0.6$	8.70	3.21	12.50	50.0
Bayesian Prior 1	$\gamma = 0.7$	8.76	3.39	12.51	63.3
Bayesian Prior 1	$\gamma = 0.8$	8.75	3.46	12.48	70.0
Bayesian Prior 2	$\gamma = 0.0$	22.58	16.00	29.26	0.0
Bayesian Prior 2	$\gamma = 0.1$	21.16	14.76	27.67	0.0
Bayesian Prior 2	$\gamma = 0.2$	20.10	13.93	26.50	0.0
Bayesian Prior 2	$\gamma = 0.3$	19.14	12.92	25.29	0.0
Bayesian Prior 2	$\gamma = 0.4$	18.26	12.04	24.39	0.0
Bayesian Prior 2	$\gamma = 0.5$	17.37	11.31	23.27	0.0
Bayesian Prior 2	$\gamma = 0.6$	16.42	10.50	22.31	0.0
Bayesian Prior 2	$\gamma = 0.7$	15.44	9.90	21.29	3.3
Bayesian Prior 2	$\gamma = 0.8$	14.34	8.87	20.42	6.7
Moment DRO	$\kappa = 0.0$	39.84	29.66	51.20	100.0
Moment DRO	$\kappa = 0.1$	36.83	26.91	47.13	100.0
Moment DRO	$\kappa = 0.2$	33.96	24.35	43.27	100.0
Moment DRO	$\kappa = 0.3$	31.23	21.96	39.78	100.0

Table 5 continued

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
Moment DRO	$\kappa = 0.4$	28.65	19.73	37.02	100.0
Moment DRO	$\kappa = 0.5$	26.22	17.74	34.40	100.0
Moment DRO	$\kappa = 0.6$	23.95	15.93	31.92	100.0
Moment DRO	$\kappa = 0.7$	21.82	14.31	29.69	100.0
Moment DRO	$\kappa = 0.8$	19.86	12.82	27.74	100.0
SAA	—	10.68	4.45	17.00	30.0
FMO	—	9.54	2.68	14.06	40.0

Table 6: Complete fixed mixed-log-normal results for $N = 1000$.

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
NP-APUB	$\gamma = 0.0$	1.98	0.71	2.84	30.0
NP-APUB	$\gamma = 0.1$	1.96	0.72	2.87	36.7
NP-APUB	$\gamma = 0.2$	1.95	0.72	2.94	43.3
NP-APUB	$\gamma = 0.3$	1.96	0.75	3.00	50.0
NP-APUB	$\gamma = 0.4$	1.94	0.78	3.04	56.7
NP-APUB	$\gamma = 0.5$	1.95	0.82	3.16	63.3
NP-APUB	$\gamma = 0.6$	1.94	0.86	3.19	70.0
NP-APUB	$\gamma = 0.7$	1.94	0.88	3.21	76.7
NP-APUB	$\gamma = 0.8$	1.96	1.00	3.33	80.0
PB-APUB	$\gamma = 0.0$	1.72	0.46	2.59	36.7
PB-APUB	$\gamma = 0.1$	1.68	0.48	2.70	40.0
PB-APUB	$\gamma = 0.2$	1.67	0.49	2.81	43.3
PB-APUB	$\gamma = 0.3$	1.65	0.50	2.88	53.3
PB-APUB	$\gamma = 0.4$	1.64	0.49	2.97	56.7
PB-APUB	$\gamma = 0.5$	1.62	0.52	3.02	60.0
PB-APUB	$\gamma = 0.6$	1.61	0.54	3.06	70.0
PB-APUB	$\gamma = 0.7$	1.60	0.57	3.13	76.7
PB-APUB	$\gamma = 0.8$	1.60	0.61	3.24	83.3
WLB-APUB	$\gamma = 0.0$	1.67	0.35	2.86	33.3
WLB-APUB	$\gamma = 0.1$	1.64	0.37	2.94	33.3
WLB-APUB	$\gamma = 0.2$	1.62	0.39	3.00	40.0
WLB-APUB	$\gamma = 0.3$	1.60	0.41	3.10	43.3
WLB-APUB	$\gamma = 0.4$	1.59	0.42	3.13	46.7
WLB-APUB	$\gamma = 0.5$	1.59	0.45	3.17	53.3
WLB-APUB	$\gamma = 0.6$	1.58	0.48	3.23	60.0
WLB-APUB	$\gamma = 0.7$	1.59	0.51	3.40	70.0
WLB-APUB	$\gamma = 0.8$	1.59	0.54	3.52	76.7
Bayesian Prior 1	$\gamma = 0.0$	2.12	0.47	4.15	30.0
Bayesian Prior 1	$\gamma = 0.1$	2.13	0.48	4.20	36.7
Bayesian Prior 1	$\gamma = 0.2$	2.15	0.49	4.18	36.7
Bayesian Prior 1	$\gamma = 0.3$	2.16	0.49	4.18	40.0
Bayesian Prior 1	$\gamma = 0.4$	2.18	0.50	4.17	43.3
Bayesian Prior 1	$\gamma = 0.5$	2.19	0.50	4.17	50.0
Bayesian Prior 1	$\gamma = 0.6$	2.22	0.52	4.05	53.3
Bayesian Prior 1	$\gamma = 0.7$	2.25	0.54	3.83	63.3
Bayesian Prior 1	$\gamma = 0.8$	2.29	0.58	3.75	70.0
Bayesian Prior 2	$\gamma = 0.0$	3.68	1.17	6.80	0.0
Bayesian Prior 2	$\gamma = 0.1$	3.64	1.15	6.72	0.0
Bayesian Prior 2	$\gamma = 0.2$	3.61	1.11	6.67	0.0
Bayesian Prior 2	$\gamma = 0.3$	3.59	1.09	6.63	3.3
Bayesian Prior 2	$\gamma = 0.4$	3.57	1.06	6.58	6.7
Bayesian Prior 2	$\gamma = 0.5$	3.56	1.02	6.52	13.3
Bayesian Prior 2	$\gamma = 0.6$	3.55	1.02	6.54	23.3

Table 6 continued

Method	Level	Mean regret	20th pct.	80th pct.	Coverage (%)
Bayesian Prior 2	$\gamma = 0.7$	3.52	1.00	6.51	23.3
Bayesian Prior 2	$\gamma = 0.8$	3.52	1.00	6.47	26.7
Moment DRO	$\kappa = 0.0$	32.66	27.47	36.13	100.0
Moment DRO	$\kappa = 0.1$	31.50	26.24	34.91	100.0
Moment DRO	$\kappa = 0.2$	30.22	25.04	33.71	100.0
Moment DRO	$\kappa = 0.3$	28.97	23.86	32.54	100.0
Moment DRO	$\kappa = 0.4$	27.75	22.72	31.40	100.0
Moment DRO	$\kappa = 0.5$	26.56	21.62	30.28	100.0
Moment DRO	$\kappa = 0.6$	25.39	20.54	29.18	100.0
Moment DRO	$\kappa = 0.7$	24.26	19.50	28.11	100.0
Moment DRO	$\kappa = 0.8$	23.15	18.48	27.07	100.0
SAA	—	1.98	0.71	2.83	30.0
FMO	—	1.69	0.46	2.47	36.7