LEARNING CONTEXT MATCHED UNCERTAINTY REPRESENTATIONS FOR ROBUST OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Two-stage robust optimization is shaped by the adversarial scenarios admitted by its uncertainty set. When side information is observed before a decision, learning this set from paired context–scenario data requires both identifying relevant historical uncertainty patterns and representing their nearby variation. However, a learned uncertainty set may still admit scenarios poorly matched to the observed context. The adversarial problem can return these scenarios when they maximize recourse cost. Existing methods restrict such scenarios through context-matching constraints in the adversarial problem, which can complicate its solution and require problem-specific reformulations. To address this issue, we propose CUB-RO, a nonparametric uncertainty-set learner that incorporates context matching into set construction, avoiding additional context-matching constraints in the adversarial problem. Given a new context, CUB-RO retrieves the K nearest historical contexts using a prespecified dissimilarity measure and constructs the uncertainty set as a union of local regions around their paired complete scenarios. The matched centers and local radii jointly determine the uncertainty-set boundary. We establish sufficient conditions for the learned set to converge to the true conditional support as the sample size grows. Moreover, we derive finite-sample bounds on the excess worst-case cost of the resulting decisions relative to the true conditional optimum. Experiments on Feature-Based Newsvendor (FBNV), Robust Portfolio Optimization (RPO), and Security-Constrained Unit Commitment (SCUC) show improvements over the evaluated baselines on the respective primary metrics. Compared with robust optimization using only the retrieved historical scenarios as the uncertainty set, CUB-RO reduces value-at-risk and conditional value-at-risk on RPO by 10.63% and 12.92%, respectively, and lowers mean realized costs by 4.78% on FBNV and 0.58% on SCUC.

1 INTRODUCTION

Two-stage robust optimization (2RO) converts uncertainty information into decisions through an uncertainty set: the decision is optimized against adverse scenarios admitted by that set (Ben-Tal et al., 2004; Bertsimas & Sim, 2004). In many applications, the decision maker also observes side information before choosing the decision. This information is not the uncertain quantity to be hedged; it indicates which past uncertainty outcomes are relevant to the current instance (Sadana et al., 2025). Weather forecasts help a grid operator identify relevant regimes of renewable forecast error, while calendar, price, and promotion information identify relevant demand regimes in inventory planning (Ban & Rudin, 2019; Zhang et al., 2024). The uncertainty set should therefore account for the observed context.

Recent research has developed data-driven methods that use paired observations of contexts and uncertain outcomes to construct conditional uncertainty sets. Existing approaches include jointly learning context partitions and their associated uncertainty regions (Chenreddy et al., 2022), and predicting uncertain outcomes followed by region calibration (Sun et al., 2023). Sample robust optimization with side information (SRO-SI) uses contextual information to weight historical scenarios, assigning lower weights to less context-relevant scenarios in the empirical distribution (Bertsimas et al., 2023a). Although context enters these formulations differently, their worst-case searches operate over the entire constructed regions or sample neighborhoods.

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.

However, in a standard 2RO formulation, the hedge is determined by an inner adversarial search. For a given first-stage decision, the adversarial problem maximizes recourse cost over every scenario admitted by the uncertainty set. A region constructed to encompass context-relevant observations may also contain scenarios poorly matched to the observed context. These scenarios remain feasible for the adversarial problem and can be returned as worst-case scenarios when they maximize recourse cost (see Section 4).

Existing methods restrict such scenarios by introducing context-matching constraints into the adversarial problem (Esteban-Pérez & Morales, 2022; Kannan et al., 2024; Nguyen et al., 2025; Yang et al., 2026). The adversarial problem then enforces context compatibility while maximizing robust loss, which complicates its solution and may require problem-specific reformulations. This motivates a simple uncertainty-set learning method that limits the inclusion of context-incompatible scenarios during set construction, without requiring additional, problem-specific matching constraints in the adversarial problem.

We propose CUB-RO, which learns a local conditional uncertainty representation before the adversarial query. From paired context–scenario observations, a prespecified context dissimilarity measure retrieves the K histories nearest to a query s. CUB-RO retains their paired complete trajectories and assigns each a radius using a shared scaling matrix. Its output is a union of local uncertainty sets whose matched centers and radii determine the uncertainty-set boundary. Thus context determines the feasible domain of the adversarial query: an unmatched history does not define a local component, and every admitted scenario lies in a local region around a matched trajectory. Context selects and localizes the components, while the adversarial problem only maximizes the recourse cost. Because the learned object is an explicit scenario set, it can be used with continuous or mixed-integer two-stage models through their standard robust optimization interfaces; solver details are deferred to the appendix.

Our contributions are:

- We propose CUB-RO, which constructs a union of local uncertainty regions around context-matched historical scenarios, separating context matching from adversarial optimization.
- We bound decision suboptimality under set approximation and inexact adversarial optimization, analyze the role of local radii, and establish sufficient conditions for set and decision-value consistency.
- We analyze how uncertainty-set geometry affects worst-case recourse costs and use finite-reference experiments on a benchmark with real-world data to examine the effects of context matching and local expansion on representation error and decision quality.

2 RELATED WORK

Learning uncertainty-set geometry and locality. Data-driven robust optimization learns uncertainty sets from observations using statistically validated prediction regions (Tulabandhula & Rudin, 2014; Hong et al., 2021; Bertsimas et al., 2021), support-vector clustering (Shang et al., 2017), and deep nonconvex representations (Goerigk & Kurtz, 2023). AGRO (Brenner et al., 2025) maps a latent ball through a variational autoencoder decoder to obtain a marginal uncertainty set. Conditional uncertainty-set methods incorporate observed covariates through KNN-based ellipsoids (Ohmori, 2021), prediction and calibration (Sun et al., 2023), conformal prediction regions (Johnstone & Cox, 2021; Patel et al., 2024), or joint context partitioning and region learning as in IDCC (Chenreddy et al., 2022). These methods use context to guide region construction, but the resulting regions are defined by geometric constraints rather than solely by the observations used to construct them. Their worst-case searches therefore also consider additional scenarios admitted by these constraints, whose relevance to the current context may differ from that of the underlying observations. Sample-path robust optimization constructs uncertainty neighborhoods around historical trajectories and averages their worst-case costs to approximate multistage stochastic optimization (Bertsimas et al., 2023b). Sample Robust Optimization with Side Information (SRO-SI) incorporates context-dependent weights into this formulation to approximate conditional expected cost (Bertsimas et al., 2023a). Here, contextual information determines how the neighborhood-wise worst-case costs are weighted in the objective, while each worst-case cost is evaluated over the corresponding sample neighborhood. In contrast, CUB-RO uses contextual retrieval to select historical

scenarios as local centers and determines their radii from distances among the selected scenarios. The matched centers and local radii define the uncertainty set for the adversarial problem.

Contextual matching in the adversarial problem. To restrict scenarios poorly matched to the observed context, existing methods impose context-matching constraints in the adversarial problem. In transport-based formulations, these constraints require candidate distributions to concentrate probability mass on the conditioning event or allocate a prescribed minimum mass to a neighborhood of the observed covariates, while satisfying a transport budget (Esteban-Pérez & Morales, 2022; Nguyen et al., 2025). Matching requirements also concern how probability mass can be transported: causal transport requires covariate perturbations to be independent of the paired original outcomes conditional on the original covariates (Yang et al., 2026). The inner problem thus evaluates worst-case costs while enforcing restrictions on probability allocation or transport. Tractable reformulations exploit the structure of the loss function, support set, or decision rules. In contrast, CUB-RO specifies the uncertainty set through contextual retrieval and local region construction before solving the adversarial problem, which maximizes recourse cost over the learned set without additional context-matching constraints.

Recourse approximation and solution methods. Once the uncertainty set is specified, solving a two-stage robust model requires accounting for adaptive recourse and identifying worst-case scenarios. Decision-rule methods approximate recourse decisions as functions of uncertainty, with affine rules providing a simple representation and deep lifting enabling more flexible piecewise-linear policies (Kuhn et al., 2011; Rahal et al., 2022). Column-and-constraint generation (C&CG) instead coordinates a master problem with a worst-case scenario subproblem, iteratively adding scenarios and their associated recourse variables and constraints; nested C&CG extends this decomposition to mixed-integer recourse (Zeng & Zhao, 2013; Zhao & Zeng, 2012). Learning can support this solution process by predicting strategies from previously solved instances or approximating recourse values within C&CG, as in Bertsimas & Kim (2024) and Neur2RO (Dumouchelle et al., 2024), respectively. Our CRG-PGA component addresses the worst-case scenario subproblem for linear programming recourse under right-hand-side uncertainty. It uses affine pieces obtained from the original recourse problem to guide multi-start projected gradient ascent and evaluates the returned scenario with that same recourse model. This component supports the solution of the robust model defined by the learned uncertainty set; its implementation is detailed in Appendix A.3.

3 PRELIMINARIES

Contextual two-stage robust optimization. Before the uncertainty $z \in \mathbb{R}^d$ is realized, a planner observes side information (covariates) $s \in \mathbb{R}^p$ and commits a first-stage decision $x \in \mathcal{X}$ at cost $c^\top x$. After z is revealed, a recourse action incurs cost $Q(x, z) = \min_{y \in \mathcal{Y}(x,z)} q^\top y$, the optimal value of a second-stage program. We seek a *policy* $x(s)$ that, for each context, hedges the recourse cost against an uncertainty set $\mathcal{U}(s)$ that may depend on s:

$$\min_{x \in \mathcal{X}} \; c^\top x \; + \; \max_{z \in \mathcal{U}(s)} \; Q(x, z). \tag{1}$$

The first stage $\mathcal{X}$ may be *non-convex*, e.g. contain binary commitments (mixed-integer). We are given i.i.d. historical pairs $\{(s_i, z_i)\}_{i=1}^n$ and evaluate a policy by its realized cost: for a held-out instance with context s, observed z^{act}, and decision $x(s)$, the realized cost is $c^\top x(s) + Q(x(s), z^{\mathrm{act}})$.

Target and learned uncertainty sets. For the analysis, let $\mathcal{C}(s)$ be a nonempty compact set of scenarios judged relevant under context s; it may be a conditional level set, a prediction region, or a domain-defined target. Define its worst-case recourse cost and total cost as

$$H_{\mathcal{C}}(x, s) := \sup_{z \in \mathcal{C}(s)} Q(x, z), \qquad F_{\mathcal{C}}(x, s) := c^\top x + H_{\mathcal{C}}(x, s). \tag{2}$$

Let $\mathcal{U}_\theta(s)$ denote the learned uncertainty set. Section 5.2 analyzes how differences between the learned and target sets affect worst-case costs and the quality of the resulting decisions.

Scenario generation. For a fixed context and first-stage decision, we call the inner maximization over the uncertainty set the adversarial problem (AP). In C&CG, a master problem optimizes the first-stage decision over a finite subset of scenarios, and the AP searches the uncertainty set for a worst-case scenario under that decision. The returned scenario and its associated recourse variables and constraints are added to the master problem, and the procedure repeats (Zeng & Zhao,

2013). Projected gradient ascent (PGA) is a heuristic for the AP that updates search variables along recourse-cost gradients and projects each update onto their feasible domain (Brenner et al., 2025).

4 Motivation: Learning Context-Matched Uncertainty Sets

Learned uncertainty representations are combined with geometric constraints to define feasible domains for robust optimization. IDCC uses a calibrated ellipsoid in the learned feature space for each context cluster (Chenreddy et al., 2022), while cVAE maps a calibrated latent region through a decoder conditioned on the observed context. In both cases, historical data determine the extent of the search domain. Although contextual information guides scenario selection or generation, the boundaries calibrated from historical data may still admit scenarios poorly matched to the current context.

For a fixed context s and first-stage decision x, the worst-case recourse cost is

$$V_s(x) = \sup_{z \in \mathcal{U}(s)} Q(x, z), \tag{3}$$

Scenarios poorly matched to the current context remain feasible if admitted by the set and may be returned when they maximize recourse cost.

We examine this mechanism on the Security-Constrained Unit Commitment (SCUC) benchmark. Figure 1 compares IDCC, cVAE, and CUB-RO under the same observed context and fixed first-stage decision. Green diamonds denote context-matched scenarios selected or generated by each method, blue regions depict the uncertainty sets, and pink stars mark the returned adversarial scenarios. In panels (a) and (b), the returned scenarios lie outside the matched scenario clouds in the projection, with recourse costs of 3298 and 259, respectively. In panel (c), the returned scenario lies within a local region around a matched historical trajectory, with a recourse cost of 108. All costs are evaluated using the original 144-dimensional recourse model; the common PCA projection is used only for visualization.

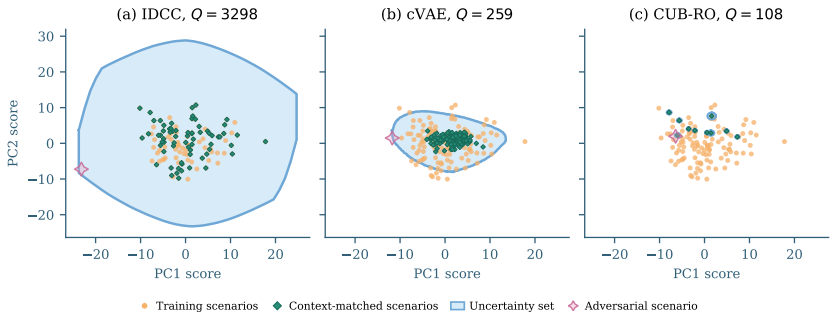

Figure 1: Uncertainty-set representations on the SCUC benchmark.

These observations motivate controlling the permitted variation around matched trajectories during set construction. CUB-RO retains context-matched historical trajectories as centers and assigns a radius to each local region. Their union defines the adversarial feasible domain, so that every admitted scenario lies within a specified neighborhood of a matched trajectory. For the convex scenario-wise losses considered here, the union and its convex hull have the same worst-case value; a larger enclosing region can admit additional extreme scenarios and increase that value. Section 5 formalizes this construction and analyzes its implications for decision quality.

5 Methodology

5.1 A context-matched union of local scenario sets

CUB-RO constructs a conditional uncertainty set by retrieving historical observations with similar contexts and forming local regions around their paired scenarios. Context dissimilarity determines the selected centers, while distances between the selected scenarios determine the local radii.

Given paired observations $\mathcal{D} = \{(s_i, z_i)\}_{i=1}^n$ and an observed context s, let Δ_S denote the context dissimilarity measure. We select the indices of the K nearest historical contexts:

$$\mathcal{N}_K(s) \in \operatorname*{arg\,min}_{\substack{I \subseteq \{1,\dots,n\} \\ |I|=K}} \sum_{i \in I} \Delta_S(s, s_i). \tag{4}$$

Their paired complete scenarios serve as the centers of the local uncertainty sets.

Let B_s be a nonsingular scaling matrix estimated from training data and shared by all local sets for context s. Define the distance between scenarios as $\delta_s(z, z') = \|B_s^{-1}(z - z')\|_2$. For each selected center, let $d_i(s) = \min_{j \in \mathcal{N}_K(s) \setminus \{i\}} \delta_s(z_i, z_j)$ for $K \geq 2$, and set $d_i(s) = 0$ for $K = 1$. The radius around each center is its distance to the nearest other selected center, multiplied by a nonnegative factor $\widehat{\beta}$. We define

$$\mathcal{U}(s) = \bigcup_{i \in \mathcal{N}_K(s)} \left\{ z_i + B_s u : \|u\|_2 \leq \widehat{\beta} d_i(s) \right\}. \tag{5}$$

Each local set is an affine image of a Euclidean ball centered at a selected historical scenario. Appendix B specifies the context preprocessing, dissimilarity measure, scaling matrix, and parameter selection for each benchmark.

By construction, every admitted scenario lies within the prescribed radius of at least one selected center:

$$z \in \mathcal{U}(s) \quad \Longrightarrow \quad \delta_s(z, z_i) \leq \widehat{\beta} d_i(s) \quad \text{for some } i \in \mathcal{N}_K(s). \tag{6}$$

The selected centers and their radii therefore determine the feasible domain of the AP, whose objective remains the recourse cost.

Proposition 1 (Convex-hull invariance)**.** *Fix a first-stage decision x, and let $\mathcal{U}(s)$ be the finite union in equation 5. If $Q(x, \cdot)$ is continuous and convex, then*

$$\max_{z \in \mathcal{U}(s)} Q(x, z) = \max_{z \in \operatorname{conv}(\mathcal{U}(s))} Q(x, z), \tag{7}$$

and a maximizer over the convex hull can be chosen from $\mathcal{U}(s)$. If the recourse function has the finite convex piecewise-affine representation $Q(x, z) = \max_p \{a_p(x) + b_p^\top z\}$, then

$$H_{\mathcal{U}}(x, s) = \max_{p,i} \left\{ a_p(x) + b_p^\top z_i + \widehat{\beta} d_i(s) \|B_s^\top b_p\|_2 \right\}. \tag{8}$$

Equation 8 makes explicit how the selected centers and local radii determine the worst-case recourse cost. Under the same piecewise-affine representation, any nonempty compact comparison set $\mathcal{G}(s)$ satisfies

$$H_{\mathcal{G}}(x, s) = \max_p \{a_p(x) + \sigma_{\mathcal{G}(s)}(b_p)\}. \tag{9}$$

where $\sigma_A(b) = \max_{z \in A} b^\top z$ is the support function of A. Thus, $\mathcal{G}(s)$ has a higher worst-case recourse cost if an affine piece attains a value on $\mathcal{G}(s)$ greater than $H_{\mathcal{U}}(x, s)$.

5.2 Decision bounds under uncertainty-set approximation

The local construction in Section 5.1 controls the variation admitted around selected historical scenarios. With finite data, however, the resulting set may omit relevant uncertainty or include additional scenarios. Restricting the set can reduce conservatism but may also leave the decision insufficiently protected. We therefore analyze how these two types of approximation error affect worst-case costs and the suboptimality of the resulting decision relative to the target uncertainty set. To measure scenario differences in terms of their effects on recourse costs, define the pseudometric

$$d_Q(z, z') := \sup_{x \in \mathcal{X}} |Q(x, z) - Q(x, z')|. \tag{10}$$

Let $\mathcal{C}(s)$ be the target context-relevant set in equation 2. For any learned representation $\mathcal{U}(s)$, define its directed recourse distortions as the following set-approximation errors under d_Q:

$$\varepsilon_{\text{miss}}(s) := \sup_{z \in \mathcal{C}(s)} \inf_{u \in \mathcal{U}(s)} d_Q(z, u), \qquad \varepsilon_{\text{spur}}(s) := \sup_{u \in \mathcal{U}(s)} \inf_{z \in \mathcal{C}(s)} d_Q(u, z). \tag{11}$$

They respectively quantify relevant uncertainty omitted by the learned set and included scenarios with no recourse-equivalent counterpart in the target.

Theorem 2 (Worst-case cost and decision bounds). *For nonempty compact $\mathcal{C}(s)$ and $\mathcal{U}(s)$ with finite directed distortions, suppose the robust minima over $\mathcal{X}$ are attained. Then every $x \in \mathcal{X}$ satisfies*

$$H_{\mathcal{C}}(x, s) \le H_{\mathcal{U}}(x, s) + \varepsilon_{\mathrm{miss}}(s), \tag{12}$$

$$H_{\mathcal{U}}(x, s) \le H_{\mathcal{C}}(x, s) + \varepsilon_{\mathrm{spur}}(s). \tag{13}$$

Consequently, the exact robust solution $x_{\mathcal{U}}$ obeys

$$F_{\mathcal{C}}(x_{\mathcal{U}}, s) - \min_{x} F_{\mathcal{C}}(x, s) \le \varepsilon_{\mathrm{miss}}(s) + \varepsilon_{\mathrm{spur}}(s). \tag{14}$$

The result applies to convex and nonconvex decision and scenario sets.

For the local construction in Section 5.1, we express the decision bound in terms of the selected centers and their radii. Write $r_i(\beta) = \beta d_i(s)$ and let $\mathcal{U}_\beta(s)$ denote equation 5 with multiplier β. Define

$$\Delta_{\mathrm{hole}}(\beta; s) := \sup_{z \in \mathcal{C}(s)} \min_i [\delta_s(z, z_i) - r_i(\beta)]_+,$$

$$m_i(s) := \inf_{z \in \mathcal{C}(s)} \delta_s(z_i, z). \tag{15}$$

Proposition 3 (Finite-sample radius tradeoff). *For $0 \le \beta_1 \le \beta_2$, $\mathcal{U}_{\beta_1}(s) \subseteq \mathcal{U}_{\beta_2}(s)$, and hence*

$$\varepsilon_{\mathrm{miss}}(\beta_2; s) \le \varepsilon_{\mathrm{miss}}(\beta_1; s), \qquad \varepsilon_{\mathrm{spur}}(\beta_2; s) \ge \varepsilon_{\mathrm{spur}}(\beta_1; s). \tag{16}$$

If Q is uniformly $L_{Q,s}$-Lipschitz under δ_s, then the exact robust solution over $\mathcal{U}_\beta(s)$ has target decision-value gap at most

$$L_{Q,s} \left\{ \Delta_{\mathrm{hole}}(\beta; s) + \max_i [m_i(s) + r_i(\beta)] \right\}. \tag{17}$$

The first term is nonincreasing in the radius and is exactly the directed δ_s-distance from the target to the union; the second bounds admitted uncertainty beyond the target and is nondecreasing.

At $\beta = 0$, the uncertainty set consists of the selected historical scenarios. Expanding the local regions tightens the bound when the reduction in the distance from the target support to the union exceeds the increase in the term for additional uncertainty. Appendix A provides the proof and relates the approximation errors to context mismatch and geometric coverage of the target support.

To characterize how these errors change with sample size, Theorem 4 gives the following rate for the target decision-value gap under its stated conditions, with $K_n \to \infty$ and $K_n/n \to 0$:

$$O_p\Big(L_Q \left[(K_n/n)^{1/m_S} + (\log K_n/K_n)^{1/m_Z} + R_n(s) \right] + \delta_{\mathrm{sep},n} + \tau_n \Big), \tag{18}$$

where m_S, m_Z are the exponents in the context-mass and conditional-support assumptions, $R_n(s)$ is the largest physical local-ball radius, $\delta_{\mathrm{sep},n}$ is the terminal AP error, and τ_n is the outer optimization tolerance. Exact separation sets $\delta_{\mathrm{sep},n} = 0$. When $R_n(s) = o_p(1)$, the learned set is conditionally support-consistent. When also $\delta_{\mathrm{sep},n} + \tau_n \to 0$, the target decision-value gap converges to zero in probability.

5.3 Solving the problem within the C&CG framework

The proposed uncertainty set can be incorporated into column-and-constraint generation (C&CG) (Zeng & Zhao, 2013). Let $\mathcal{U}_i(s)$ denote the local set centered at z_i in equation 5. At a C&CG iteration with first-stage decision x_t, the adversarial problem over the proposed uncertainty set takes the form

$$\max_{z \in \mathcal{U}(s)} Q(x_t, z) = \max_{i \in \mathcal{N}_K(s)} \max_{z \in \mathcal{U}_i(s)} Q(x_t, z). \tag{19}$$

We search each local component and retain the feasible candidate with the highest evaluated recourse cost. For linear programming recourse under right-hand-side uncertainty, we develop critical-region-guided projected gradient ascent (CRG-PGA). It uses the active affine pieces of the recourse value function to guide projected ascent and re-solves the recourse problem when a candidate leaves the current piece's validity region. The selected scenario is returned to the master problem. Appendix A.3 provides algorithmic details and benchmark-specific solution methods.

6 EXPERIMENTS

6.1 EXPERIMENTAL SETUP

Benchmarks. We use three benchmarks. Feature-Based Newsvendor uses the Genpact meal-delivery benchmark data supplied with Yang et al. (2026). Robust Portfolio Optimization follows the rolling stock-market benchmark of Chenreddy et al. (2022). For Security-Constrained Unit Commitment (SCUC), we combine the California Test System with CAISO load and renewable trajectories (Estrada-Garcia et al., 2026; California Independent System Operator (CAISO), n.d.). SCUC uses 120 training days and 30 held-out days; uncertainty is a 144-dimensional renewable forecast-residual trajectory, and the first stage contains binary commitment decisions. The newsvendor and portfolio problems are continuous, whereas SCUC has linear recourse after commitment.

Baselines. We compare with **sample average approximation (SAA)**, **RO**, **DRO**, zero-radius **KNN-RO**, unconditional **AGRO**, conditional **AGRO (cAGRO)**, and **contextual KNN global-ball RO (cKNN-GB)**. On the two continuous benchmarks, we also compare with **Sample Robust Optimization with Side Information (SRO-SI)**, **IDCC**, and **causal-transport DRO (CT-DRO)**. For SCUC, we report the size of the SRO-SI formulation and evaluate IDCC in the fixed-decision representation study in Section 6.3. SCUC lies outside the decision classes covered by CT-DRO's published tractable reformulations. Appendix B.2 gives the corresponding formulation details.

Evaluation. We retain each benchmark's original train–test split and evaluate every resulting decision on held-out data with the same settlement model. Mean cost is the main metric for Feature-Based Newsvendor and SCUC. For Robust Portfolio Optimization, we follow the source paper and use value-at-risk (VaR); conditional value-at-risk (CVaR) is also reported. On Portfolio, KNN-RO uses the same context-matched centers as CUB-RO.

Repeated runs and implementation. Main-table results for randomized learned sets and the SCUC CUB-RO separator are averaged over five complete runs. Appendices C.1–C.3 report paired time-block analyses for all three benchmarks, and Appendix C.4 reports dispersion across model seeds. Preprocessing and parameter choices are in Appendix B.2; optimization backends, sensitivity analyses, and implementation details are in Appendices B and E.1.

6.2 HELD-OUT DECISION PERFORMANCE

Table 1: Held-out performance on the two continuous decision problems. Randomized learned-set rows average five complete seeded evaluations; Portfolio entries average 30 year–asset-seed evaluations.

(a) Feature-Based Newsvendor ($n = 92$).

Method	mean ↓	VaR_{95} ↓	CVaR_{95} ↓
SAA	87.04	231.68	472.66
RO	133.61	410.40	667.38
DRO	87.04	231.68	472.66
KNN-RO	90.15	219.20	419.38
SRO-SI	87.66	**215.75**	426.05
cKNN-GB	109.29	319.01	455.13
CT-DRO	87.13	231.06	482.21
IDCC	104.11	216.97	**405.81**
AGRO	93.54	238.77	482.18
cAGRO	92.26	284.86	452.30
CUB-RO	**85.84**	231.60	430.67

(b) Robust Portfolio Optimization.

Method	VaR_{95} ↓	CVaR_{95} ↓
SAA	2.5270%	3.6634%
RO	1.4822%	2.2997%
DRO	1.1763%	1.7552%
KNN-RO	1.1472%	1.7420%
SRO-SI	2.3102%	3.5723%
cKNN-GB	1.0729%	1.5915%
CT-DRO	1.6632%	2.3377%
IDCC	1.5588%	2.4258%
AGRO	1.6359%	2.3986%
cAGRO	1.8877%	2.8005%
CUB-RO	**1.0253%**	**1.5170%**

Feature-Based Newsvendor. Table 1(a) evaluates each order quantity against the demand observed in its test week. Mean held-out loss is therefore the primary measure of realized performance, while VaR and CVaR describe losses in the worst weeks. **CUB-RO** obtains the lowest mean loss, 85.84. Relative to **SAA**, it lowers mean loss by 1.38% and CVaR by 8.88%, with essentially unchanged VaR. The methods with lower tail losses pay for them in the average outcome. Compared with **KNN-RO**, CUB-RO reduces mean loss by 4.78%, while its VaR and CVaR are 5.66% and 2.69% higher. Compared with **SRO-SI**, its mean is 2.08% lower, while VaR and CVaR are 7.35% and 1.08% higher. The contrast is largest for **IDCC**: CUB-RO reduces mean loss by 17.55%, while its VaR and CVaR are 6.74% and 6.13% higher. This is the empirical signature of conservatism in

this benchmark: the lower upper-tail losses of KNN-RO, SRO-SI, and IDCC coincide with higher average realized losses. CUB-RO occupies a less conservative point in this comparison without reverting to the nonrobust result: relative to SAA, it improves both mean loss and CVaR.

Robust Portfolio Optimization. The source benchmark evaluates portfolios by held-out VaR, with CVaR measuring the losses beyond that quantile. **CUB-RO** obtains the lowest values on both metrics. Relative to **SAA**, it reduces VaR by 59.43% and CVaR by 58.59%; the reductions relative to **DRO** are 12.84% and 13.57%. The controlled comparison with **KNN-RO** isolates local expansion: both methods retrieve the same context-matched returns, but KNN-RO retains only the observed points. Expanding those points locally reduces VaR from 1.1472% to 1.0253% and CVaR from 1.7420% to 1.5170%, improvements of 10.63% and 12.92%.

Localization therefore does not weaken tail protection: after restricting the set to neighborhoods of matched returns, CUB-RO records lower held-out tail losses. Relative to **cKNN-GB**, context-conditioned **cAGRO**, and **SRO-SI**, it reduces VaR by 4.44%, 45.69%, and 55.62%, and CVaR by 4.68%, 45.83%, and 57.53%, respectively. On this benchmark, local matched neighborhoods provide more effective tail protection than the corresponding regional, generative, or weighted representations.

Table 2: Held-out SCUC cost on days 420–449, in hundreds of USD per day; mean settled cost is primary. CUB-RO, AGRO, cAGRO, and cKNN-GB are five-seed means.

Method	first stage	realized Q	mean ↓	CVaR_{95} ↓
Perfect information (LB)	—	—	1896.35	—
SAA (10 scenarios)	1886.45	58.75	1945.20	2354.83
RO	1921.38	118.67	2040.05	3086.77
DRO	1917.94	38.30	1956.24	2266.31
KNN-RO	1891.12	55.92	1947.04	2356.33
cKNN-GB	1959.40	31.42	1990.83	2352.15
AGRO	1913.88	51.19	1965.07	2320.23
cAGRO	1917.86	53.57	1971.43	2429.33
CUB-RO	1899.30	36.44	**1935.74**	**2260.25**

All decisions are settled with the same held-out SCUC model. Appendix B reports the optimization backends and search budgets; Appendix D reports computational costs and stopping states. The SCUC SAA row follows the project protocol with ten K-means representatives of the 120 training residual trajectories.

Security-Constrained Unit Commitment. Table 2 settles every commitment against the same realized renewable trajectory; costs are reported in hundreds of USD per day. **CUB-RO** obtains the lowest mean cost, 1935.74, or $193,574 per day, and the lowest CVaR_{95}, 2260.25. Relative to **SAA**, it reduces total cost by $946 per day, or 0.49%, as lower realized recourse more than offsets the higher first-stage cost. The absolute margin is small because SAA is already only $4,885 per day above the perfect-information lower bound; CUB-RO recovers 19.36% of this remaining improvement space. Compared with zero-radius **KNN-RO**, mean cost is 0.58% lower, or $1,130 per day. Relative to **AGRO**, CUB-RO saves $1,458 per day in the first stage and $1,475 in realized recourse, lowering total cost by $2,933, or 1.49%. The corresponding savings over context-conditioned **cAGRO** are $1,856, $1,713, and $3,569 per day, for a 1.81% lower total. In both cases, higher commitment cost does not translate into lower realized recourse. Regional **cKNN-GB** exhibits a different form of conservatism: it spends $6,011 more per day in the first stage to save only $502 in recourse, so CUB-RO lowers its total by $5,509, or 2.77%. These differences motivate examining which adversarial directions the conditional representations admit.

Separation accuracy on SCUC. In a separate diagnostic on 150 previously saved CUB-RO decisions, evaluating the separator yields a mean separation-error upper bound of 0.54% of the evaluated objective, including first-stage cost. Appendix D.2 gives the fixed-decision evaluation and normalization procedure; this diagnostic does not evaluate the main-table decisions in Table 2.

6.3 Conditional uncertainty representations under adversarial querying

Adversarial scenarios versus realized outcomes. Figure 1 compares the uncertainty-set representations and the recourse costs of the returned adversarial scenarios. To examine how these scenarios

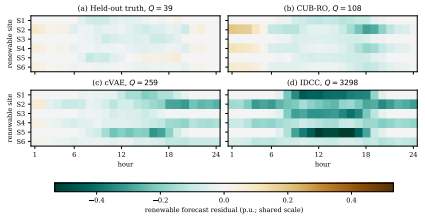

Figure 2: Fixed-decision SCUC residuals; titles report original-LP recourse for the returned cVAE and IDCC scenarios.

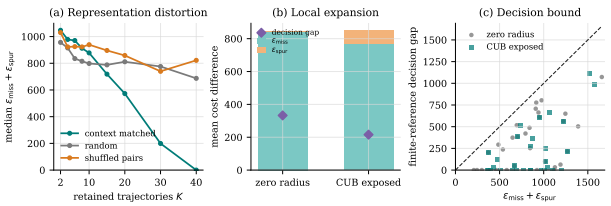

Figure 3: SCUC finite-reference diagnostics: context matching, local-ball query scenarios, and the finite-library decision bound.

differ from the realized trajectory, Figure 2 displays their renewable forecast residuals across sites and hours under the same context and fixed first-stage decision. The IDCC scenario exhibits substantially more negative residuals than the realization, particularly around midday at sites S1, S3, and S5, implying much lower renewable generation over those periods. The cVAE scenario also exhibits larger negative deviations at several sites, whereas the CUB-RO scenario remains closer to the realized residual pattern. Appendix F.4 specifies the shared first-stage decision, uncertainty-set configurations, and scenario-search procedures.

Context matching and local expansion. Figure 3(a) compares context-matched, random, and shuffled-pair selections against a reference set of the 40 weather-nearest historical trajectories. As more trajectories are retained, the median total distortion decreases more sharply for context-matched selection, falling below both controls at 15, 20, and 30 trajectories. Figure 3(b) then fixes the ten matched centers and compares the centers alone with a finite set that also includes scenarios returned by local-ball queries. Adding these scenarios reduces mean missing distortion from 841.20 to 771.02 and the mean finite-reference decision gap from 332.70 to 215.79, while introducing mean spurious distortion of 80.12. These results show how the two construction steps contribute: context matching improves the representation of the reference uncertainty as more centers are retained, while local expansion supplements the selected centers with nearby scenarios and lowers the average decision gap.

Uncertainty-set approximation and decision quality. Figure 3(c) plots the finite-reference decision gap against the sum of missing and spurious distortions. For both the zero-radius representation and the locally expanded finite scenario set, all evaluated gaps lie below the corresponding bounds, consistent with the analysis in Section 5.2. The two distortion terms also explain the tradeoff in Figure 3(b): local expansion reduces omitted uncertainty but introduces additional uncertainty, so the mean bound increases slightly even as the mean decision gap decreases. The bound therefore limits the decision gap without requiring its changes to track improvements in decision quality. Appendices F.2 and F.3 provide the reference-set construction, finite decision class, and evaluation procedures; Appendix F.4 reports the supplementary IDCC analysis.

7 Conclusion

This work learns context-dependent uncertainty sets from paired historical data. CUB-RO retrieves matched trajectories and constructs local regions around them. For the convex losses considered here, convexification preserves the worst-case value. Our analysis bounds the effects of omitted and introduced uncertainty on decision quality and characterizes the finite-sample tradeoff associated with local radii. Across three benchmarks, the learned local representation improves Portfolio tail risk and attains the lowest mean realized cost on Feature-Based Newsvendor and SCUC. Fixed-decision diagnostics show that adversarial searches return substantially higher recourse costs in unified conditional regions than in the locally anchored representation, supporting context matching during set construction.

AI USE STATEMENT

In this work, we used generative AI tools to assist with language polishing and manuscript formatting. The authors performed the final data checks and reference verification manually. We take responsibility for the final content of this work, including text and artifacts prepared with the assistance of generative AI.

REFERENCES

Gah-Yi Ban and Cynthia Rudin. The big data newsvendor: practical insights from machine learning. *Operations Research*, 67(1):90–108, 2019.

A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Transactions on Power Systems*, 28(1): 52–63, 2013.

Dimitris Bertsimas and Cheol Woo Kim. A machine learning approach to two-stage adaptive robust optimization. *European Journal of Operational Research*, 319(1):16–30, 2024.

Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

Dimitris Bertsimas, Dick den Hertog, and Jean Pauphilet. Probabilistic guarantees in robust optimization. *SIAM Journal on Optimization*, 31(4):2893–2920, 2021.

Dimitris Bertsimas, Christopher McCord, and Bradley Sturt. Dynamic optimization with side information. *European Journal of Operational Research*, 304(2):634–651, 2023a.

Dimitris Bertsimas, Shimrit Shtern, and Bradley Sturt. A data-driven approach to multistage stochastic linear optimization. *Management Science*, 69(1):51–74, 2023b.

Aron Brenner, Rahman Khorramfar, Jennifer Sun, and Saurabh Amin. A deep generative learning approach for two-stage adaptive robust optimization. In *International Conference on Learning Representations*, 2025.

California Independent System Operator (CAISO). *Business Practice Manual for Market Operations, Version 106*, 2026. URL `https://bpmcm.caiso.com/Pages/BPMDetails.aspx?BPM=Market+Operations`.

California Independent System Operator (CAISO). Open access same-time information system (OASIS), n.d. URL `https://oasis.caiso.com/`.

M. Carrion and J. M. Arroyo. A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem. *IEEE Transactions on Power Systems*, 21(3):1371–1378, 2006.

Abhilash Reddy Chenreddy, Nymisha Bandi, and Erick Delage. Data-driven conditional robust optimization. In *Advances in Neural Information Processing Systems*, volume 35, pp. 9525–9537. Curran Associates, Inc., 2022.

Justin Dumouchelle, Esther Julien, Jannis Kurtz, and Elias Khalil. Neur2RO: Neural two-stage robust optimization. In *International Conference on Learning Representations*, 2024.

Adrián Esteban-Pérez and Juan M. Morales. Distributionally robust stochastic programs with side information based on trimmings. *Mathematical Programming*, 195(1):1069–1105, 2022.

Juan-Alberto Estrada-Garcia, Ruiwei Jiang, and Alexandre Moreira. Dynamic transmission line switching amid wildfire-prone weather under decision-dependent uncertainty. *INFORMS Journal on Computing*, 2026. Articles in Advance.

Matteo Fischetti and Jason Jo. Deep neural networks and mixed integer linear optimization. *Constraints*, 23(3):296–309, 2018.

Marc Goerigk and Jannis Kurtz. Data-driven robust optimization using deep neural networks. *Computers & Operations Research*, 151:106087, 2023.

L. Jeff Hong, Zhiyuan Huang, and Henry Lam. Learning-based robust optimization: Procedures and statistical guarantees. *Management Science*, 67(6):3447–3467, 2021.

Chancellor Johnstone and Bruce Cox. Conformal uncertainty sets for robust optimization. In *Proceedings of the Tenth Symposium on Conformal and Probabilistic Prediction and Applications*, volume 152 of *Proceedings of Machine Learning Research*, pp. 72–90. PMLR, 2021.

Rohit Kannan, Güzin Bayraksan, and James R. Luedtke. Residuals-based distributionally robust optimization with covariate information. *Mathematical Programming*, 207(1):369–425, 2024.

Daniel Kuhn, Wolfram Wiesemann, and Angelos Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

Viet Anh Nguyen, Fan Zhang, Shanshan Wang, Jose Blanchet, Erick Delage, and Yinyu Ye. Robustifying conditional portfolio decisions via optimal transport. *Operations Research*, 73(5): 2801–2829, 2025.

Shunichi Ohmori. A predictive prescription using minimum volume K-nearest neighbor enclosing ellipsoid and robust optimization. *Mathematics*, 9(2):119, 2021.

Yash P. Patel, Sahana Rayan, and Ambuj Tewari. Conformal contextual robust optimization. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 2485–2493. PMLR, 2024.

Said Rahal, Zukui Li, and Dimitri J. Papageorgiou. Deep lifted decision rules for two-stage adaptive optimization problems. *Computers & Chemical Engineering*, 159:107661, 2022.

Utsav Sadana, Abhilash Chenreddy, Erick Delage, Alexandre Forel, Emma Frejinger, and Thibaut Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025.

Chao Shang, Xiaolin Huang, and Fengqi You. Data-driven robust optimization based on kernel learning. *Computers & Chemical Engineering*, 106:464–479, 2017.

Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.

Chunlin Sun, Linyu Liu, and Xiaocheng Li. Predict-then-calibrate: A new perspective of robust contextual LP. In *Advances in Neural Information Processing Systems*, volume 36, pp. 17713–17741. Curran Associates, Inc., 2023.

Theja Tulabandhula and Cynthia Rudin. Robust optimization using machine learning for uncertainty sets. *arXiv preprint arXiv:1407.1097*, 2014.

Jincheng Yang, Luhao Zhang, Ningyuan Chen, Rui Gao, and Ming Hu. Decision making with side information: A causal transport robust approach. *Operations Research*, 74(4):2249–2269, 2026.

Bo Zeng and Long Zhao. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461, 2013.

Luhao Zhang, Jincheng Yang, and Rui Gao. Optimal robust policy for feature-based newsvendor. *Management Science*, 70(4):2315–2329, 2024.

Long Zhao and Bo Zeng. An exact algorithm for two-stage robust optimization with mixed integer recourse problems. Technical report, Optimization Online, 2012. URL `https://optimization-online.org/2012/01/3310/`.

A SUPPLEMENTARY METHOD AND THEORY

A.1 REPRESENTATION GUARANTEES AND OUTER C&CG BOUNDS

For a metric d, write $h_d(A,B) = \sup_{a\in A}\inf_{b\in B} d(a,b)$ and $d_{H,d}(A,B) = \max\{h_d(A,B), h_d(B,A)\}$. The Euclidean versions omit the subscript. Let $I_n(s) = \mathcal{N}_{K_n}(s)$ and let $\mathcal{U}_{n,\beta}(s)$ be equation 5 with $K = K_n$ and multiplier β. Because $\delta_s(z,z') = \|B_s^{-1}(z-z')\|_2$, its ith component is exactly the δ_s-ball of radius $r_i(\beta) = \beta d_i(s)$ centered at z_i. Consequently,

$$h_{\delta_s}(\mathcal{C}(s), \mathcal{U}_{n,\beta}(s)) = \sup_{z\in\mathcal{C}(s)} \min_{i\in I_n(s)} [\delta_s(z,z_i) - r_i(\beta)]_+ = \Delta_{\text{hole}}(\beta; s). \tag{20}$$

For the growing-sample analysis, let d_S be a metric on the context representation space that induces the same neighbor ordering as Δ_S, with ties resolved by training order. Contexts S_i, s and the conditional support $\mathcal{C}(s)$ in this analysis refer to this representation. Squared Euclidean dissimilarities, including positive coordinate weights, induce the metric $d_S = \sqrt{\Delta_S}$. Cosine dissimilarity on nonzero context vectors induces Euclidean distance on their unit-normalized representations, $d_S = \sqrt{2\Delta_S}$. The local-mass and support-regularity conditions below are stated with respect to d_S. Define the Euclidean scenario-space quantities

$$\Delta_{\text{cover,n}}(s) := \sup_{z\in\mathcal{C}(s)} \min_{i\in I_n(s)} \|z - z_i\|_2, \tag{21}$$

$$\Delta_{\text{match,n}}(s) := \max_{i\in I_n(s)} h(\mathcal{C}(s_i), \mathcal{C}(s)), \tag{22}$$

$$R_n(s) := \max_{i\in I_n(s)} \|B_s\|_{\text{op}} \widehat{\beta} d_i(s). \tag{23}$$

If $z_i \in \mathcal{C}(s_i)$ for every selected observation, these quantities give the radius-insensitive but asymptotically convenient bounds

$$h(\mathcal{C}(s), \mathcal{U}_{n,\beta}(s)) \le \Delta_{\text{cover,n}}(s), \qquad h(\mathcal{U}_{n,\beta}(s), \mathcal{C}(s)) \le \Delta_{\text{match,n}}(s) + R_n(s). \tag{24}$$

For the growing-sample analysis, equation 24 provides a coarser radius-insensitive cover term and an explicit physical-radius term.

Theorem 4 (Conditional support consistency)**.** *Let $(S_i, Z_i)_{i=1}^n$ be independent draws with conditional support $\mathcal{C}(s') = \text{supp}(Z \mid S = s')$. Assume locally around s: (i) $\Pr\{d_S(S,s) \le r\} \ge c_S r^{m_S}$; (ii) the support map satisfies $d_H(\mathcal{C}(s'), \mathcal{C}(s'')) \le L_S d_S(s', s'')$ for nearby s', s''; (iii) $\mathcal{C}(s)$ has covering number at most $C_Z r^{-m_Z}$ and, for all nearby s', all $z \in \mathcal{C}(s)$, and all sufficiently small r,*

$$\Pr\{\|Z - z\|_2 \le r + L_S d_S(s', s) \mid S = s'\} \ge c_Z r^{m_Z}. \tag{25}$$

If $K_n \to \infty$ and $K_n/n \to 0$, then

$$d_H(\mathcal{U}_n(s), \mathcal{C}(s)) = O_p\Big((K_n/n)^{1/m_S} + (\log K_n/K_n)^{1/m_Z} + R_n(s)\Big). \tag{26}$$

If $R_n(s) = o_p(1)$, the set is conditionally support-consistent. If, in addition, Q is uniformly L_Q-Lipschitz and the conditions of Theorem 7 hold for each n, with terminal separation error $\delta_{\text{sep},n}$ and outer tolerance τ_n, then

$$\begin{aligned} F_{\mathcal{C}}(\hat{x}_n, s) - \min_x F_{\mathcal{C}}(x,s) = O_p\big(L_Q[(K_n/n)^{1/m_S} + (\log K_n/K_n)^{1/m_Z} + R_n(s)] \\ + \delta_{\text{sep},n} + \tau_n\big). \end{aligned} \tag{27}$$

Proposition 5 (Comparison of worst-case recourse costs)**.** *Fix x and let $Q(x,z) = \max_{p\in\mathcal{P}}\{a_p(x) + b_p^\top z\}$, with finite $\mathcal{P}$. For any compact set A, let $\sigma_A(b) = \max_{z\in A} b^\top z$. Then*

$$H_A(x) = \max_{p\in\mathcal{P}}\{a_p(x) + \sigma_A(b_p)\}. \tag{28}$$

For a locally anchored set $\mathcal{U}(s)$ and any unified conditional region $\mathcal{G}(s)$, if some recourse piece p satisfies

$$a_p(x) + \sigma_{\mathcal{G}(s)}(b_p) > H_{\mathcal{U}}(x,s) + \varepsilon, \tag{29}$$

then every ε-accurate separator on $\mathcal{G}(s)$ returns a value satisfying

$$\widehat{H}_{\mathcal{G}} > H_{\mathcal{U}}(x,s). \tag{30}$$

The relevant AP mechanism is extra support in a direction exposed by an actual recourse piece, irrespective of set containment or average fit. The comparison concerns worst-case recourse costs at a fixed decision; held-out performance is evaluated separately using the decision induced by each representation.

At iteration t, the learned scenarios enter the restricted master

$$\min_{x\in\mathcal{X},\eta} c^\top x+\eta \quad \text{s.t.} \quad Q(x,z_j)\le\eta, \quad z_j\in\mathcal{S}_t. \tag{31}$$

For an exact AP call at outer iteration t, let

$$L_t := c^\top x_t+\eta_t, \qquad A_t := c^\top x_t+Q(x_t,\widehat{z}_t), \tag{32}$$

where L_t is the restricted-master value and $\widehat{z}_t \in \arg\max_{z\in\mathcal{U}(s)} Q(x_t,z)$. Define $U_t := \min_{0\le j\le t} A_j$.

Proposition 6 (Outer C&CG bounds)**.** *If the restricted master and AP are solved exactly, then $L_t\le v_{\mathcal{U}}\le U_t$. Consequently, $0\le U_t-L_t\le\tau$ certifies a τ-optimal incumbent.*

Proof. The restricted master relaxes the robust problem, so $L_t\le v_{\mathcal{U}}$. Exact AP solution makes every A_j the robust objective of a feasible first-stage decision, hence $v_{\mathcal{U}}\le U_t$. The stated certificate follows immediately. □

Theorem 7 (Combined representation and approximate separation)**.** *Suppose the restricted master is solved globally and the AP returns a feasible $\hat{z}\in\mathcal{U}(s)$. Define the separation error*

$$\delta_{\mathrm{sep}}(\hat{x},s) := H_{\mathcal{U}}(\hat{x},s)-Q(\hat{x},\hat{z})\ge 0. \tag{33}$$

If the algorithm stops when $Q(\hat{x},\hat{z})\le\hat{\eta}+\tau$, then

$$F_{\mathcal{C}}(\hat{x},s)-\min_{x\in\mathcal{X}} F_{\mathcal{C}}(x,s)\le\varepsilon_{\mathrm{miss}}(s)+\varepsilon_{\mathrm{spur}}(s)+\delta_{\mathrm{sep}}(\hat{x},s)+\tau. \tag{34}$$

Exact separation is the special case $\delta_{\mathrm{sep}}=0$. CUB-RO uses this case on Feature-Based Newsvendor and Robust Portfolio Optimization. For SCUC, Section D.2 quantifies δ_{sep} through fixed-decision evaluations of the multi-start separator.

A.2 Convex-hull invariance and support geometry

Proof of Proposition 1. Every component in equation 5 is compact, so their finite union is compact and continuity of $Q(x,\cdot)$ gives an attaining $z^\star\in\mathcal{U}(s)$. For any $z=\sum_j\lambda_j u_j\in\mathrm{conv}(\mathcal{U}(s))$, convexity gives

$$Q(x,z)\le\sum_j\lambda_j Q(x,u_j)\le\max_{u\in\mathcal{U}(s)} Q(x,u).$$

Together with $\mathcal{U}(s)\subseteq\mathrm{conv}(\mathcal{U}(s))$, this proves equation 7; the original-union maximizer $z^\star$ therefore also attains the convex-hull value. For the upper-envelope form, interchange the finite maxima over pieces and components and apply the ellipsoid support formula to obtain equation 8. □

For an upper-envelope LP recourse $Q(x,z)=\max_{p\in\mathcal{P}}\{a_p(x)+b_p^\top z\}$, this result can be written directly in terms of support functions:

$$\sup_{z\in\mathcal{U}} Q(x,z)=\max_{p\in\mathcal{P}}\{a_p(x)+\sigma_{\mathcal{U}}(b_p)\}, \qquad \sigma_{\mathcal{U}}=\sigma_{\mathrm{conv}(\mathcal{U})}. \tag{35}$$

For the union in equation 5, writing $r_i=\widehat{\beta}d_i(s)$, the support function is

$$\sigma_{\mathcal{U}(s)}(b)=\max_{i\in\mathcal{N}_K(s)}\left\{b^\top z_i+r_i\|B_s^\top b\|_2\right\}. \tag{36}$$

Hence its union and convex hull have the same robust value for convex scenario-wise recourse, while the AP over the union still returns a scenario belonging to one of its locally anchored components. A larger unified set can nevertheless increase the robust value by extending the support boundary beyond that convex hull in a recourse-relevant direction. Feature-Based Newsvendor has convex piecewise-linear scenario loss, Robust Portfolio Optimization has affine scenario loss, and the SCUC recourse used here has the upper-envelope form, so this interpretation applies to all three reported benchmarks.

A.3 BENCHMARK-SPECIFIC ADVERSARIAL SOLVERS

The three benchmarks use solver interfaces matched to their value functions. Feature-Based Newsvendor maximizes a scalar convex piecewise-linear loss over a union of intervals, so it evaluates the two endpoints of every interval. Robust Portfolio Optimization has affine scenario loss, so every ellipsoidal component is queried by its closed-form support function; in the main experiments, the outer problem is solved as a convex conic program. These two CUB-RO adversarial problems are solved exactly.

SCUC has fixed-objective RHS-parametric LP recourse,

$$Q(x, z) = \min_{y \geq 0} \{q^\top y : Wy \geq h - Tx - Dz\}, \tag{37}$$

whose dual solutions expose affine pieces

$$q_p(x, z) = a_p(x) + b_p^\top z, \qquad Q(x, z) = \max_{p \in \mathcal{P}} q_p(x, z). \tag{38}$$

For SCUC, the two projected-gradient methods are standard projected gradient ascent (PGA) and our critical-region-guided projected gradient ascent (CRG-PGA). For fixed x, each affine piece is active on the region where $q_p(x, z) = Q(x, z)$. An exact recourse LP query supplies both the true value and the gradient b_p of the currently active affine piece. CRG-PGA follows that piece through the uncertainty-set parameterization, projects every step back to its feasible search domain, and re-solves the original recourse LP whenever the proposed scenario leaves the current piece. The newly active piece then guides the next segment. After the initial searches, an independently initialized trajectory explores additional portions of the set; newly discovered pieces are added to the search queue. The returned scenario is the best feasible candidate after a final evaluation by the original recourse LP. For CUB-RO, this procedure is run on each selected affine ball; in the solver comparison, it is run in the latent balls of AGRO and cAGRO. The search operates in parameter space and therefore also applies when a neural decoder image is nonconvex. The step budget J applies separately to each queued piece–start pair and to the additional trajectory. Discovering another piece can enqueue a new search task.

Algorithm 1 CRG-PGA for the SCUC adversarial problem

Require: Decision x; feasible parameterization $z = G_s(w)$ and projection Π_s; recourse LP; M initial starts; step budget J

Solve the recourse LP at M feasible starts; initialize the incumbent, store active pieces, and enqueue piece–start pairs.
while the search queue is nonempty **do**
  Remove a piece–start pair (p, w) from the queue.
  for $j = 1, \ldots, J$ **do**
    Update the incumbent using $q_p(x, G_s(w))$ and take a projected ascent step to w'; stop this task if $w' = w$.
    if $G_s(w')$ leaves the validity region of piece p **then**
      Solve the recourse LP at $G_s(w')$, update the incumbent, and enqueue the resulting piece–start pair within the storage limits; end this task.
    else
      Set $w \leftarrow w'$.
    end if
  end for
end while
Start an independent trajectory with a recourse LP solve; take at most J projected steps, querying the LP when the active piece changes; update the incumbent and queue the resulting pairs.
Process the queued pairs as above.
Re-evaluate the incumbent with the original recourse LP and return it.

CRG-PGA is a multi-start nonconvex search. Its restart guarantee is stated directly. For $\epsilon \geq 0$, let $\mathcal{B}_\epsilon(x, s)$ be the feasible initial states from which one CRG-PGA trajectory returns a scenario $\hat{z}$ satisfying

$$Q(x, \hat{z}) \geq \sup_{z \in \mathcal{U}(s)} Q(x, z) - \epsilon,$$

and let $\mu_\epsilon(x,s) = \nu_s(\mathcal{B}_\epsilon(x,s))$ under the restart distribution ν_s.

Proposition 8 (Multi-start recovery)**.** *If $\mu_\epsilon(x,s) > 0$, then M independent CRG-PGA starts return an ϵ-global AP solution with probability at least*

$$1 - (1 - \mu_\epsilon(x,s))^M,$$

which converges to one as $M \to \infty$.

Proof. One start fails to enter $\mathcal{B}_\epsilon(x,s)$ with probability $1-\mu_\epsilon(x,s)$. Independence gives probability $(1-\mu_\epsilon(x,s))^M$ that all starts fail. The complementary event contains at least one ϵ-global trajectory, and CRG-PGA returns the best LP-validated candidate across starts. □

In particular, setting $\epsilon = 0$ gives probabilistic global convergence whenever the basin of a global AP maximizer has positive restart mass.

Corollary 9 (Probabilistic approximate end-to-end bound)**.** *Fix s and $\mathcal{U}(s)$. Suppose C&CG makes at most T SCUC AP calls for a fixed T, solves each restricted master globally, and uses M starts drawn independently from ν_s and independently of the preceding search history at each call. If $\inf_{x\in\mathcal{X}} \mu_\epsilon(x,s) \geq \underline{\mu}_\epsilon > 0$, then, with probability at least $\max\{0, 1 - T(1-\underline{\mu}_\epsilon)^M\}$, every queried restricted-master solution $(\hat{x},\hat{\eta})$ whose returned feasible AP scenario satisfies $Q(\hat{x},\hat{z}) \leq \hat{\eta} + \tau$ obeys*

$$F_{\mathcal{C}}(\hat{x},s) - \min_x F_{\mathcal{C}}(x,s) \leq \varepsilon_{\mathrm{miss}}(s) + \varepsilon_{\mathrm{spur}}(s) + \epsilon + \tau. \tag{39}$$

Proof. Let E_t be the event that call t is made and its separation error exceeds ϵ. Conditional on the history before that call, its decision is fixed, so Proposition 8 and the uniform mass bound give a failure probability at most $q = (1-\underline{\mu}_\epsilon)^M$. Thus $\Pr(E_t) \leq q$, and a union bound gives $\Pr(\bigcup_{t=1}^T E_t) \leq Tq$. On the complementary event, every executed call has separation error at most ϵ. Applying Theorem 7 to any queried iterate satisfying the stopping condition proves the claim. □

The proposition characterizes recovery as the restart count increases. The SCUC experiments report their initialization counts and step budgets in Section B.

A.4 Proofs and supporting corollaries

Proof of Theorem 2. Fix x and $z \in \mathcal{C}(s)$. For every $\delta > 0$, there is $u \in \mathcal{U}(s)$ satisfying $d_Q(z,u) \leq \varepsilon_{\mathrm{miss}}(s) + \delta$. Hence

$$Q(x,z) \leq H_{\mathcal{U}}(x,s) + \varepsilon_{\mathrm{miss}}(s) + \delta.$$

Taking the supremum over z and letting $\delta \downarrow 0$ proves equation 12. Interchanging the sets proves equation 13. Therefore

$$\begin{aligned} F_{\mathcal{C}}(x_{\mathcal{U}},s) &\leq F_{\mathcal{U}}(x_{\mathcal{U}},s) + \varepsilon_{\mathrm{miss}}(s) \\ &\leq F_{\mathcal{U}}(x_{\mathcal{C}},s) + \varepsilon_{\mathrm{miss}}(s) \\ &\leq F_{\mathcal{C}}(x_{\mathcal{C}},s) + \varepsilon_{\mathrm{miss}}(s) + \varepsilon_{\mathrm{spur}}(s). \end{aligned}$$

□

Corollary 10 (Ambient Hausdorff bound)**.** *If $|Q(x,z) - Q(x,z')| \leq L_Q\|z - z'\|$ uniformly over x, the two distortions are bounded by L_Q times their corresponding directed Hausdorff distances.*

Proof. Substitute the Lipschitz inequality into equation 10 and equation 11. □

Corollary 11 (Recourse-induced support discrepancy)**.** *Suppose $Q(x,z) = \max_{\pi\in\Pi(x)}\{a_\pi(x)^\top z + b_\pi(x)\}$. Then*

$$H_{\mathcal{A}}(x,s) = \max_{\pi\in\Pi(x)}\{b_\pi(x) + \sigma_{\mathcal{A}(s)}(a_\pi(x))\},$$

and

$$|H_{\mathcal{A}}(x,s) - H_{\mathcal{B}}(x,s)| \leq \sup_{\pi\in\Pi(x)} |\sigma_{\mathcal{A}(s)}(a_\pi(x)) - \sigma_{\mathcal{B}(s)}(a_\pi(x))|.$$

Proof. Interchange the finite piece maximum and set supremum, then use $|\max_i r_i - \max_i t_i| \le \max_i |r_i - t_i|$. □

Proof of Proposition 3. Increasing β expands every component, so $\mathcal{U}_{\beta_1}(s) \subseteq \mathcal{U}_{\beta_2}(s)$. Taking an infimum over a larger approximating set cannot increase the target-to-set directed distortion, while taking a supremum over a larger admitted set cannot decrease the set-to-target distortion. This proves equation 16 directly for the recourse pseudometric d_Q.

In the δ_s metric, the distance from z to the ith component is exactly $[\delta_s(z, z_i) - r_i(\beta)]_+$. Taking the minimum over components and supremum over $z \in \mathcal{C}(s)$ proves equation 20. Conversely, for any $u \in \mathcal{U}_\beta(s)$, some i satisfies $\delta_s(u, z_i) \le r_i(\beta)$, and the triangle inequality gives

$$\inf_{z \in \mathcal{C}(s)} \delta_s(u, z) \le r_i(\beta) + m_i(s) \le \max_j \{r_j(\beta) + m_j(s)\}.$$

Uniform $L_{Q,s}$-Lipschitz continuity transfers these two directed distances to $\varepsilon_{\text{miss}}$ and $\varepsilon_{\text{spur}}$; applying Theorem 2 proves equation 17.

For completeness, every selected center belongs to the union, so its Euclidean target-to-set distance is at most $\Delta_{\text{cover,n}}$. If $z_i \in \mathcal{C}(s_i)$, the same triangle argument in Euclidean norm bounds the reverse distance by $\Delta_{\text{match,n}} + R_n$. This proves the coarse bounds in equation 24 used by the consistency theorem. □

Proof of Theorem 4. Let $\rho_n(s) = \max_{i \in I_n(s)} d_S(S_i, s)$ be the K_nth-neighbor context radius under d_S, using the neighbor ordering in equation 4. The local context-mass assumption and a binomial concentration bound give

$$\rho_n(s) = O_p\Big((K_n/n)^{1/m_S}\Big).$$

The Hausdorff-Lipschitz support map then gives $\Delta_{\text{match,n}}(s) \le L_S \rho_n(s)$.

Condition on the selected contexts and cover $\mathcal{C}(s)$ by at most $C_Z r^{-m_Z}$ balls of radius r. Uniform conditional lower mass implies that the probability that any one covering ball has no corresponding selected outcome is at most $\exp(-cK_n r^{m_Z})$, after enlarging its radius by $L_S \rho_n(s)$ to account for support movement. A union bound over the cover, with r a sufficiently large constant multiple of $(\log K_n/K_n)^{1/m_Z}$, yields

$$\Delta_{\text{cover,n}}(s) = O_p\Big(\rho_n(s) + (\log K_n/K_n)^{1/m_Z}\Big).$$

Proposition 3 now bounds the two directed Hausdorff distances, proving equation 26. The consistency claim follows when $R_n(s) = o_p(1)$. Finally combine the Lipschitz implication, Proposition 3, and Theorem 7 to obtain equation 27. □

Proof of Theorem 7. Let $v_{\mathcal{U}} = \min_x F_{\mathcal{U}}(x, s)$ and let $\widehat{L} = c^\top \hat{x} + \hat{\eta}$ be the final restricted-master value. Because the restricted master contains only a subset of scenarios, $\widehat{L} \le v_{\mathcal{U}}$. Feasibility of $\hat{z}$, the definition of δ_{sep}, and the stopping rule give

$$F_{\mathcal{U}}(\hat{x}, s) = c^\top \hat{x} + Q(\hat{x}, \hat{z}) + \delta_{\text{sep}}(\hat{x}, s) \le \widehat{L} + \tau + \delta_{\text{sep}}(\hat{x}, s) \le v_{\mathcal{U}} + \tau + \delta_{\text{sep}}(\hat{x}, s).$$

Combining this with the two directed inequalities in Theorem 2 proves equation 34. □

Corollary 12 (Modular decision consistency)**.** *Under the conditions of Theorem 7 for each n, if $\varepsilon_{\text{miss},n} + \varepsilon_{\text{spur},n} + \delta_{\text{sep},n} + \tau_n \to 0$, the decision-value gap relative to the target conditional problem converges to zero.*

Proof of Proposition 5. Interchanging the two finite maxima gives $H_A(x) = \max_p \{a_p(x) + \sigma_A(b_p)\}$. If equation 29 holds, then $H_{\mathcal{G}}(x, s) > H_{\mathcal{U}}(x, s) + \varepsilon$. Separator accuracy therefore gives $\widehat{H}_{\mathcal{G}} \ge H_{\mathcal{G}} - \varepsilon > H_{\mathcal{U}}(x, s)$. □

B BENCHMARK MODELS AND EXPERIMENTAL PROTOCOL

B.1 BENCHMARK DATA AND OPTIMIZATION MODELS

Security-Constrained Unit Commitment. The transmission network is the Northern-California partition of the California Test System, with thermal-generator and power-transfer distribution factor (PTDF) data from the accompanying materials of Estrada-Garcia et al. (2026). Load and renewable forecast/realization trajectories follow the California Independent System Operator Open Access Same-time Information System (CAISO OASIS) data (California Independent System Operator (CAISO), n.d.). The model is a cost-based day-ahead SCUC (Carrion & Arroyo, 2006; Bertsimas et al., 2013) aligned with the CAISO Market Operations Business Practice Manual (California Independent System Operator (CAISO), 2026). Each uncertainty observation is a 24-hour, six-site renewable forecast-residual trajectory in $\mathbb{R}^{144}$. We train on days 300–419 and settle the saved first-stage decisions on days 420–449. Power is represented in per unit with 1 p.u. equal to 100 MW and each period is one hour. Consequently, one reported SCUC cost unit equals 100 USD; all table entries are daily costs. The perfect-information lower bound re-solves the same model separately for each held-out day with that day's renewable deviation revealed before commitment and is reported as a reference lower bound.

Here is the optimization model used in the common problem interface. Let $g \in \mathcal{G}$, $t \in \mathcal{T}$, $b \in \mathcal{B}$, $\ell \in \mathcal{L}$, and $j \in \mathcal{J}$ index generators, hours, buses, lines, and renewable sites. The known day-ahead quantities are system load D_t and renewable forecast $\widehat{r}_{tj}$; the uncertainty z_{tj} is the subsequent renewable deviation. Let H be the PTDF matrix, M_G and M_R map generators and renewable sites to buses, and define $H_G = HM_G$ and $H_R = HM_R$. The first-stage variables are commitment $u_{tg} \in \{0,1\}$, bounded startup/shutdown indicators $v_{tg}, w_{tg} \in [0,1]$, base generation above minimum $0 \le \delta_{tg} \le (\overline{P}_g - \underline{P}_g)u_{tg}$, energy-cost epigraph $e_{tg} \ge 0$, and nonnegative base slacks $(\ell^0, o^0, c^0, \phi^{0,+}, \phi^{0,-})$. Write $p^0_{tg} = \underline{P}_g u_{tg} + \delta_{tg}$. Their feasible region $\mathcal{X}(D, \widehat{r})$ is defined by

$$u_{tg} - u_{t-1,g} = v_{tg} - w_{tg}, \tag{40}$$

$$\sum_{k=t-U_g+1}^{t} v_{kg} \le u_{tg}, \qquad \sum_{k=t-D_g+1}^{t} w_{kg} \le 1 - u_{tg}, \tag{41}$$

$$e_{tg} \ge a_{gk}\delta_{tg} + b_{gk}, \quad k = 1, \dots, K_g, \tag{42}$$

$$\sum_g p^0_{tg} - \mathbf{1}^\top c^0_t - \mathbf{1}^\top o^0_t + \mathbf{1}^\top \ell^0_t = D_t - \mathbf{1}^\top \widehat{r}_t, \tag{43}$$

$$\underline{F} - \phi^{0,-}_t \le f^0_t \le \overline{F} + \phi^{0,+}_t, \tag{44}$$

$$p^0_{tg} - p^0_{t-1,g} \le 60\rho_g u_{t-1,g} + (\underline{P}_g + 30\rho_g)v_{tg}, \tag{45}$$

$$p^0_{t-1,g} - p^0_{tg} \le 60\rho_g u_{tg} + (\underline{P}_g + 30\rho_g)w_{tg}, \tag{46}$$

where the minimum-up/down sums apply only when their index windows are available, $v_{1g} \ge u_{1g}$ implements the daily cold start, and

$$f^0_t = H_G p^0_t + H\ell^0_t - H_R c^0_t - Ho^0_t + H(M_R \widehat{r}_t - D_t\lambda). \tag{47}$$

Here λ is the nodal load-distribution vector; c, ℓ, and o denote renewable curtailment, load shedding, and over-generation spill, respectively. The coefficients (a_{gk}, b_{gk}) form the implemented convex piecewise-linear energy curve: three equal-width segments with marginal-cost multipliers $(0.9, 1.0, 1.1)$.

After z is revealed, the committed fleet uses nonnegative within-band upward redispatch q^f, penalized upward redispatch q^p, downward redispatch q^d, and recourse slacks $(\ell^1, o^1, c^1, \phi^{1,+}, \phi^{1,-})$.

For fixed $x \in \mathcal{X}(D, \widehat{r})$, its recourse value is

$$Q(x,z) = \min \quad \sum_{t,g}\big(c_g q^f_{tg} + \kappa c_g q^p_{tg}\big) + P_{\mathrm{bal}} \sum_t \mathbf{1}^\top(\ell^1_t + o^1_t) + P_{\mathrm{curt}} \sum_t \mathbf{1}^\top c^1_t + P_{\mathrm{tx}} \sum_t \mathbf{1}^\top(\phi^{1,+}_t + \phi^{1,-}_t) \tag{48}$$

$$\text{s.t.} \quad q^f_{tg} \le \beta_g u_{tg}, \tag{49}$$

$$q^f_{tg} + q^p_{tg} \le (\overline{P}_g - \underline{P}_g) u_{tg} - \delta_{tg}, \qquad q^f_{tg} + q^p_{tg} \le 60\rho_g u_{tg}, \tag{50}$$

$$q^d_{tg} \le \delta_{tg}, \qquad q^d_{tg} \le 60\rho_g u_{tg}, \tag{51}$$

$$\mathbf{1}^\top(q^f_t + q^p_t - q^d_t) - \mathbf{1}^\top c^1_t - \mathbf{1}^\top o^1_t + \mathbf{1}^\top \ell^1_t = -\mathbf{1}^\top z_t, \tag{52}$$

$$f^1_t = H_G(q^f_t + q^p_t - q^d_t) + H\ell^1_t - H_R c^1_t - H o^1_t + H_R z_t, \tag{53}$$

$$\underline{F} - \phi^{0,-}_t - \phi^{1,-}_t \le f^0_t + f^1_t \le \overline{F} + \phi^{0,+}_t + \phi^{1,+}_t. \tag{54}$$

All variables in this recourse problem are continuous. With startup cost S_g and minimum-load cost M_g, define

$$C^0_{\mathrm{slack}} = P_{\mathrm{bal}} \sum_t \mathbf{1}^\top(\ell^0_t + o^0_t) + P_{\mathrm{curt}} \sum_t \mathbf{1}^\top c^0_t + P_{\mathrm{tx}} \sum_t \mathbf{1}^\top(\phi^{0,+}_t + \phi^{0,-}_t). \tag{55}$$

The implementation uses $P_{\mathrm{bal}} = 1000$, $P_{\mathrm{curt}} = 500$, $P_{\mathrm{tx}} = 2500$, $\kappa = 1.2$, and $\beta_g = \max\{0.05, 0.03\overline{P}_g\}$ in the stated per-unit cost convention. The reported robust model is exactly

$$\min_{x \in \mathcal{X}(D,\widehat{r})} \left\{ \sum_{t,g}(S_g v_{tg} + M_g u_{tg} + e_{tg}) + C^0_{\mathrm{slack}} + \max_{z \in \mathcal{U}(s)} Q(x,z) \right\}. \tag{56}$$

Its binary commitment variables make the first stage a mixed-integer linear program (MILP), while the recourse problem is linear.

Feature-Based Newsvendor. Following Example 5 of Yang et al. (2026), we use the Genpact Food Demand Forecasting benchmark data supplied with that paper. Context contains the released price and promotion features and uncertainty is subsequent demand. For order y and demand z, loss is $h(y-z)_+ + b(z-y)_+$ with $h = 0.25$ and $b = 1$. A 52-week rolling training window is followed by four held-out weeks in each of 23 folds, yielding 92 settled decisions. For each context s, every robust-set method solves

$$\min_{y \ge 0} \max_{z \in \mathcal{U}(s)} \left\{ h(y-z)_+ + b(z-y)_+ \right\}. \tag{57}$$

Equivalently, for a fixed scenario z, its common-interface recourse LP is

$$Q(y,z) = \min_{q \ge 0}\{q : q \ge h(y-z),\ q \ge b(z-y)\}. \tag{58}$$

Robust Portfolio Optimization. We follow Section 5.2 and the released repository of Chenreddy et al. (2022). The decision is a fully invested, long-only allocation over 15 selected assets and uncertainty is the vector of daily returns. Evaluation follows the released rolling construction for the three held-out years 2017–2019 and ten asset-selection seeds per year. The source benchmark formulates this experiment as risk minimization and evaluates the resulting policies by held-out loss VaR across confidence levels. We retain its VaR_{95} endpoint and add CVaR_{95} of the same saved daily losses. For return vector $z \in \mathbb{R}^{15}$ and allocation p, every robust-set method solves the single-stage problem

$$\min_{p \in \mathbb{R}^{15}} \max_{z \in \mathcal{U}(s)} -z^\top p \quad \text{s.t.} \quad \mathbf{1}^\top p = 1, \quad 0 \le p_i \le 1. \tag{59}$$

For a held-out realized return z^{out}, settlement loss is $-(z^{\mathrm{out}})^\top p$; the reported VaR and CVaR are computed over these saved daily losses.

B.2 Frozen preprocessing and method configurations

All fitted transformations use the training split of the corresponding fold. Fixed numerical settings and data-dependent selection rules are distinguished below. Feature-Based Newsvendor uses the

four released context variables (checkout price, base price, e-mail promotion, and homepage promotion). CUB-RO uses weighted squared Euclidean distance after training-window coordinatewise standardization, with fixed coordinate weights recorded in the accompanying experiment configuration. Within each 52-week training window, its first 36 observations fit candidate representations and the remaining 16 select $(K, \widehat{\beta})$ by validation CVaR_{80} subject to validation mean loss no more than 5% above SAA. The candidate grid is $K \in \{5, 8, 10, 12, 15, 18, 20, 24, 30, 36\}$ and $\widehat{\beta} \in \{0, 0.025, 0.05, 0.075, 0.1, 0.15, 0.2\}$. The selected pair is refitted on all 52 observations before settling the next four weeks. Across the 23 windows, $\widehat{\beta} = 0.2$, 0, and 0.05 are selected 12, 10, and one time, respectively. KNN-RO and cKNN-GB retain their frozen $K = 15$ configurations. CT-DRO uses the released weighted food distance on the four unstandardized coordinates. SCUC uses $K = 10$ and cosine dissimilarity between complete 24-hour decision-time context trajectories. The implementation floors the product of the two context norms at 10^{-12} in the cosine denominator. For each weather variable, its training mean and standard deviation are pooled across training days and all 24 hours before standardization. The same context tensor is supplied to cAGRO; CUB-RO uses it in the context dissimilarity measure and cAGRO uses it as the generator condition. The renewable forecast also enters the SCUC optimization instance. Robust Portfolio Optimization uses $K = 40$, the published 21-dimensional contemporaneous market covariates, and squared Euclidean distance after the same training-only standardization. A feature with training standard deviation below 10^{-12} is assigned unit scale. Stable sorting resolves dissimilarity ties by training order.

Feature-Based Newsvendor is scalar; its B_s is the demand standard deviation estimated from the complete training window and is independent of the query. For the training-selected $\widehat{\beta}$, each center radius is $\widehat{\beta}$ times its nearest-center distance in training-standardized demand units. For SCUC, CUB-RO uses $\widehat{\beta} = 0.05$ and $B_s = \mathrm{diag}(\max\{\widehat{\sigma}_j, 10^{-8}\})$, where $\widehat{\sigma}_j$ is estimated from the complete training split. This factor is also independent of the query. The radius of center z_i is $\widehat{\beta}$ times its nearest-center distance in these standardized scenario coordinates. This diagonal choice avoids estimating an unrestricted 144×144 covariance matrix from $K = 10$ matched days; joint temporal and spatial patterns remain in each complete trajectory center. The SCUC value $\widehat{\beta} = 0.05$ is a fixed dimensionless radius multiplier. The portfolio experiment instead uses query-dependent local joint covariance estimated from the K matched training returns, together with a training-only leave-one-out multiplier: for each training observation we remove that observation, repeat the K-neighbor construction with context standardization fixed, and record the smallest multiplier at which its scenario enters the resulting union. We set the radius multiplier $\widehat{\beta}$ to the higher empirical 0.95 quantile of these leave-one-out calibration scores. For each query, $\widehat{\Sigma}_s$ is the Ledoit–Wolf shrinkage estimate of those matched returns. We symmetrize it and compute an eigendecomposition $\widehat{\Sigma}_s = V\,\mathrm{diag}(\lambda)V^\top$, and use

$$B_s = V\,\mathrm{diag}\left(\sqrt{\max\{\lambda_j, 10^{-10}\max(\lambda_{\max}, 10^{-12})\}}\right).$$

The eigenvalue floor makes B_s full rank when $K < d$, and all Mahalanobis computations use this regularized factor.

The frozen baseline settings are as follows. RO uses the coordinatewise 2.5% and 97.5% training quantiles; its SCUC budget is $\Gamma = 144$. The SCUC SAA project protocol clusters all 120 training residual trajectories into ten K-means representatives and solves the empirical-average model on those representatives; every resulting decision is then settled at the held-out residual in the common SCUC model. SCUC DRO uses the fixed radius $r = 0.5$. In Feature-Based Newsvendor, the DRO radius is selected from $\{0, 0.5, 2, 5\}$ and the CT-DRO radius from $\{0, 0.1, 0.3, 1\}$ using only the final 20% of each rolling training window. The selected DRO radius is zero in all 23 rolling folds, so its decisions and held-out losses coincide with SAA in Table 1(a). A supplementary fixed-radius sweep is reported in Appendix E.2. In the portfolio benchmark, the DRO radius is one quarter of the median nearest-neighbor distance in standardized return space, converted back to return units; CT-DRO uses $\rho = 0.1$ and ten clusters. IDCC uses two clusters, target coverage 0.95, $(\alpha_k, \alpha_s, \beta) = (0.5, 0.5, 0.1)$, 800 epochs, learning rate 0.01, weight decay 0.01, and randomized assignment. cKNN-GB uses four K-means scenario regions (10 initializations) and the empirical 0.95 within-region covariance radius. On Portfolio, KNN-RO uses the same K, context features, and distance as CUB-RO, with zero ball radius.

The three contextual comparison methods have different SCUC formulations. The tractable CT-DRO decision models used in its source paper optimize affine decision rules or one-dimensional continuous decisions under convexity conditions (Yang et al., 2026); they do not provide a binary unit-commitment counterpart. IDCC represents its learned ReLU/ellipsoid region by a mixed-integer model using a ReLU encoding (Fischetti & Jo, 2018), but its published adversarial formulation assumes that the cost is linear in the uncertainty (Chenreddy et al., 2022). SCUC recourse is instead a convex piecewise-affine LP value function. We therefore use the fitted IDCC region in the fixed-decision piecewise-affine diagnostic of Appendix F.4, rather than report it as the published end-to-end SCUC method. SRO-SI has a third, computational, limitation described next.

SRO-SI follows the samplewise robust objective with k-nearest-neighbor weights from Bertsimas et al. (2023a). Recourse models use the paper's multi-policy linear-decision-rule approximation and its robust counterpart; the objective-linear Portfolio model reduces directly to weighted local support functions. In Feature-Based Newsvendor, $K \in \{14, 20, 27\}$ and the scalar-neighborhood radius multiplier lies in $\{0, 0.05, 0.1, 0.2, 0.4\}$, selected on the final 20% of each rolling training window; in Robust Portfolio Optimization, $K \in \{20, 40, 80\}$ and the radius multiplier lies in $\{0, 0.1, 0.25, 0.5, 1\}$, selected on the final 25% of the current training window. In both cases the test period remains withheld. The portfolio radius is expressed relative to the median training-sample nearest-neighbor ℓ_1 return distance; in the scalar Newsvendor problem the norm choices coincide.

For two-stage recourse, the published SRO-SI computation assigns every positive-weight neighborhood its own affine policy $y_i(z) = y_{i,0} + Y_i(z - z_i)$ and robustifies all feasibility and cost expressions over that neighborhood. It therefore introduces Kn_yd_z policy-slope coefficients before dualizing the robust constraints. CUB-RO instead makes context matching part of the uncertainty-set construction. Its adversarial call returns an explicit scenario, and the restricted master adds the corresponding native recourse block without introducing a decision rule for every matched neighborhood. Hence its model size grows with generated adversarial scenarios rather than with the dense product Kn_yd_z.

Applying the same published multi-policy formulation to the full SCUC model is not computationally viable. SCUC has 11,280 continuous recourse variables and 144 uncertainty coordinates; Table 3 lists the deterministic-counterpart size before presolve. We therefore report SRO-SI on the two continuous benchmarks and do not replace its SCUC formulation by a different scenario-generation algorithm.

Table 3: Size of the published multi-policy LDR counterpart when applied to the full 144-dimensional SCUC model.

K	policy intercepts	policy slopes	support epigraphs	support inequalities
5	56,400	8,121,600	129,155	37,196,640
10	112,800	16,243,200	258,310	74,393,280
20	225,600	32,486,400	516,620	148,786,560

As a formulation check on the scalar Newsvendor model, we used two local centers at 2 and 4 with equal weights and radius 1. Explicit maximization over the two interval endpoints and the multi-policy robust counterpart both return the order decision 4.6 and robust objective 0.65. Repeating the calculation with centers 0.5 and 4 and intersecting the radius-one neighborhoods with nonnegative demand gives decision 4.6 and objective 0.775 under both explicit endpoint evaluation and the support-dual formulation.

For the neural baselines, Feature-Based Newsvendor uses a two-dimensional latent space, 80 epochs, two restarts, and 12 PGA steps per restart. The portfolio runs use latent dimension four, 120 epochs, four restarts, and 30 steps. SCUC AGRO uses an unconditional VAE with latent dimension eight, hidden width 32, and 80 epochs; cAGRO uses a conditional VAE (Sohn et al., 2015) with latent dimension four, width 64, and 400 epochs, conditioned on the same standardized 24-hour context trajectory used by CUB-RO and without a support filter. Their latent-ball radius is the training 0.95 quantile of encoded latent norms. The AGRO encoder and decoder each have two width-32 ReLU–batch-normalization layers; the cAGRO encoder and decoder each have two width-64 ReLU layers. Their main-table SCUC rows use PGA in latent space. At each PGA iteration, the recourse LP is solved at the decoded scenario, and its cost gradient is backpropagated through the decoder (Brenner et al., 2025). In the solver comparison in Table 8, CRG-PGA is applied to the same generators and latent balls, using the budget in Section B.4. The cVAE scenario comparison in Figures 1 and 2 uses

the forecast-conditioned model and search settings specified in Appendix F.4. The SCUC PGA runs use three restarts and 24 steps.

Table 4: Adversarial and outer solvers for the main experiments.

Benchmark	CUB-RO AP	Outer optimization	AGRO/cAGRO AP
Feature-Based Newsvendor	exact scalar endpoints	scalar convex program	PGA in latent space
Robust Portfolio Optimization	exact ellipsoid support	convex conic program	PGA in latent space
Security-Constrained Unit Commitment	CRG-PGA on local balls	MILP C&CG with LP recourse	PGA in latent space

B.3 Evaluation protocol

All methods use the same benchmark formulation and common mathematical interface, and maintable methods settle their saved first-stage decisions on held-out truth. The IDCC–SCUC adapter supplies the separately reported mechanism experiment using affine pieces from the same SCUC recourse LP.

Tail-risk aggregation. For a finite loss sample $\ell_1, \ldots, \ell_m$, we use the linearly interpolated empirical 0.95 quantile for VaR and the standard empirical CVaR

$$\widehat{\mathrm{CVaR}}_{.95} = \min_{\tau \in \mathbb{R}} \left\{ \tau + \frac{1}{0.05m} \sum_{i=1}^{m} (\ell_i - \tau)_+ \right\}.$$

Feature-Based Newsvendor applies this definition to the 92 held-out losses from 23 nonoverlapping four-week test blocks; SCUC applies it to the 30 daily realized total costs. Their primary evaluation statistic is the arithmetic mean of the same held-out settled costs.

Portfolio aggregation is hierarchical. For each held-out-year–asset-seed run, VaR and CVaR are first computed from that run's approximately 251 daily losses, including the fractional boundary order statistic dictated by that run's sample size. Table 1 then averages each risk statistic over the 30 evaluations: three held-out years and ten asset-selection seeds per year. Asset-selection runs within a year share the same market calendar. The reported values therefore summarize the benchmark's hierarchical year–asset-seed construction. Boldface marks the smallest reported point estimate in each column.

B.4 Optimization budgets, software, and randomness

CUB-RO has an exact scalar separation problem on Feature-Based Newsvendor, and the portfolio CUB-RO and cKNN-GB local-ball problems in the main experiments are solved directly as convex conic models. For SCUC, CUB-RO and cKNN-GB in the main experiments and AGRO and cAGRO in the CRG-PGA solver comparison use four feasible initializations, one additional independently initialized search trajectory, and step budget $J = 30$ as defined in Algorithm 1. All four methods use five seeds, a 12-iteration outer cap, and a returned-scene residual threshold of 10^{-4}. Here the outer termination residual is the absolute difference between the restricted-master value and the objective obtained from the returned LP-validated scene. Table 9 reports, for CRG-PGA, the numbers of runs reaching the inner search cap and satisfying the outer stopping criterion, each out of 150 runs per method. Exact scalar and conic separators, candidate-generation work, stopping status, and wall-clock time are recorded in their corresponding units.

Optimization uses Gurobi 13.0.2, with Gurobi 13.0.1 for the five-seed SCUC CUB-RO evaluation, on a central processing unit (CPU). Each SCUC job uses eight Gurobi threads on AMD EPYC 7763 processors, and independent days are batch-parallelized; neural models use PyTorch 2.11 without graphics-processing-unit (GPU) acceleration. Randomized methods use five model seeds in the main performance tables. The portfolio experiment separately uses the ten released asset-selection seeds in each held-out year; these define asset universes and are not model seeds. IDCC's rolling Feature-Based Newsvendor split is indexed by the training-window start. Outer C&CG uses a returned-scene residual threshold of 10^{-4}; IDCC mixed-integer piece queries use relative mixed-integer-programming (MIP) tolerance 10^{-4}.

The main tables use the configurations listed above. Code, preprocessing scripts, optimization model implementations, experiment configurations, and random seeds will be released with the paper.

C SUPPLEMENTARY PERFORMANCE ANALYSIS

C.1 FEATURE-BASED NEWSVENDOR TIME-BLOCK ANALYSIS

The 92 held-out decisions form 23 consecutive four-week test folds. We use 20,000 paired circular moving-block bootstrap draws with a primary block length of four folds (16 weeks); every method and model seed receives the same sampled fold indices. For randomized methods, each draw computes mean cost and CVaR separately for each model seed, then averages each statistic over the five seeds. Table 5 reports percentile intervals for the primary mean-cost and supplementary CVaR summaries.

Table 5: Feature-Based Newsvendor paired time-block analysis. Brackets give the 2.5% and 97.5% bootstrap percentiles.

Method	mean cost [interval]	CVaR_{95} [interval]
SAA	87.04 [63.87, 114.83]	472.66 [203.99, 696.78]
RO	133.61 [95.81, 187.03]	667.38 [325.55, 934.57]
DRO	87.04 [63.87, 114.83]	472.66 [203.99, 696.78]
KNN-RO	90.15 [73.55, 116.92]	419.38 [206.00, 630.03]
SRO-SI	87.66 [65.66, 117.96]	426.05 [177.20, 695.78]
cKNN-GB	109.29 [88.68, 138.39]	455.13 [283.34, 559.49]
CT-DRO	87.13 [64.57, 114.62]	482.21 [224.78, 696.54]
IDCC	104.11 [75.29, 135.62]	405.81 [199.63, 658.54]
AGRO	93.54 [69.63, 124.47]	482.18 [220.41, 710.50]
cAGRO	92.26 [68.93, 124.07]	452.30 [249.68, 635.66]
CUB-RO	**85.84** [65.75, 117.10]	430.67 [224.91, 660.57]

For the primary mean-cost statistic, CUB-RO minus SAA is -1.20 with interval $[-16.97, 14.60]$, and CUB-RO minus the frozen KNN-RO baseline is -4.31 with interval $[-12.38, 3.02]$. For supplementary CVaR, the corresponding differences are -41.99 with interval $[-134.04, 48.03]$ and 11.28 with interval $[3.10, 30.55]$, respectively. IDCC minus CUB-RO has mean difference 18.28 and CVaR difference -24.86: its lower tail estimate is accompanied by a markedly larger average realized loss. The CVaR difference SRO-SI minus SAA is -46.61 with the 16-week-block interval $[-129.75, 1.30]$; using four- and eight-week blocks changes the upper endpoint to 23.87 and 12.47, respectively. Thus the upper confidence limit is sensitive to the temporal block length.

C.2 YEAR-STRATIFIED PORTFOLIO RISK AND TIME-BLOCK ANALYSIS

The 30 year–asset-seed evaluations are not 30 independent repetitions: the ten asset seeds in a held-out year share every market date. Table 6 therefore separates the three years and averages only over the ten asset universes within each year. The overall column is the descriptive mean of the 30 within-evaluation risk statistics used in Table 1. The 2017, 2018, and 2019 panels contain 251, 251, and 252 trading days, respectively. Equality of the saved date sequence across all methods and seeds within each year is verified before aggregation.

For paired uncertainty quantification, each bootstrap draw independently samples circular blocks within each year and applies the same sampled date indices to every method and all ten asset seeds in that year. With 10,000 draws and a 20-trading-day block, the mean CVaR_{95} difference in percentage points is -0.0745 for **CUB-RO** minus **cKNN-GB**, with interval $[-0.1332, -0.0183]$, and -0.2249 for **CUB-RO** minus **KNN-RO**, with interval $[-0.2774, -0.1707]$. For **CUB-RO** minus **SRO-SI**, the difference is -2.0553 percentage points with interval $[-2.2722, -1.8094]$. The cKNN-GB statistic averages its five model-seed values inside each paired resample. The resampling keeps the three observed years as fixed strata and measures stability to temporal variation within those years.

C.3 SCUC PAIRED TIME-BLOCK ANALYSIS

The 30 held-out SCUC days are consecutive. We use 20,000 paired circular moving-block bootstrap draws with five-day blocks. Each draw applies the same day indices to every method and seed;

Table 6: Year-stratified held-out portfolio risk (percent loss). Each randomized-method cell additionally averages five model seeds after computing risk within each of the ten asset-seed evaluations; "overall" averages the three year-level strata.

(a) VaR_{95}.

Method	2017	2018	2019	overall
SAA	2.2079	3.0822	2.2910	2.5270
RO	1.0391	1.9641	1.4433	1.4822
DRO	0.6105	1.8452	1.0733	1.1763
KNN-RO	0.7568	1.6163	1.0685	1.1472
SRO-SI	2.0374	2.6458	2.2475	2.3102
cKNN-GB	0.5922	1.6141	1.0123	1.0729
CT-DRO	1.1801	2.2497	1.5599	1.6632
IDCC	1.0341	2.1376	1.5048	1.5588
AGRO	1.1531	2.1606	1.5941	1.6359
cAGRO	1.4736	2.5095	1.6800	1.8877
CUB-RO	**0.5859**	**1.5449**	**0.9450**	**1.0253**

(b) CVaR_{95}.

Method	2017	2018	2019	overall
SAA	3.0029	4.5631	3.4241	3.6634
RO	1.6612	2.9491	2.2890	2.2997
DRO	0.9382	2.5694	1.7581	1.7552
KNN-RO	1.1811	2.3880	1.6568	1.7420
SRO-SI	3.4101	3.9079	3.3990	3.5723
cKNN-GB	**0.8407**	2.3718	1.5621	1.5915
CT-DRO	1.6815	3.0931	2.2385	2.3377
IDCC	1.5696	3.4221	2.2858	2.4258
AGRO	1.6385	3.2219	2.3353	2.3986
cAGRO	2.2111	3.6565	2.5340	2.8005
CUB-RO	0.8893	**2.2482**	**1.4137**	**1.5170**

for CUB-RO, AGRO, cAGRO, and cKNN-GB, the mean cost is first computed within each of five complete seeds and then averaged. The KNN-RO row uses its 30-day cost vector. This analysis holds seeds fixed. Table 7 uses the same hundreds-of-USD units as Table 2. SRO-SI is not included because its published multi-policy counterpart has the SCUC-scale dimensions reported in Table 3.

Table 7: SCUC paired five-day-block analysis of mean settled cost. Brackets are 2.5% and 97.5% bootstrap percentiles. The final column is the paired difference CUB-RO minus the indicated method.

Method	mean cost [interval]	paired difference [interval]
SAA	1945.20 [1846.06, 2038.39]	-9.46 [-30.35, 8.42]
RO	2040.05 [1920.55, 2165.57]	-104.31 [-182.61, -50.47]
DRO	1956.24 [1877.19, 2034.54]	-20.50 [-39.19, -3.77]
KNN-RO	1947.04 [1847.14, 2040.64]	-11.30 [-37.37, 8.21]
cKNN-GB	1990.83 [1903.93, 2079.99]	-55.09 [-73.16, -36.21]
AGRO	1965.07 [1882.82, 2043.80]	-29.33 [-43.64, -15.53]
cAGRO	1971.43 [1887.81, 2056.79]	-35.69 [-57.85, -14.83]
CUB-RO	**1935.74** [1844.30, 2022.22]	—

The paired intervals against SAA and KNN-RO cross zero. The intervals against DRO, RO, cKNN-GB, AGRO, and cAGRO remain below zero under this five-day block choice. Repeating the analysis with one-, three-, and seven-day blocks gives the same pattern of intervals crossing or remaining below zero. These differences are interpreted relative to the perfect-information mean of 1896.35. The available SAA-to-perfect-information improvement is 48.85 cost units, of which CUB-RO realizes 9.46, or 19.36%. The corresponding reductions in the gap to the perfect-information lower bound, relative to KNN-RO and DRO, are 22.29% and 34.23%. The paired intervals characterize variation across held-out day blocks, whereas these ratios measure the fraction of each baseline's gap to the perfect-information lower bound eliminated by CUB-RO on this evaluation period.

C.4 Seed stability

We independently rerun the machine-learning (ML) uncertainty-set pipelines with five seeds while holding the training/test split, optimization model, hyperparameters, and held-out settlement fixed. Seeds control randomness in uncertainty-set construction and numerical search. Each seed first produces one complete benchmark statistic: mean cost over the 92 Feature-Based Newsvendor decisions, mean VaR over the 30 Robust Portfolio year–asset-seed evaluations, and mean settled cost over the 30 SCUC days. Figure 4 shows percentage deviations from each method's five-seed mean. The large marker is zero, and the bar shows the sample standard deviation of these percentage deviations; translucent points correspond to the five complete runs.

On Feature-Based Newsvendor, the sample standard deviation of mean cost is 0.20 for cKNN-GB, 7.25 for IDCC, 4.37 for AGRO, 2.82 for cAGRO, and zero for deterministic CUB-RO; these correspond to coefficients of variation of 0.18%, 6.96%, 4.68%, 3.05%, and zero. The corresponding CVaR standard deviations are 1.24, 7.87, 43.06, 24.98, and zero, respectively. On Robust Portfo-

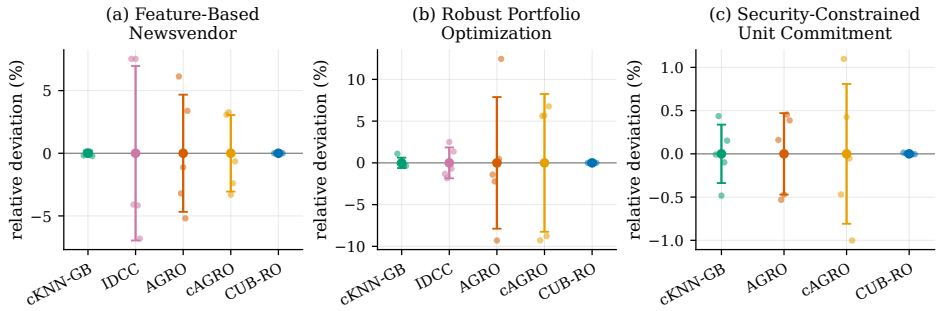

Figure 4: Seed stability under fixed data and optimization settings. Panel (a) uses mean held-out Feature-Based Newsvendor cost; panel (b) uses mean Robust Portfolio VaR; and panel (c) uses mean held-out SCUC settled cost. Values are percentage deviations from each method's five-seed mean; error bars show their sample standard deviations.

lio, the coefficients of variation of mean VaR are 0.62% for cKNN-GB, 1.85% for IDCC, 7.88% for AGRO, 8.25% for cAGRO, and zero for CUB-RO. On SCUC, the standard deviations of mean settled cost are 6.73 for cKNN-GB, 9.25 for AGRO, 15.94 for cAGRO, and 0.169 for CUB-RO, corresponding to 0.34%, 0.47%, 0.81%, and 0.00872% of their means. CUB-RO uses the main-table runs, with fixed retrieval and randomized separation. Across all three benchmarks, CUB-RO has less reported seed variation than cKNN-GB, AGRO, and cAGRO; cAGRO has the largest SCUC variation, whereas IDCC has the largest Feature-Based Newsvendor variation.

D SCUC COMPUTATION AND SEPARATION ACCURACY

D.1 COMPUTATIONAL COST AND STOPPING STATUS

Table 8 compares standard PGA with CRG-PGA on the same 30 SCUC days. CUB-RO uses the main-table and Figure 4 evaluations. Runtime includes the outer MILP C&CG loop. An AP call is one outer separation call; search LPs count original recourse LP solves during separation search, excluding terminal validation and separate diagnostic audits. The Stored pieces column reports the number of affine-piece records retained during search. Within each search, records are deduplicated by LP basis indices and recourse gradients rounded to ten decimal places. Counts are summed over component searches and AP calls.

Table 8: SCUC computational cost on 30 held-out days. All CRG-PGA rows use the shared search budget in Section B.4, including the 12-iteration outer cap. Means average the per-seed 30-day means; medians pool all runs in a row. Entries a/b are mean/median; AP calls report the mean.

Set	AP solver	Seeds	Time (s)	Outer it.	AP calls	Search LPs	Stored pieces
CUB-RO	CRG-PGA	5	771.5/647.4	7.2/6	7.2	1914.1/1621	1410.8/1176
cKNN-GB	CRG-PGA	5	1255.5/1194.3	10.1/11	10.1	3279.4/3279	2936.3/2912.5
AGRO	PGA	5	92.6/74.9	3.0/3	3.0	213.6/216	–
AGRO	CRG-PGA	5	821.0/736.3	12.0/12	12.0	1425.1/1494.5	1292.5/1360
cAGRO	PGA	5	85.0/74.0	3.0/3	3.0	213.6/216	–
cAGRO	CRG-PGA	5	768.2/727.2	11.8/12	11.8	1570.6/1647.5	1410.8/1488

Table 9: SCUC stopping status for CRG-PGA under the common 12-iteration outer cap. Inner cap hit counts runs reaching the inner search cap; Outer stop counts runs satisfying the outer stopping criterion. Both columns use 150 runs per method: five seeds on the same 30 held-out days.

Set	Outer cap	Inner cap hit	Outer stop
CUB-RO	12	0/150	134/150
cKNN-GB	12	150/150	80/150
AGRO	12	150/150	2/150
cAGRO	12	150/150	6/150

For any feasible scenes accumulated by a budget-limited separator,

$$\widehat{H}_t(x) := \max_{z \in \mathcal{S}_t} Q(x, z) \leq \sup_{z \in \mathcal{U}(s)} Q(x, z) = H(x). \tag{60}$$

Since $\mathcal{S}_t \subseteq \mathcal{S}_{t+1}$, the restricted-master optimal value is monotone nondecreasing for a nested scenario sequence. Both stopping counts use the same 150 runs per method under the common 12-iteration outer cap. Proposition 6 gives the exact-AP certificate; Theorem 7 relates δ_{sep} to decision quality at a terminal iterate satisfying its stopping condition.

D.2 Separation error

We quantify the separation error δ_{sep} in equation 33 on a separate diagnostic batch of 150 CUB-RO decisions (30 test days and five seeds). Holding each decision and its context fixed, we reevaluate CRG-PGA using the shared configuration in Section B.4. Each returned scenario $\hat{z}$ is checked for membership in the local-ball union and evaluated with the recourse LP.

For a fixed decision, the SCUC recourse LP decomposes into 24 hourly convex functions, each depending on six site-level errors. For each local ball, we apply an orthogonal rotation in normalized coordinates and enclose each hourly block in a box. We solve the hourly LP at its 64 vertices and construct an affine majorant of their values. Convexity extends this domination to the entire box. Summing the hourly majorants and maximizing over the ball gives an upper bound on its worst-case recourse cost. Box subdivision and support-function bounds on the resulting box–ball intersections tighten this calculation while preserving coverage of every local ball. Subtracting $Q(\hat{x}, \hat{z})$ from the resulting bound on $H_{\mathcal{U}}(\hat{x}, s)$ bounds δ_{sep}. The computation uses numerical feasibility checks for the LP solutions, affine majorants, and subdivision coverage.

Across all 150 decisions, the mean upper bound on δ_{sep} is **10.5541**, in the benchmark's cost units (hundreds of USD per day). Dividing each decision's separation-error bound by its evaluated objective $c^\top \hat{x} + Q(\hat{x}, \hat{z})$ and then averaging over all days and seeds gives **0.54%**. These fixed-decision evaluations quantify the separation-error term in Theorem 7 at the experimental search budget.

D.3 Inner search budget

Table 10 increases only the inner search budget on days 429 and 445, whose 4×30 settings are marked "no", with day 420 included as a comparison marked "yes". The outer cap remains 12 in every cell; Table 2 retains the frozen main configuration. A budget $a \times b$ denotes $M = a$ feasible-scene initializations and step budget $J = b$ in Algorithm 1.

Table 10: SCUC inner-budget stability on selected days. "Stop" is yes only if all five seeds meet the returned-scene outer criterion at 10^{-4}. Values are means for yes and worst-case results among non-stopping runs for no; H_{ret} is the returned LP-validated worst recourse value; Δu is commitment Hamming distance from the same day's 4×30 solution.

Day	Budget	Stop	Residual	H_{ret}	Settled	Δu	LPs
420	4×30	yes	1.72e−5	93.047	1998.057	0	1592
420	8×30	yes	1.72e−5	93.047	1998.057	0	2657
420	4×60	yes	5.45e−5	93.045	1998.057	48	2035
420	8×60	yes	6.93e−5	93.045	1998.057	48	3417
429	4×30	no	2.04e−4	71.429	1622.035	0	2520
429	8×30	no	2.04e−4	71.429	1622.035	0	4300
429	4×60	yes	7.64e−6	71.540	1621.931	48	1460
429	8×60	yes	7.64e−6	71.540	1621.931	48	2443
445	4×30	no	2.48e−3	99.262	1823.754	0	2990
445	8×30	yes	0	98.899	1825.154	48	4543
445	4×60	yes	4.55e−13	100.066	1823.876	20	3058
445	8×60	yes	2.27e−13	98.564	1824.496	68	5095

For day 420, reported settled cost varies by less than 0.0005. Increasing the step budget J changes "Stop" to "yes" for day 429 and changes the reported settled cost by -0.105 (-0.0065%). Each expanded-budget setting for day 445 is marked "yes"; the reported settled-cost change is between $+0.122$ and $+1.400$ ($+0.0067\%$ to $+0.0768\%$). The mean of 48 commitment changes on days

420 and 429 reflects full-day exchanges between generator pairs with identical first-stage objective coefficients.

The CRG-PGA rows in Table 8 use the shared configuration in Section B.4. The PGA rows use standard PGA for the same AGRO and cAGRO representations, with C&CG as the outer algorithm.

D.4 Scaling with the number of local components

Table 11 isolates the representation-size parameter K for CUB-RO at fixed $d = 144$, $\widehat{\beta} = 0.05$, four starts, and the same four preselected SCUC days used by Figure 5. Every component uses the same CRG-PGA backend. Fifteen of the 16 outer runs meet the stopping tolerance; the remaining $K = 5$ run has absolute residual 0.0072. Component AP calls scale as K times the number of outer AP calls. From $K = 10$ to $K = 40$, mean recourse-LP work grows from 1339 to 5885 solves per day. The $K = 5$ mean is nonmonotone because its difficult run requires substantially more outer iterations.

Table 11: CUB-RO SCUC scaling with the number of context-matched components. Entries a/b are mean/median over days 420, 430, 440, and 449.

K	Time (s)	Outer it.	Component APs	Recourse LPs	Pieces
5	2262.2/1451.3	14.8/12	73.8/60	2869/1592	2238/1120
10	563.0/589.4	5.3/5.5	52.5/55	1339/1342	966/970
20	1521.5/1376.5	7.5/7	150/140	3747/3446	2726/2519
40	2192.9/2311.7	6.8/7	270/280	5885/6138	4109/4345

The SCUC benchmark fixes the six-site, 24-hour uncertainty dimension and the physical unit-commitment system. The scaling experiment therefore varies the method's representation size K while preserving the benchmark optimization model.

E Sensitivity analysis

E.1 Representation and search parameters

Figure 5 separates representation hyperparameters $(K, \widehat{\beta})$ from separator budgets. Feature-Based Newsvendor uses all 23 rolling folds (92 settled decisions). The portfolio grid uses the training-side validation year for each of 30 year–asset-seed repeats, reserving the held-out returns for settlement.[1] For SCUC, we preselect the calendar-spaced held-out days 420, 430, 440, and 449 and run the complete Cartesian grid $K \in \{5, 10, 20, 40\}$ and $\widehat{\beta} \in \{0, 0.025, 0.05, 0.1\}$. These SCUC panels use the same CRG-PGA implementation as Table 11, separating representation parameters from the initialization count M and step budget J. The 30-day main-table configuration is fixed independently of this four-day diagnostic grid.

The response is problem dependent. On Feature-Based Newsvendor, increasing K from 5 to 40 raises CVaR from 396.7–402.4 to 411.4–414.7; at $K = 10$, changing the radius moves CVaR by 0.46. In Figure 5(b), mean portfolio validation CVaR decreases as K and $\widehat{\beta}$ increase over the evaluated grid.

On the four SCUC days, mean realized total cost ranges from 1983.1 to 2011.9 over all 16 configurations, and from 1988.0 to 2011.9 over the 13 configurations meeting the returned-scene criterion on all four days. Moderate increases can reach a plateau: for day 440, $(K, \widehat{\beta}) = (10, 0.025)$ and $(20, 0.025)$ return the same decision and settled cost even though the latter doubles the number of ball queries. At $(5, 0.05)$, $(5, 0.1)$, and $(40, 0.1)$, one of four outer problems reaches the 30-iteration limit; their maximum absolute residuals are 0.0072, 0.0023, and 9.647, respectively. Mean runtime at $(40, 0.1)$ is 5289.6 seconds versus 56.2 seconds at $(5, 0)$. These iteration-limited decisions remain in the reported means.

[1] The portfolio sensitivity grid uses C&CG with a 30-iteration cap and an absolute stopping tolerance of 10^{-5}. Validation metrics include decisions returned at the iteration cap without meeting the tolerance.

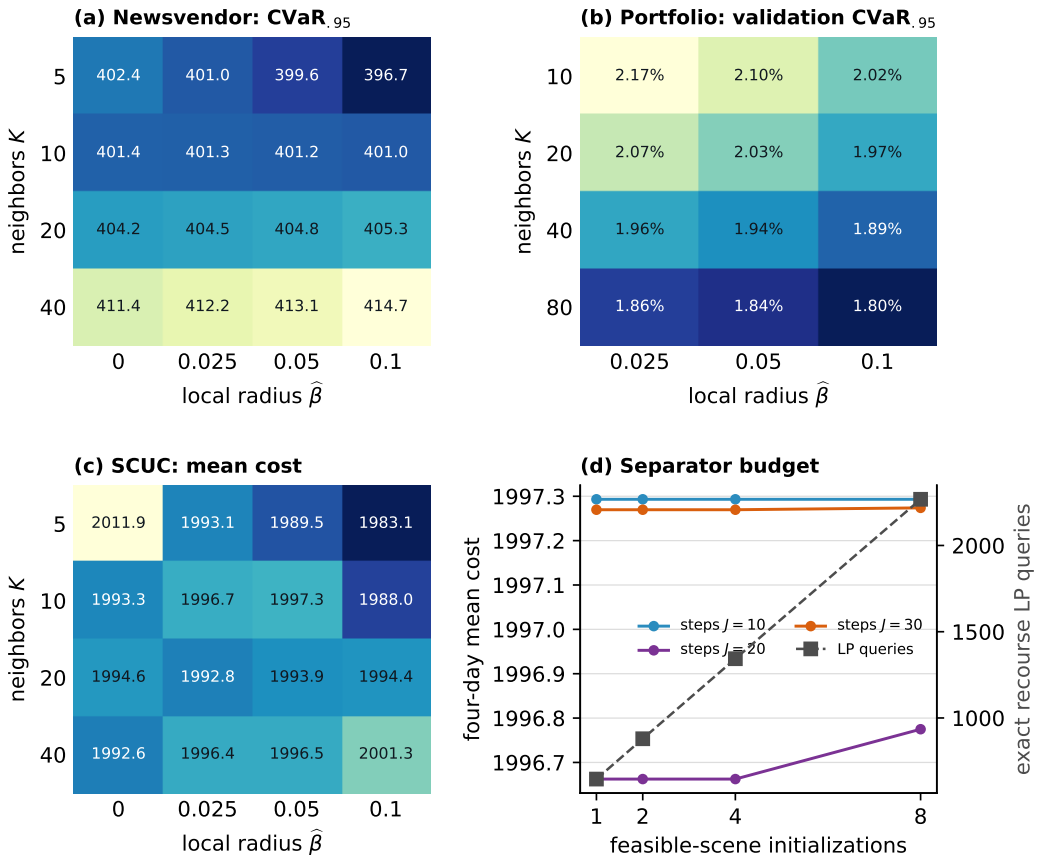

Figure 5: Hyperparameter sensitivity. Panels (a)–(c) report lower-is-better risk or cost; panel (b) uses training-only validation, and panel (c) uses four preselected SCUC days. Outer stopping status for panel (c) is reported in the accompanying text. Panel (d) varies M and J from Algorithm 1 at $K = 10$ and $\widehat{\beta} = 0.05$.

For the CRG-PGA search-budget sensitivity analysis, we fix $(K, \widehat{\beta}) = (10, 0.05)$ on the same four SCUC days and cross feasible-scene initializations $\{1, 2, 4, 8\}$ with step budgets $J \in \{10, 20, 30\}$. All 48 runs meet the outer termination residual criterion. Across the 12 settings, four-day mean settled cost lies in the narrow interval $[1996.663, 1997.293]$. At $J = 30$, raising initializations from one to four changes mean cost from 1997.270 to 1997.270 after rounding, while mean exact recourse-LP queries rise from 653 to 1339; eight initializations require 2498 queries. Thus the additional CRG-PGA search changes this four-day point estimate by much less than one cost unit. The frozen main configuration uses four initializations and $J = 30$.

E.2 Feature-Based Newsvendor DRO radius sensitivity

Table 12 reports an archived fixed-radius sweep with holding cost $h = 0.25$, shortage cost $b = 1$, 52-week training windows, and four-week test blocks. For each radius, the DRO problem is solved in each of the 23 rolling windows, and realized newsvendor losses are averaged over the 92 held-out decisions. The sweep evaluates $r \in \{0, 10, 50, 100, 300, 1000\}$.

At the reported precision, mean cost is unchanged for radii from 0 to 100 and increases to 235.50 and 241.60 at radii 300 and 1000, respectively. The smaller positive radii tested provide no improvement in mean cost, while the larger radii substantially increase it.

Table 12: Archived Feature-Based Newsvendor DRO radius sweep. Values are held-out mean costs, reported to the precision retained in the experiment record. The $r = 0$ row is the SAA reference for this sweep.

Radius r	Mean cost
0	87.04
10	87.04
50	87.04
100	87.04
300	235.50
1000	241.60

F Representation diagnostics

F.1 Convex-hull invariance sanity check

We test Proposition 1 on all 92 held-out Feature-Based Newsvendor decisions using the main-table nested-validation rule, its selected $(K, \widehat{\beta})$ in each rolling window, diagonal geometry, and weighted squared Euclidean distance between standardized contexts. Because demand is scalar, each local ball is an interval and its convex hull is the interval between the smallest lower endpoint and largest upper endpoint. We solve the interval union and this convex hull separately using the same Newsvendor formulation and exact endpoint-separation oracle. At the CUB-RO decision y we also compare the local-interval oracle with independent endpoint enumeration and report

$$g(y,s) = H_{\mathrm{exact}}(y,s) - H_{\mathrm{local}}(y,s), \qquad g_{\mathrm{rel}}(y,s) = \frac{|g(y,s)|}{\max\{1, |H_{\mathrm{exact}}(y,s)|\}}.$$

Here H_{exact} enumerates every local-interval endpoint, while H_{local} is the value returned by the local-interval oracle, which evaluates both endpoints of each interval using the recourse LP.

Table 13: Feature-Based Newsvendor convex-hull and fixed-decision AP sanity check over 92 held-out decisions.

Quantity	Result
Disconnected interval unions	92/92
Mean connected components	14.51
Mean interval length added by convexification	610.30
Union vs. hull: decision difference	0
Union vs. hull: AP-value difference	0
Union vs. hull: robust-objective difference	0
CUB-RO vs. exact union: objective difference	0
Local-interval oracle vs. exact AP: maximum absolute gap	0
Local-interval oracle vs. exact AP: mean / median gap	0/0
Local-interval oracle vs. exact AP: 95% / 99% gap quantile	0/0
Local-interval oracle vs. exact AP: maximum relative gap	0
Cases with $H_{\mathrm{local}} - H_{\mathrm{exact}} > 10^{-7}$ / 10^{-5}	0/0
Recourse LP calls: mean / median / maximum	31.22/20/72

Thus the equality is observed on sets for which convexification adds substantial interval length, rather than only on already connected unions. The 184 union and hull solves all meet the 10^{-8} C&CG tolerance. The CUB-RO implementation and the independent exact endpoint solve agree in both decision and robust objective. At the returned decision, the local-interval oracle also agrees with independent endpoint enumeration on all 92 queries. The experiment checks the identity in equation 35: for convex recourse, the relevant representation is the context-matched support boundary, not the omitted interior intervals.

F.2 FINITE-TARGET SCUC BOUND-TIGHTNESS PROTOCOL

The bound-tightness study uses the same SCUC formulation as Table 2 and retains the original 144-dimensional renewable-residual trajectories. Days 300–419 form the training pool and days 420–449 supply query contexts. Results are averaged over five random seeds for each test day and configuration before computing statistics across days. For each query, the finite reference set contains the 40 complete training trajectories with greatest cosine similarity in standardized hourly weather space. Target construction and representation selection use training residuals exclusively. Weather channels are standardized using means and standard deviations pooled over the training days and hours. The standardized hourly trajectories are flattened for cosine-similarity ranking against each test-day context.

The audit decision class contains the distinct feasible first-stage decisions already saved by weather-cosine, weather root-mean-square-error (RMSE), random, AGRO, and cVAE runs; day 420 additionally includes the available project-SAA, region-set, and IDCC outer decisions. For every day we solve the original recourse LP for every audit decision and all 120 training trajectories. This gives 18,240 recourse solves. On this finite class, $d_Q^{\mathrm{audit}}(z,u) = \max_{x\in\mathcal{X}_{\mathrm{audit}}} |Q(x,z) - Q(x,u)|$ and both directed distortions are computed by exhaustive pairwise comparison. Thus d_Q^{audit} is exact on the stated finite decision library. The population theorem uses the corresponding supremum over the complete feasible set $\mathcal{X}$.

Matched representations retain the first K target trajectories. Random representations retain a nested random ordering of the 120-day pool, while the shuffled-pair control preserves context ranks but randomly permutes their paired residual trajectories. These orderings remain fixed as K increases. We use $K \in \{2, 4, 6, 8, 10, 15, 20, 30, 40\}$. Both the representation-optimal and reference-optimal decisions are selected from the same finite decision library, with objectives comprising first-stage cost and worst-case recourse over the corresponding scenario set. The representation-optimal decision is then evaluated over the finite reference set. Figure 3(a) reports the median sum of missing and spurious distortions across the 30 test days. The study records the target decision-value gap and $\varepsilon_{\mathrm{miss}} + \varepsilon_{\mathrm{spur}}$. Among positive bounds, the gap-to-bound ratio has median 0.048, mean 0.207, and maximum 0.975. At $K = 40$, the matched representation equals the 40-trajectory reference set, so its directed distortion and target decision-value gap are zero by definition.

For the separation audit, the uncertainty set is fixed to the 40-trajectory reference set and the restricted AP inspects nested subsets of size $2, 4, 8, 16, 24, 32, 40$. Its exact error is the difference between the full and restricted worst recourse at the returned decision. The study compares the restricted-AP shortfall with the target decision-value gap to assess bound tightness. Their mean values decrease from $663.00/336.14$ at two inspected trajectories to $45.82/16.24$ at 32 and $0/0$ at 40. At 40, the restricted AP has inspected every trajectory in the reference set.

This experiment is a finite-reference-set bound-tightness study. It quantifies representation distortion and restricted-AP shortfall in the original SCUC recourse model and compares matched, random, and shuffled-pair constructions. The local-mass, support-regularity, uniform-Lipschitz, and growing-sample conditions characterize the population results in Section 5.2; the audit characterizes the fixed $n = 120$ empirical problem on $\mathcal{X}_{\mathrm{audit}}$.

F.3 CUB-RO LOCAL-BALL FINITE-REFERENCE AUDIT

Figures 3(b) and 3(c) use the same 30 query days, 40-trajectory finite targets, and saved decision libraries as Figure 3(a). For each day, the first ten matched trajectories are used both as the zero-radius representation and as the centers of the main SCUC CUB-RO construction with $\widehat{\beta} = 0.05$ and diagonal training scale. For every saved decision, the affine-ball oracle uses four initial starts and step budget $J = 30$ on each component, as defined in Algorithm 1. Every returned scenario is re-evaluated by the original recourse LP. The finite scenario set consists of the ten matched centers and the distinct LP-validated scenarios returned by local-ball queries, with duplicates removed. Every scenario in this finite subset of the continuous local-ball union is evaluated under every decision in the library to compute the directed distortions and the finite-reference decision gap. Figure 3(b) reports means across the 30 days, while Figure 3(c) plots one decision-gap–bound pair per day for each representation.

Across 30 days, the zero-radius mean/median/maximum finite-reference decision gap is $332.70/272.94/1072.51$; the corresponding values for this scenario set are $215.79/40.45/1110.86$. Mean missing distortion decreases from 841.20 to 771.02, while mean spurious distortion increases from zero to 80.12. Consequently the mean total bound changes from 841.20 to 851.14, even as the mean decision gap decreases by 35.14%. Optimization over this scenario set selects the same audit decision as the zero-radius representation on 15 days, lowers the gap on 13, and raises it on two. Every audited decision gap is bounded by its corresponding sum of finite-library distortions.

F.4 Fixed-decision SCUC comparisons and IDCC diagnostics

Fixed-decision scenario comparison. Figures 1 and 2 use test day 420, with days 300–419 as the training period. Each scenario contains renewable forecast residuals for six sites over 24 hours. All three methods share a first-stage decision obtained by sample average approximation (SAA). The SAA problem uses ten complete historical residual trajectories selected by weather similarity to the test day. Each weather channel is standardized using its mean and standard deviation pooled over the training days and hours, and the resulting hourly trajectories are ranked by cosine similarity. The relative MIP gap is set to 0.003, with a time limit of 1800 seconds. The first-stage decision, load, and renewable forecasts remain fixed when querying the three uncertainty sets. Figure 2 displays the returned and observed residuals using a shared color scale, with 1 p.u. equal to 100 MW.

IDCC configuration. IDCC is fitted to days 300–419 for 800 epochs with two regions, target weighted coverage 0.95, and its released Linear–ReLU–Linear set map. Directly supplying the full 504-dimensional hourly context causes one soft-cluster mass to underflow to zero. The diagnostic therefore averages each of six forecast and fifteen weather channels over the day, while retaining the complete 144-dimensional uncertainty trajectory. The resulting region masses are 0.433 and 0.567, with weighted empirical coverages 0.962 and 0.956. The context classifier selects the region used for scenario search in Figures 1 and 2.

For a fixed first-stage decision and selected region, every contained training trajectory is evaluated by the original SCUC recourse LP. Each discovered dual-affine piece is then globally maximized over the IDCC region. ReLU states use indicator constraints and the projected ellipsoid is imposed as a convex quadratic constraint. Every cut uses the returned scenario's value from the original recourse LP. Across the 30-day diagnostic all generated piece subproblems reached optimal solver status with relative MIP tolerance 10^{-4}.

CUB-RO configuration. CUB-RO uses ten weather-matched historical trajectories as centers. Each hour–weather coordinate is standardized across the training days before computing cosine similarity. The local sets share a diagonal scaling matrix formed from the training residual standard deviations, and each radius is 0.05 times the distance to the nearest other selected center in the scaled scenario space. CRG-PGA queries each component with four initial starts and an iteration budget of 30. The returned scenario is the candidate with the highest recourse cost across components, evaluated using the original recourse linear program.

cVAE configuration. The cVAE is conditioned on the 144-dimensional renewable forecast trajectory and uses latent dimension four, hidden width 64, and 400 training epochs. The latent-ball radius is the training 0.95 quantile of encoded latent norms, equal to 6.2301. Because the cVAE adversarial problem is more difficult to solve, we allocate a larger search budget to it. CRG-PGA uses four initial starts, a regional iteration budget of 60, and up to four pricing rounds, each with four starts and 60 iterations. The affine-piece limit is 512. Three independently initialized searches are performed, and the candidate with the highest recourse cost, evaluated using the original recourse linear program, is displayed.

Additional IDCC diagnostics. The input-box audit is performed in IDCC's scaled coordinates. A coordinate is declared near a box bound when its distance to the nearer bound is at most 1% of that coordinate's box span. Joint distance to training data is reported as standardized RMSE and as a Ledoit–Wolf Mahalanobis distance. These diagnostics localize the effect to joint recombination: the marginal extrapolation rate is 0%.

The one-piece query starts from the recourse piece active at the worst contained training trajectory and maximizes that affine function over the unchanged fitted IDCC region. It already attains a mean of 4619.91, while continuing recourse-piece generation increases the mean by only 60.50 (1.29%

of the iterated-query mean). Hence the mechanism is not explained by an unusually long search: a single exact argmax query over the unified region exposes almost all of the reported amplification. This IDCC adapter is reported as a mechanism experiment, separately from the SCUC baseline ranking.

For day 420, scanning all complete historical trajectories in the selected region gives recourse 434.50. Maximizing the first generated affine recourse piece raises the validated value to 3264.76; 2, 4, 8, and 12 generated pieces give 3292.77, 3295.47, 3298.32, and 3298.32. The diagnostic outer loop meets the prescribed stopping criterion in two iterations. Relative to a saved SAA control near 2000.6 total cost, it raises first-stage cost to 2010.02 and settles at 2036.15: substantial upfront protection buys little realized recourse reduction.

Table 14: IDCC–SCUC mechanism audit over 30 fixed held-out decisions. Results are averaged over five random seeds for each day before computing statistics across days. Metrics are grouped by the mechanism they diagnose; the experiment is reported separately from baseline rankings.

Diagnostic	Reported statistic	Value
Adversarial recourse value		
Reference query	Historical-trajectory maximum, mean over days	1105.60
One-piece query	First affine support query, mean over days	4619.91
Iterated query	Returned recourse value, mean over days	4680.41
Exposure	One-piece/iterated mean-value ratio	98.71%
Amplification	Iterated/reference ratio, median over days	4.21×
Amplification	Iterated/reference ratio, range over days	2.68–7.59×
Separation work		
One-piece query	Affine-piece maximization time, mean	0.238 s
Iterated query	Total solve time, mean/median	4.389/4.185 s
Iterated query	Generated pieces, mean/median	10.53/10.50
Iterated query	Original recourse-LP evaluations, mean/median	19.53/19.50
Proximity to box boundary		
Day-level count	Days with at least 95% of coordinates near a box bound	30/30
Coordinate fraction	Coordinates near a box bound, mean over days	99.63%
Marginal extrapolation	Coordinates outside training min–max, mean over days	0.00%
Distance and ellipsoid position		
Nearest trajectory	Standardized RMSE to nearest full trajectory, mean	2.03
Ellipsoid position	Fraction of projected-ellipsoid radius used, mean	64.48%