REFLECTED FENCHEL–YOUNG CONFORMAL CALIBRATION FOR DECISION-AWARE CONTEXTUAL ROBUST OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Conformal prediction is commonly used to construct uncertainty sets for contextual robust optimization. However, residual-based calibration can enlarge these sets in response to prediction errors that do not affect the optimal decision, inducing unnecessary conservatism. This motivates calibrating the regret of the nominal decision directly. However, we show that, under regularity conditions, minimizing worst-case regret over the resulting regret set uniquely recovers the nominal decision, so robust optimization leaves the decision unchanged. We address this limitation using reflected Fenchel–Young (FY) conformal scores for convex optimization with linear objective uncertainty. The score provides a richer characterization of decision performance, admitting an exact decomposition into nominal regret, the suboptimality of the hindsight optimizer under the prediction, and a reflected stability gap. Split conformal calibration gives finite-sample marginal guarantees. We further characterize when the calibrated set improves worst-case regret and develop a master–pricing algorithm with optimality certificates. Experiments confirm the predicted self-centering phenomenon and show that reflected FY refinement improves decision quality over residual-based robust baselines in portfolio and real-data solar-grid allocation problems.

1 INTRODUCTION

Predictive models are increasingly embedded in modern decision-making systems. Under the common predict-then-optimize (PTO) paradigm, a predictive model maps contextual features to estimates of unknown problem parameters, which are then used in a deterministic optimization problem. This separation reflects the distinct forecasting and operational decision functions found in practice. However, point predictions alone do not quantify forecast uncertainty, and decisions optimized for those predictions can cause substantial regret when the parameters are realized.

Robust optimization addresses this issue by optimizing against the worst-case realization in a prescribed uncertainty set (Ben-Tal et al., 2009). Recent work on contextual robust optimization uses contextual information and predictive models to construct these sets. Many of these approaches use conformal prediction to obtain finite-sample marginal coverage under exchangeability, using a quantile of conformal scores evaluated on held-out calibration data (Angelopoulos & Bates, 2023; Shafer & Vovk, 2008). The score determines which parameter realizations enter the uncertainty set considered by the downstream optimization problem. Existing methods often calibrate prediction residuals or other prediction-centered discrepancies (Johnstone & Cox, 2021; Sun et al., 2023; Patel et al., 2024). Scores based only on deviation from the predicted parameter $\widehat{\theta}$ do not distinguish errors by their effect on the decision. A large residual may leave the optimal decision unchanged, whereas a smaller perturbation can change the decision and incur substantial regret (Elmachtoub & Grigas, 2022; Smith & Winkler, 2006). Figure 1 illustrates this mismatch for a linear optimization problem.

With the predictor held fixed, we ask how the calibration score can support both regret certification and refinement of the nominal decision through minimax-regret optimization. A direct alternative to residual-based calibration is to calibrate the realized regret of the nominal decision $\widehat{x}$, which measures its loss relative to the optimal decision under the realized parameters. Our first result shows that, under the stated regularity conditions, the full regret sublevel set is *self-centering*: $\widehat{x}$

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.

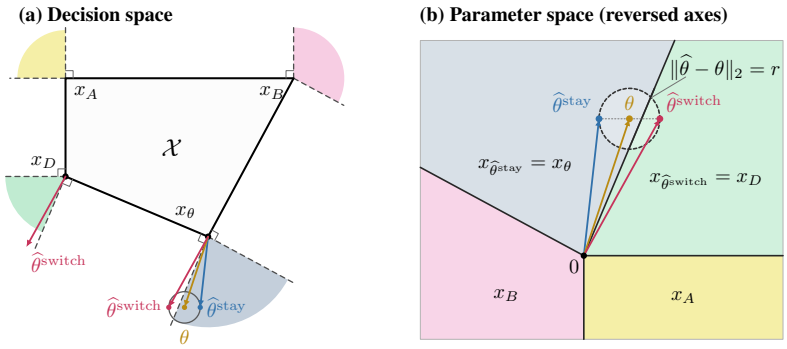


Figure 1: Prediction-centered calibration fails to distinguish decision-sensitive regions. The two predictions have the same ℓ_2 score relative to θ, but only $\widehat{\theta}^{\text{switch}}$ changes the optimal decision.

remains its unique minimax-regret solution. Thus, exact-regret calibration bounds nominal regret without changing the decision through reoptimization. This reveals a distinction between certifying a fixed decision and constructing uncertainty for choosing among alternative decisions.

We therefore consider two-sided information: both the regret of $\widehat{x}$ under the realized parameter θ and the suboptimality of the hindsight optimizer x_θ under the prediction $\widehat{\theta}$. We introduce a reflected Fenchel–Young (FY) score that incorporates these comparisons for convex optimization with uncertainty entering linearly in the objective. The score admits an exact three-gap decomposition: nominal regret, the reverse optimality gap of the hindsight optimizer under the prediction, and the regret of the nominal decision under a reflected parameter. For unregularized linear optimization, symmetric reflection ($\lambda = 2$) recovers the SPO+ loss (Elmachtoub & Grigas, 2022). We calibrate the score using split conformal prediction and choose a decision by minimizing worst-case regret over the resulting uncertainty set. Our contributions are summarized as follows.

1. Self-centering and reflected decision-aware calibration: We identify a structural limitation of exact-regret calibration: a full regret sublevel can provide a tight nominal certificate while uniquely retaining the nominal decision under minimax-regret optimization. We introduce a reflected FY score that augments exact regret with reverse-optimality and stability information, preserving additional decision-relevant geometry for robust optimization.
2. Guaranteed and certified robust refinement: We establish finite-sample marginal regret guarantees for decisions obtained from the reflected FY set. We characterize when the nominal decision remains robust-optimal and when reoptimization strictly reduces worst-case regret over the same set, with quantitative bounds on this reduction. The outer minimization is convex, although finding worst-case parameters can require nonconvex optimization. We derive a finite-scenario representation and a master–pricing procedure with verifiable optimality certificates when global pricing bounds are available.
3. Empirical validation: Experiments on a simplex allocation problem verify the predicted self-centering phenomenon and when reflected calibration enables refinement. Across synthetic portfolio and real-data solar-grid allocation problems, the reflected FY construction improves decision quality over residual-based robust baselines.

Related work. Prediction-driven optimization uses contextual information to estimate unknown parameters before solving the corresponding optimization problem (Bertsimas & Kallus, 2020; Liu et al., 2021). However, predictive accuracy need not translate into decision quality: small prediction errors can change the optimal decision and lead to substantial regret (Smith & Winkler, 2006). The smart predict-then-optimize framework trains predictors for linear optimization using the convex SPO+ surrogate for decision regret (Elmachtoub & Grigas, 2022). Related work studies risk guarantees and calibration (Ho-Nguyen & Kılınç-Karzan, 2022), as well as integrated estimation–optimization formulations (Qi et al., 2026). Decision-focused learning pursues this objective through differentiable optimization (Donti et al., 2017; Wilder et al., 2019), ranking losses aligned with downstream cost (Mandi et al., 2022), and solver-free or robust-constraint surrogates (Berden et al., 2025; Im et al., 2025). Complementary work debiases policy-performance estimates to guide policy

selection (Gupta et al., 2024). These approaches primarily modify how the predictor is trained. Our work keeps the predictor fixed and constructs decision-aware uncertainty sets for the downstream robust optimization problem.

Conformal prediction provides finite-sample marginal coverage under exchangeability (Shafer & Vovk, 2008; Angelopoulos & Bates, 2023), with extensions to risk control (Angelopoulos et al., 2024). In contextual robust optimization, uncertainty sets are calibrated from prediction residuals (Johnstone & Cox, 2021; Sun et al., 2023) or learned through downstream decision loss (Chenreddy & Delage, 2024; Yeh et al., 2025). Related applications include robustness control, decision risk assessment, and inverse optimization (Hu et al., 2026; Zhou et al., 2026; Lin et al., 2024). Other work asks how prediction sets should serve decisions: Cortes-Gomez et al. (2025) optimize their downstream costs, Kiyani et al. (2025) derive max–min policies for risk-averse decisions, and Wang & Dobriban (2026) minimize worst-case expected loss while accounting for outcomes outside the set. We study how the calibration score affects the resulting decision, showing why nominal-regret calibration can leave it unchanged and how the reflected Fenchel–Young score permits worst-case regret refinement.

Fenchel–Young losses connect regularized prediction with convex conjugacy (Blondel et al., 2020) and have been used in inverse optimization (Li et al., 2025; Sakaue et al., 2025). The related literature on Bregman centers provides dual representations for finite sets (Nock & Nielsen, 2005) and characterizes optimality through active farthest points under Legendre assumptions (Bauschke et al., 2010). We extend these characterizations to nonsmooth objectives and constrained decisions, and use them to study when reflected-FY calibration supports minimax-regret refinement.

2 From Exact Regret to Reflected Fenchel–Young Calibration

2.1 Problem Setup and Preliminaries

Let $\mathcal{Z}$ denote the context space and Θ the uncertain parameter space. The context z includes all model inputs known at decision time, and a fixed predictor $\widehat{f} : \mathcal{Z} \to \Theta$ may use any subset of them. We consider contextual convex optimization problems in which uncertainty enters linearly in the objective, while known convex regularization and feasibility constraints are collected in a proper, closed, convex function Ψ. These components may depend on z. We suppress this dependence in Ψ, its conjugate A and the optimal solution x_θ:

$$x_\theta \in \arg\max_x \{\langle \theta, x\rangle - \Psi(x)\}, \qquad A(\theta) := \Psi^*(\theta) = \sup_x \{\langle \theta, x\rangle - \Psi(x)\}. \tag{1}$$

This formulation covers many standard models, including linear programs ($\Psi = I_{\mathcal{X}}$), regularized allocation problems, and constrained quadratic programs, while feasibility can be encoded by an indicator term in Ψ (Boyd & Vandenberghe, 2004). The conjugate A is the optimal-value function, and Predict-then-Optimize (PTO) simply implements the nominal decision $\widehat{x} := x_{\widehat{\theta}}$.

Split conformal calibration. We use split conformal prediction to construct the uncertainty sets (Angelopoulos & Bates, 2023). Fix the predictor $\hat{f}$ and the conformal score $s(z, \theta)$ measuring the discrepancy, and evaluate the score on held-out calibration pairs: $s_i := s(z_i, \theta_i), i \in [n]$, with order statistics $s_{(1)} \le \cdots \le s_{(n)}$. For a miscoverage level $\alpha \in (0, 1)$, set

$$k = \lceil (n+1)(1-\alpha) \rceil, \qquad \eta_\alpha := s_{(k)}, \qquad s_{(n+1)} := +\infty. \tag{2}$$

If the calibration and test pairs are exchangeable, the standard results give $\mathbb{P}\{s(z_{\text{test}}, \theta_{\text{test}}) \le \eta_\alpha\} \ge 1-\alpha$, so the sublevel set $\{\theta : s(z, \theta) \le \eta_\alpha\}$ is a valid $(1-\alpha)$-uncertainty set. The score determines both the shape of the set and its relevance to the downstream decision.

Fenchel–Young gaps. We express decision suboptimality through the Fenchel–Young gap (Blondel et al., 2020) in this paper:

$$L_\Psi(u; x) := A(u) + \Psi(x) - \langle u, x\rangle \ge 0. \tag{3}$$

For $u \in \operatorname{dom} A$ and $x \in \operatorname{dom} \Psi$, this gap measures the suboptimality of the decision x under parameter u and is zero exactly when x is optimal. When A is differentiable at v and finite at u, define the Bregman divergence $D_A(u, v) := A(u) - A(v) - \langle \nabla A(v), u - v\rangle$.

2.2 Exact-Regret Calibration Certifies but Can Preclude Refinement

We measure the performance of the nominal decision $\widehat{x}$ by its regret under the realized parameter:

$$\mathrm{Reg}(\widehat{x};\theta) := \big[\langle\theta, x_\theta\rangle - \Psi(x_\theta)\big] - \big[\langle\theta,\widehat{x}\rangle - \Psi(\widehat{x})\big] = L_\Psi(\theta;\widehat{x}). \tag{4}$$

For an unregularized linear program, it reduces to $\langle\theta, x_\theta - \widehat{x}\rangle$. If A is differentiable at $\widehat{\theta}$, it equals $D_A(\theta,\widehat{\theta})$. Therefore, a direct approach is to use the nominal regret in (4) as the conformal score. The split conformal calibration then bounds the nominal regret by η_α with marginal probability at least $1-\alpha$. We ask whether minimizing worst-case regret over this regret set can improve on $\hat{x}$.

Theorem 1 (Self-centering of exact-regret conformal sets)**.** *For $\eta > 0$, let $B_\eta(\widehat{\theta}) := \{\theta : L_\Psi(\theta;\widehat{x}) \le \eta\}$ be the full regret sublevel over the entire positive-dimensional effective parameter space. Suppose either (i) A is differentiable at $\widehat{\theta}$, $B_\eta(\widehat{\theta})$ is compact with $\widehat{\theta}$ in its interior, and A is finite and continuous on a neighborhood of this set; or (ii) $\Psi = I_{\mathcal{X}}$ for a compact convex decision set $\mathcal{X}$. Then*

$$\operatorname*{arg\,min}_{x\in\mathrm{dom}\,\Psi} \sup_{\theta\in B_\eta(\widehat{\theta})} L_\Psi(\theta;x) = \{\widehat{x}\}, \qquad \sup_{\theta\in B_\eta(\widehat{\theta})} L_\Psi(\theta;\widehat{x}) = \eta. \tag{5}$$

Theorem 1 answers this question negatively: (5) identifies $\widehat{x}$ as the unique minimax-regret decision and gives worst-case regret η. To understand whether this limitation stems from a conservative nominal certificate, we next ask whether another set with the same coverage can give a tighter nominal regret bound.

Proposition 1 (Nominal regret certification)**.** *Fix a context and let θ follow its conditional distribution. Define $\eta^* := q_{1-\alpha}(\mathrm{Reg}(\widehat{x};\theta))$, where $q_{1-\alpha}(Y) := \inf\{t : \mathbb{P}(Y \le t) \ge 1-\alpha\}$ is the left quantile. For every Borel set U satisfying $\mathbb{P}(\theta\in U) \ge 1-\alpha$, we have $\sup_{\theta\in U}\mathrm{Reg}(\widehat{x};\theta) \ge \eta^*$. Equality is attained by $U = \{\theta : \mathrm{Reg}(\widehat{x};\theta) \le \eta^*\}$. A nondecreasing transformation of nominal regret produces open or closed regret sublevels. At positive finite cutoffs satisfying Theorem 1, these sets retain the same minimax-regret decision.*

Together, Theorem 1 and Proposition 1 distinguish nominal certification from decision refinement: in the fixed-context population comparison, an exact-regret sublevel attains the smallest nominal certificate, yet under the conditions of Theorem 1, reoptimization over that set uniquely retains $\widehat{x}$. This motivates retaining decision comparisons beyond nominal regret alone.

2.3 Decision-Aware Calibration with the Reflected Fenchel–Young Score

The self-centering result motivates examining both sides of the decision comparison. Nominal regret measures how $\widehat{x}$ performs under θ; the reverse comparison asks how the hindsight optimizer x_θ performs under the prediction $\widehat{\theta}$. We measure the latter by the *reverse gap*, $\Delta_{\widehat{\theta}}(x_\theta) := L_\Psi(\widehat{\theta};x_\theta)$. If A is differentiable at both θ and $\widehat{\theta}$, the two gaps are Bregman divergences with their arguments reversed:

$$\mathrm{Reg}(\widehat{x};\theta) = D_A(\theta,\widehat{\theta}), \qquad \Delta_{\widehat{\theta}}(x_\theta) = D_A(\widehat{\theta},\theta). \tag{6}$$

We combine these comparisons by evaluating the two decisions using the FY gap in (3) under a common parameter. For $\lambda > 1$, matching their scaled FY difference to $\mathrm{Reg}(\hat{x};\theta) + \lambda\Delta_{\hat{\theta}}(x_\theta)$ yields

$$\begin{aligned} r_\lambda(\theta,\widehat{\theta}) &:= \frac{\lambda\widehat{\theta}-\theta}{\lambda-1} = \widehat{\theta} + \frac{\widehat{\theta}-\theta}{\lambda-1}, \\ \mathrm{Reg}(\widehat{x};\theta) + \lambda\Delta_{\widehat{\theta}}(x_\theta) &= (\lambda-1)\big\{L_\Psi(r_\lambda;x_\theta) - L_\Psi(r_\lambda;\widehat{x})\big\}. \end{aligned} \tag{7}$$

This reflection in (7) is unique within this coefficient-matched FY comparison. Dropping the subtracted nonnegative term gives an upper bound on their weighted sum, which we use as the calibration score. We call θ admissible if $A(\theta)$ and $A(r_\lambda)$ are finite and the optimum at θ is attained by the selected x_θ. For admissible θ, define the reflected FY score and its sublevel set by

$$S_\lambda(\theta,\widehat{\theta}) := (\lambda-1)L_\Psi(r_\lambda;x_\theta), \qquad \mathcal{U}_{\lambda,\eta}(\widehat{\theta}) := \{\vartheta\in\Theta : S_\lambda(\vartheta,\widehat{\theta}) \le \eta\}. \tag{8}$$

Here $\eta \ge 0$ is a score threshold and conformal calibration uses $s(z,\theta) = S_\lambda(\theta,\widehat{f}(z))$, evaluated with the model at z, to select $\eta_{\alpha,\lambda}$. The following decomposition identifies the additional decision gap introduced by this upper bound.

Theorem 2 (Three-gap decomposition)**.** *For every admissible θ,*

$$S_\lambda(\theta, \widehat{\theta}) = \underbrace{\mathrm{Reg}(\widehat{x}; \theta)}_{\text{nominal regret}} + \lambda \underbrace{\Delta_{\widehat{\theta}}(x_\theta)}_{\text{reverse gap}} + (\lambda - 1) \underbrace{L_\Psi(r_\lambda; \widehat{x})}_{\text{reflected stability gap}} . \tag{9}$$

If A is differentiable at θ, then $S_\lambda(\theta, \widehat{\theta}) = (\lambda - 1) D_A(r_\lambda, \theta)$.

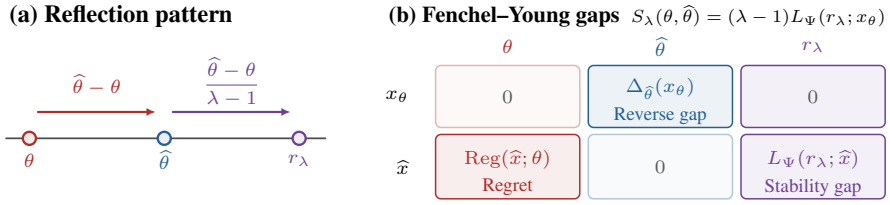


Figure 2: Reflection pattern and the three-gap identity.

All three gaps are nonnegative. The reverse gap $\Delta_{\widehat{\theta}}(x_\theta)$ evaluates the hindsight decision x_θ under the prediction $\hat{\theta}$, while the reflected stability gap $L_\Psi(r_\lambda; \widehat{x})$ evaluates the nominal decision $\hat{x}$ under the reflected parameter r_λ. The nominal-regret term underlies regret certification, while the reverse-gap term also restricts the hindsight optimizers admitted by the calibrated set, as quantified by the localization bound below and the screening rule in Corollary 4.

Reflection geometry and the connection to SPO+ loss. Equation (7) places θ and r_λ on opposite sides of the prediction $\widehat{\theta}$. At $\lambda = 2$, $r_2 = 2\widehat{\theta} - \theta$, so the reflection is symmetric about the prediction (Figure 2). For linear optimization, this special case of $\lambda = 2$ recovers SPO+, the convex surrogate of decision regret used to train predictors in Elmachtoub & Grigas (2022).

Corollary 1 (Reduction to the SPO+ loss)**.** *Consider the linear program $\min_{x \in \mathcal{X}} c^\top x$ with compact feasible set $\mathcal{X}$, realized cost vector c, and prediction $\hat{c}$. For $\Psi = I_{\mathcal{X}}$ (the indicator of $\mathcal{X}$), $\theta = -c$ and $\widehat{\theta} = -\widehat{c}$, the same optimizer selection in both losses gives*

$$S_2(-c, -\widehat{c}) = \ell_{\mathrm{SPO+}}(\widehat{c}, c). \tag{10}$$

This algebraic identity also holds for nonconvex $\mathcal{X}$.

The symmetric choice $\lambda = 2$ also has a local interpretation. When A is twice continuously differentiable near $\widehat{\theta}$, the local quadratic coefficient of S_λ is proportional to $\lambda^2/(\lambda - 1)$, which is uniquely minimized at $\lambda = 2$ (Appendix A.2). Thus, $\lambda = 2$ provides a principled default, while other values of $\lambda > 1$ may be selected on a separate tuning split and fixed before conformal calibration.

The nonnegative decomposition in (9) allows a single conformal coverage event to control all three gaps simultaneously. Corollary 2 subsequently transfers the resulting regret certificate to the robust decision.

Theorem 3 (Joint conformal decision guarantee)**.** *Let $\eta_{\alpha,\lambda}$ be the standard conformal threshold in (2) computed from n calibration scores. Suppose the predictor, λ, and score are fixed before calibration, and the calibration and test pairs (z, θ) are exchangeable. Then*

$$\mathbb{P}\Big\{ S_\lambda(\theta_{\mathrm{test}}, \widehat{\theta}_{\mathrm{test}}) \le \eta_{\alpha,\lambda} \Big\} \ge \frac{\lceil (n+1)(1-\alpha) \rceil}{n+1} \ge 1 - \alpha.$$

At the test context, let $\widehat{x} = x_{\widehat{\theta}_{\mathrm{test}}}$ and $r_\lambda = r_\lambda(\theta_{\mathrm{test}}, \widehat{\theta}_{\mathrm{test}})$. If $\eta_{\alpha,\lambda} < \infty$ almost surely, the following bounds hold simultaneously with the same probability:

$$\mathrm{Reg}(\widehat{x}; \theta_{\mathrm{test}}) \le \eta_{\alpha,\lambda}, \qquad \Delta_{\widehat{\theta}_{\mathrm{test}}}(x_{\theta_{\mathrm{test}}}) \le \frac{\eta_{\alpha,\lambda}}{\lambda}, \qquad L_\Psi(r_\lambda; \widehat{x}) \le \frac{\eta_{\alpha,\lambda}}{\lambda - 1}. \tag{11}$$

When Ψ is μ-strongly convex with respect to $\|\cdot\|$ uniformly across contexts and $\eta_{\alpha,\lambda} < \infty$ almost surely, the coverage event also implies:

$$\|x_{\theta_{\mathrm{test}}} - \widehat{x}\| \le \sqrt{\frac{2\eta_{\alpha,\lambda}}{\lambda \mu}}. \tag{12}$$

The localization bound depends on the ratio $\eta_{\alpha,\lambda}/\lambda$ rather than on the score threshold alone. The appendix extends this bound to more general growth conditions on the reverse optimality gap.

3 Robust Refinement with Conformal Fenchel–Young Sets

We now use the calibrated sets in (8) to choose a decision x. A standard robust formulation maximizes the worst-case objective value over the uncertainty set. Instead, we minimize worst-case regret, comparing each decision with the optimum under the same parameter realization, as formalized in (13). We first establish the marginal regret guarantee for the resulting decision, then characterize when reoptimization improves on $\widehat{x}$ over the same set, quantify this gain, and develop a solution procedure with optimality certificates.

3.1 The Induced Robust Optimization Problem

Fix a prediction $\hat{\theta} = \hat{f}(z)$ and a finite threshold $\eta \geq 0$, and write $U := \mathcal{U}_{\lambda,\eta}(\widehat{\theta})$. For a decision x, define its worst-case regret and the optimal robust value as

$$\rho_U(x) := \sup_{\theta \in U} L_\Psi(\theta; x), \qquad R^*(U) := \inf_{x \in \operatorname{dom} \Psi} \rho_U(x). \tag{13}$$

The three-gap identity in (9) bounds the worst-case regret of the nominal decision $\widehat{x}$:

$$R^*(U) \leq \rho_U(\widehat{x}) \leq \eta. \tag{14}$$

We next examine whether this minimization is convex and whether its optimizer retains the marginal regret guarantee.

Proposition 2. *For every nonempty $U \subseteq \operatorname{dom} A$, the function ρ_U is convex in x and, for every $x \in \operatorname{dom} \Psi$, satisfies $\rho_U(x) = \rho_{\operatorname{conv}(U)}(x)$. Consequently, $R^*(U) = R^*(\operatorname{conv}(U))$.*

Corollary 2 (Marginal regret guarantee)**.** *Under the assumptions of Theorem 3, suppose that $\eta_{\alpha,\lambda} < \infty$ almost surely. Let $U_z = \mathcal{U}_{\lambda,\eta_{\alpha,\lambda}}(\widehat{f}(z))$ and let $x^{\mathrm{rob}}(z)$ be an exact robust optimizer. If the worst-case value is measurable in the context z and the calibration data, then*

$$\mathbb{P}\big\{ L_\Psi(\theta_{\mathit{test}}; x^{\mathrm{rob}}(z_{\mathit{test}})) \leq R^*(U_{z_{\mathit{test}}}) \leq \eta_{\alpha,\lambda} \big\} \geq 1 - \alpha. \tag{15}$$

Thus, the uncertainty set U need not be convex for the outer minimization to be convex or for the marginal guarantee to hold. In practice, exact optimization is not essential: a certified ε-suboptimal feasible policy inherits the same marginal regret guarantee with threshold $\eta_{\alpha,\lambda} + \varepsilon$ (Appendix A.3). The main computational difficulty is evaluating $\rho_U(x)$ over the generally nonconvex set U.

3.2 When Is Robustification Necessary?

We first ask whether the nominal decision $\widehat{x}$ already minimizes worst-case regret over the calibrated set. Assume that U is compact and that A is finite and continuous on U. Let $W_U(x) := \arg\max_{\theta \in U} L_\Psi(\theta; x)$ denote the active worst-case parameters for decision x.

Theorem 4 (Robust optimality)**.** *Under the preceding assumptions, for every $x \in \operatorname{dom} \Psi$,*

$$\partial \rho_U(x) = \partial \Psi(x) - \operatorname{conv} W_U(x), \tag{16}$$

$$x \in \arg\min \rho_U \iff \partial \Psi(x) \cap \operatorname{conv} W_U(x) \neq \varnothing. \tag{17}$$

If Ψ is differentiable at $\widehat{x}$, nominal optimality gives $\nabla \Psi(\widehat{x}) = \widehat{\theta}$, and the condition reduces to $\widehat{\theta} \in \operatorname{conv} W_U(\widehat{x})$.

Positive worst-case regret does not by itself justify changing the nominal decision. What matters is whether the active worst-case parameters in $W_U(\widehat{x})$ can be balanced to satisfy the nominal first-order optimality condition in (17). We measure the failure of this condition by the movement residual

$$\mathfrak{M}(\widehat{x}; U) := \operatorname{dist}_*\big(\partial \Psi(\widehat{x}), \operatorname{conv} W_U(\widehat{x})\big), \tag{18}$$

where dist_* denotes distance under the dual norm $\|\cdot\|_*$. Under the stated assumptions, $\mathfrak{M}(\widehat{x}; U) = 0$ if and only if the nominal decision is robust-optimal. We next ask how many active scenarios suffice to certify optimality and how the residual measures the rate of local descent.

Corollary 3 (Finite screening and local descent)**.** *Under the assumptions of Theorem 4, the nominal decision is robust-optimal if and only if there exist* $\theta^1, \dots, \theta^k \in W_U(\widehat{x})$ *and weights* $p_i \geq 0$ *with* $\sum_i p_i = 1$ *such that*

$$\sum_{i=1}^{k} p_i \theta^i \in \partial\Psi(\widehat{x}).$$

The certificate can be chosen with $k \leq \dim\big(\operatorname{aff} W_U(\widehat{x})\big) + 1$ *active parameters. If* Ψ *is finite on a relative neighborhood of* $\widehat{x}$ *within the feasible affine hull, then the one-sided directional derivative of* ρ_U *satisfies*

$$\min_{\|v\|\leq 1} \rho'_U(\widehat{x}; v) = -\mathfrak{M}(\widehat{x}; U). \tag{19}$$

Corollary 4 (Reverse-gap screening)**.** *For a discrete or linear program, suppose* $\Delta_{\widehat{\theta}}(x_\theta) \geq \Gamma > 0$ *for every admissible selected optimizer* $x_\theta \neq \widehat{x}$*. Then* $\eta \leq \lambda\Gamma$ *implies* $\rho_U(\widehat{x}) = 0$*.*

The preceding results determine whether the nominal decision already minimizes worst-case regret over U. We next use the same movement residual to characterize whether reoptimization has strictly positive value and to quantify the possible reduction.

3.3 The Value of Robustification

We now ask how much can be gained through robustifying the problem. For $\rho_U(\widehat{x}) < \infty$, define $\mathcal{G}_{\mathrm{rob}}(U) := \rho_U(\widehat{x}) - R^*(U)$ as the reduction in worst-case regret obtained by optimizing over the same uncertainty set rather than retaining the nominal decision $\hat{x}$.

Proposition 3 (Value of robustification)**.** *Under the assumptions of Theorem 4,* $\mathcal{G}_{\mathrm{rob}}(U) > 0$ *if and only if* $\mathfrak{M}(\widehat{x}; U) > 0$*. If* Ψ *is* μ*-strongly convex in* $\|\cdot\|$*, the robust decision* x^{rob} *is unique and*

$$\frac{\mu}{2}\|x^{\mathrm{rob}} - \widehat{x}\|^2 \leq \mathcal{G}_{\mathrm{rob}}(U) \leq \frac{\mathfrak{M}(\widehat{x}; U)^2}{2\mu}. \tag{20}$$

Proposition 3 shows that this same imbalance also determines whether robustification has strictly positive value. For a fixed strong-convexity modulus μ, a small movement residual $\mathfrak{M}(\widehat{x}; U)$ leaves little room for reducing worst-case regret. Conversely, a large change from $\hat{x}$ to x^{rob} must be accompanied by a correspondingly large reduction over the same set.

The gain from optimization is distinct from the conservatism of the uncertainty set itself. For this population comparison, fix a context z and let θ follow its conditional distribution. Define $V_\alpha := \inf_{x\in\mathrm{dom}\,\Psi} q_{1-\alpha}(L_\Psi(\theta; x))$, the smallest achievable $(1-\alpha)$-quantile of regret. Any fixed set U with coverage at least $1-\alpha$ under this distribution satisfies $V_\alpha \leq R^*(U)$. Consequently,

$$\rho_U(\widehat{x}) - V_\alpha = \underbrace{\mathcal{G}_{\mathrm{rob}}(U)}_{\text{recovered by optimization}} + \underbrace{R^*(U) - V_\alpha}_{\text{remaining set conservatism}}. \tag{21}$$

This fixed-context benchmark is distinct from the marginal guarantee in Corollary 2. Exact-regret sets attain the smallest nominal regret bound but remain self-centering under Theorem 1. Reflected FY sets may increase this bound while allowing it to be reduced by optimization. A strict reduction over U need not outweigh the increase in its nominal regret bound; Equation (30) in Appendix A.5 gives the exact trade-off between this certificate cost and the gain from reoptimization.

3.4 Characterizing and Computing the Robust Decision

When the nominal decision is not robust-optimal, we solve the minimax-regret problem in Equation (13). The outer minimization is convex, but the uncertainty set may contain infinitely many parameters, and finding a worst-case parameter can require nonconvex optimization. We first identify a finite collection of scenarios that supports a robust optimum, then use this representation to construct a scenario-generation procedure.

Under the assumptions of Theorem 4, suppose the robust minimum is attained. Let $\mathcal{P}(U)$ denote the Borel probability measures on U. The following result represents the robust value using distributions supported on active worst-case parameters.

Theorem 5 (Adversarial barycenter)**.** *Under the preceding assumptions,*

$$R^*(U) = \max_{\nu \in \mathcal{P}(U)} \{\mathbb{E}_\nu[A(\theta)] - A(\mathbb{E}_\nu[\theta])\}. \tag{22}$$

For a robust optimizer x^{rob}, there exists an optimal distribution ν^* with finite support in $W_U(x^{\mathrm{rob}})$. Its barycenter $m^* := \mathbb{E}_{\nu^*}[\theta]$ satisfies $m^* \in \partial\Psi(x^{\mathrm{rob}})$. Every support point attains the robust regret: $L_\Psi(\theta; x^{\mathrm{rob}}) = R^*(U)$ for all $\theta \in \mathrm{supp}(\nu^*)$. If A is differentiable at m^*, then $x^{\mathrm{rob}} = x_{m^*}$.

Theorem 5 expresses the active-scenario balance in Theorem 4 through an adversarial distribution. The robust decision x^{rob} is therefore optimal for the barycenter m^* of a finite set of active worst-case scenarios. The weights describe an adversarial balance between scenarios, rather than their occurrence probabilities.

Scenario generation and optimality certificates. We solve a restricted master problem over selected scenarios and use a pricing problem to search for additional worst-case parameters. For a nonempty scenario set $U_J = \{\theta^1, \ldots, \theta^J\} \subseteq U$, solve the convex restricted master

$$R_J := \min_{x,t} \left\{ t : L_\Psi(\theta^j; x) \le t,\ j = 1, \ldots, J \right\}. \tag{23}$$

Let x_J be an optimal solution of (23). The pricing problem evaluates the candidate decision over the full set by computing $\rho_U(x_J) = \sup_{\theta \in U} L_\Psi(\theta; x_J)$. If the pricing value $\rho_U(x_J)$ exceeds the master value R_J by more than the prescribed tolerance, we add a worst-case parameter to the master and repeat. The following bounds provide a stopping certificate at the desired optimality tolerance.

Proposition 4 (Scenario-generation certificate)**.** *For every scenario set $U_J \subseteq U$ and an optimal master solution x_J,*

$$R_J \le R^*(U) \le \rho_U(x_J). \tag{24}$$

For $\varepsilon \ge 0$, if pricing is solved globally and $\rho_U(x_J) - R_J \le \varepsilon$, then $\rho_U(x_J) - R^*(U) \le \varepsilon$. More generally, any global upper bound on $\rho_U(x_J)$ can replace the exact pricing value in this gap test.

For approximate solves, the stopping test based on (24) uses a certified master lower bound and a global upper bound on $\rho_U(x_J)$. Heuristic pricing can identify additional scenarios, but feasible adversarial values alone do not certify optimality. These certificates also connect computation to statistical validity: under the assumptions of Corollary 2, a feasible policy with an optimality gap certified to be at most ε almost surely inherits the marginal regret guarantee with threshold $\eta_{\alpha,\lambda} + \varepsilon$ (Appendix A.3).

4 Numerical Experiments

We evaluate decision refinement and out-of-sample regret on simplex allocation, constrained portfolio optimization, and real solar-grid data. Baselines include PTO, conformal robust optimization with residual ℓ_2, box, and ellipsoid scores (Johnstone & Cox, 2021; Sun et al., 2023), and k-nearest neighbor robust optimization (KNN RO; Ohmori 2021). Prediction-based methods share a predictor; KNN RO uses local training outcomes. The comparisons evaluate the complete FY calibration and minimax-regret framework, whereas robust baselines maximize worst-case objective value. Appendix B.1 discusses the comparison scope. We use $\lambda = 2$ and report mean regret, its 95th percentile (Q95), and CVaR_{95}, the average regret in the worst 5% of cases. Protocols, parameter settings, and detailed results are in Appendix B.

A simplex allocation problem. We distribute a unit budget over two assets to maximize total return. On 5,000 test contexts with $\alpha = 0.05$, Figure 3(a) plots allocation transferred away from PTO against the absolute predicted return difference normalized by $\eta_{\alpha,2}/2$; the horizontal scale is linear up to one and logarithmic thereafter. Direct regret calibration retains PTO throughout, whereas FY reallocates weight below the switching threshold of one. Above it, the nominally preferred asset remains optimal over the FY set. Panel (b) compares worst-case regret before and after optimization over the same FY set, with the vertical arrow measuring $\mathcal{G}_{\mathrm{rob}}(U)$. FY changes 423 decisions, with an average $\mathcal{G}_{\mathrm{rob}}$ of 0.023 overall and 0.272 among those that move. The remaining contexts lie at the origin, where PTO has zero worst-case regret. The exact QP comparison in Appendix B.3 connects this behavior to active-set changes.

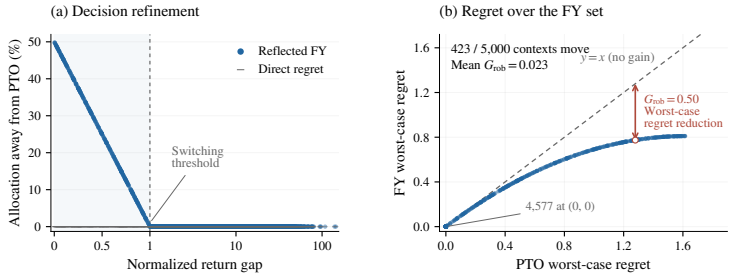


Figure 3: Decision refinement and worst-case regret in the two-asset LP.

Constrained portfolio optimization. We next consider a 20-asset portfolio QP with a quadratic penalty and position limits, using a controlled mean–variance model (Markowitz, 1952). Historical training returns are shifted relative to the current calibration and test distribution; Table 1 reports the shift $\kappa = 2$ on 1,000 common test instances at $\alpha = 0.1$. FY reduces mean regret by 8.93% and CVaR_{95} by 5.51% relative to PTO, and improves both criteria over all robust baselines. Its average Euclidean adjustment from PTO is 0.00639, with score coverage of 89.3%. These results show that the reflected-FY regret framework can improve both average and upper-tail decision quality in a problem with quadratic curvature. Replacing the heuristic decisions with globally certified references leaves the reported mean and CVaR_{95} unchanged at the displayed precision (Table 6(b)).

Table 1: Portfolio QP regret $\times 10^3$.

Method	Mean	Q95	CVaR95
PTO	3.172	6.134	7.307
FY	**2.888**	**5.755**	**6.904**
Residual ℓ_2	5.282	9.096	10.304
Residual box	3.236	6.339	7.468
Residual ellipsoid	5.274	9.083	10.290
KNN RO	5.339	9.301	10.431

Table 2: Solar-grid regret $\times 10^3$ and coverage.

Method	Mean	Q95	CVaR95	Coverage
PTO	**78.694**	260.635	371.787	–
FY	80.030	**205.938**	**246.440**	87.50%
Residual ℓ_2	80.064	206.245	247.074	82.29%
Residual box	117.741	354.090	448.768	85.42%
Residual ellipsoid	93.537	236.479	307.520	79.17%
KNN RO	93.537	236.479	307.519	79.17%

Real-data solar-grid allocation. Finally, we use the publicly released monthly solar-installation records for four Indiana substations from Zhou et al. (2026). Following their continuous relaxation, we allocate a budget of two upgrades to minimize a smoothed overload penalty. Twelve chronological training/calibration/test windows yield 96 test months, with $\alpha = 0.2$. Table 2 shows that FY reduces CVaR_{95} by 33.71% relative to PTO while increasing mean regret by 1.70%. FY and residual ℓ_2 attain similar regret, with empirical coverage of 87.50% and 82.29%, respectively, against an 80% target. The substations have equal costs and nearly identical predicted risks relative to calibrated uncertainty, so both methods allocate approximately one half of an upgrade to each station. In the symmetric limit, the two uncertainty sets induce the same allocation despite their different shapes.

5 Conclusion and Discussion

We distinguish regret certification of a nominal decision from its refinement through minimax-regret optimization. Exact-regret calibration can provide a tight nominal certificate while leaving the decision uniquely minimax optimal under the stated conditions. We develop reflected Fenchel–Young conformal scores for convex optimization with uncertainty entering linearly in the objective. The reflected score retains additional decision comparisons and gives finite-sample marginal bounds on three decision gaps after conformal calibration, with a regret guarantee for the resulting robust decision. Our analysis characterizes when reoptimization strictly reduces worst-case regret over the calibrated set and connects the active-scenario optimality condition to a finite-support adversarial representation and a master–pricing procedure with optimality certificates when global pricing bounds are available. The experiments validate our theoretical findings on self-centering and decision refinement and demonstrate gains in decision quality over residual-based robust baselines.

AI USE STATEMENT

Large language models (LLMs) were used as supporting tools throughout this work. They were primarily used to polish the writing and improve the clarity and presentation of the manuscript. During the early stages of the project, discussions with LLMs also assisted in brainstorming research ideas, and LLM-based tools were used to support code drafting and debugging. For mathematical analysis, LLMs were also used to assist with proof exploration and to check intermediate mathematical derivations. All mathematical claims and proofs were independently verified, revised, and finalized by the authors. The authors remain fully responsible for the scientific content, experimental results, and conclusions of this work.

ETHICS STATEMENT

This work uses synthetic and publicly available data and does not involve human subjects or sensitive personal information. We do not identify specific ethical concerns or direct risks of harm arising from the study. The research was conducted in accordance with the ICLR Code of Ethics.

REPRODUCIBILITY STATEMENT

Proofs of the theoretical results under the stated assumptions are provided in Appendix A. Appendix B documents the data sources and preprocessing, synthetic data generation, data splits, prediction models, baselines, hyperparameters, optimization procedures, and evaluation metrics, together with solver accuracy checks. An anonymized codebase with scripts for reproducing the reported experiments will be provided with the submission as supplementary material.

REFERENCES

Anastasios Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. Conformal risk control. In *International Conference on Learning Representations*, volume 2024, pp. 55198–55218, 2024.

Anastasios N Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning*, 16(4):494–591, 2023.

Heinz H. Bauschke, Mason S. Macklem, Jason B. Sewell, and Xianfu Wang. Klee sets and Chebyshev centers for the right Bregman distance. *Journal of Approximation Theory*, 162(6):1225–1244, 2010.

Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*. Princeton University Press, 2009.

Senne Berden, Ali İrfan Mahmutoğulları, Dimos Tsouros, and Tias Guns. Solver-free decision-focused learning for linear optimization problems. In *Advances in Neural Information Processing Systems*, volume 38, pp. 158127–158145. Curran Associates, Inc., 2025.

Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

Mathieu Blondel, André F. T. Martins, and Vlad Niculae. Learning with Fenchel–Young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Abhilash Reddy Chenreddy and Erick Delage. End-to-end conditional robust optimization. In *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence*, volume 244 of *Proceedings of Machine Learning Research*, pp. 736–748. PMLR, 2024.

Santiago Cortes-Gomez, Carlos Miguel Patiño, Yewon Byun, Steven Wu, Eric Horvitz, and Bryan Wilder. Utility-directed conformal prediction: A decision-aware framework for actionable uncertainty quantification. In *The Thirteenth International Conference on Learning Representations*, 2025.

Priya L. Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 5490–5500. Curran Associates Inc., 2017.

Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1): 9–26, 2022.

Vishal Gupta, Michael Huang, and Paat Rusmevichientong. Debiasing in-sample policy performance for small-data, large-scale optimization. *Operations Research*, 72(2):848–870, 2024.

Nam Ho-Nguyen and Fatma Kılınç-Karzan. Risk guarantees for end-to-end prediction and optimization processes. *Management Science*, 68(12):8680–8698, 2022.

Yang Hu, Jieren Tan, Changliang Zou, Yajie Bao, and Haojie Ren. Conformal robustness control: A new strategy for robust decision. In *The Fourteenth International Conference on Learning Representations*, 2026.

Hyungki Im, Wyame Benslimane, and Paul Grigas. Smart surrogate losses for contextual stochastic linear optimization with robust constraints. In *Advances in Neural Information Processing Systems*, volume 38, pp. 75149–75174. Curran Associates, Inc., 2025.

Chancellor Johnstone and Bruce Cox. Conformal uncertainty sets for robust optimization. In *Conformal and Probabilistic Prediction and Applications*, pp. 72–90. PMLR, 2021.

Shayan Kiyani, George J Pappas, Aaron Roth, and Hamed Hassani. Decision theoretic foundations for conformal prediction: Optimal uncertainty quantification for risk-averse agents. In *International Conference on Machine Learning*, pp. 30943–30965. PMLR, 2025.

Zhehao Li, Yanchen Wu, Jian Chen, and Xiaojie Mao. A Fenchel-Young loss approach to data-driven inverse optimization. *arXiv preprint arXiv:2502.16120*, 2025.

Bo Lin, Erick Delage, and Timothy C. Y. Chan. Conformal inverse optimization. In *Advances in Neural Information Processing Systems*, volume 37, pp. 63534–63564, 2024.

Sheng Liu, Long He, and Zuo-Jun Max Shen. On-time last-mile delivery: Order assignment with travel-time predictors. *Management Science*, 67(7):4095–4119, 2021.

Jayanta Mandi, Víctor Bucarey, Maxime Mulamba Ke Tchomba, and Tias Guns. Decision-focused learning: Through the lens of learning to rank. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14935–14947. PMLR, 2022.

Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

Richard Nock and Frank Nielsen. Fitting the smallest enclosing Bregman ball. In *European Conference on Machine Learning*, pp. 649–656. Springer, 2005.

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.

Meng Qi, Paul Grigas, and Zuo-Jun (Max) Shen. Integrated conditional estimation-optimization. *Operations Research*, 74(3):1604–1625, 2026.

Shinsaku Sakaue, Han Bao, and Taira Tsuchiya. Revisiting online learning approach to inverse linear optimization: A Fenchel–Young loss perspective and gap-dependent regret analysis. In Yingzhen Li, Stephan Mandt, Shipra Agrawal, and Emtiyaz Khan (eds.), *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 46–54. PMLR, 03–05 May 2025.

Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(12):371–421, 2008.

James E Smith and Robert L Winkler. The optimizer's curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, 52(3):311–322, 2006.

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, 2023.

Tao Wang and Edgar Dobriban. Optimal decision-making based on prediction sets. In *Forty-third International Conference on Machine Learning*, 2026.

Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1658–1665. AAAI Press, 2019.

Christopher Yeh, Nicolas Christianson, Alan Wu, Adam Wierman, and Yisong Yue. End-to-end conformal calibration for optimization under uncertainty. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856.

Wenbin Zhou, Agni Orfanoudaki, and Shixiang Zhu. Conformalized decision risk assessment. In *The Fourteenth International Conference on Learning Representations*, 2026.

A Proofs and Complementary Analysis

This appendix follows Sections 2 and 3, providing proofs of the main results and further analysis of calibration cost, decision movement, and the value of calibration geometry.

Throughout the regret and FY analysis, we remove parameter directions that add the same constant to every feasible objective. Interiors and dimensions are taken in the effective space $E = \operatorname{span}(\operatorname{dom}\Psi - \operatorname{dom}\Psi)$, defined by the full feasible domain at the given context. Constraints are included in Ψ, and $\partial\Psi$ denotes its full convex subdifferential.

We use a fixed measurable optimizer rule whenever the same oracle parameter occurs at a given context and prediction. In contextual statements, the predictor, score, and λ are fixed before calibration, and all maps are measurable. The reflected score is set to $+\infty$ outside the admissible domain in (8).

A.1 Exact-Regret Calibration

Fenchel equality at an oracle solution gives

$$L_\Psi(u; x_v) = A(u) - A(v) - \langle x_v, u - v\rangle. \tag{25}$$

Thus, when A is differentiable at v, this loss equals $D_A(u, v)$.

Proof of Theorem 1. Write $f(\theta) = L_\Psi(\theta; \widehat{x})$ and $B = B_\eta(\widehat{\theta})$. In case (i), $f(\theta) = D_A(\theta, \widehat{\theta})$. Continuity on a neighborhood of the full sublevel implies $f = \eta$ on the boundary of B: a boundary point with $f < \eta$ would have a neighborhood contained in B. Hence $\sup_B L_\Psi(\theta; \widehat{x}) = \eta$.

To prove uniqueness, take any $x \neq \widehat{x}$, set $d = x - \widehat{x}$, and choose $\theta_d \in \arg\min_{\theta\in B}\langle\theta, d\rangle$. Since $\widehat{\theta}$ is interior and $d \neq 0$ in effective coordinates, θ_d lies on the boundary and $\langle\theta_d, d\rangle < \langle\widehat{\theta}, d\rangle$. Nominal optimality gives $\Psi(x) - \Psi(\widehat{x}) \geq \langle\widehat{\theta}, d\rangle$, and therefore

$$L_\Psi(\theta_d; x) = \eta + \Psi(x) - \Psi(\widehat{x}) - \langle\theta_d, d\rangle > \eta.$$

Every other decision therefore has strictly larger worst-case regret.

In case (ii), A is the support function of $\mathcal{X}$, so f is finite, nonnegative, and positively homogeneous. Since the effective space has positive dimension, choose $y \in \mathcal{X} \setminus \{\widehat{x}\}$. Then $f(y - \widehat{x}) \geq \|y - \widehat{x}\|_2^2 > 0$; scaling this direction shows that $\sup_B f = \eta$. For any $x \neq \widehat{x}$, set $d = x - \widehat{x}$. If $f(-d) > 0$, take $t = \eta/f(-d)$; then $-td \in B$ and

$$L_\Psi(-td; x) = f(-td) + t\|d\|_2^2 = \eta + t\|d\|_2^2 > \eta.$$

If $f(-d) = 0$, every $-td$, $t > 0$, belongs to B, and the same identity gives $\sup_B L_\Psi(\theta; x) = +\infty$. Thus $\widehat{x}$ is the unique minimizer, including when the nominal linear objective has multiple optimizers. □

Proof of Proposition 1. Let $Y = \operatorname{Reg}(\widehat{x}; \theta)$ and $M = \sup_{\theta\in U}\operatorname{Reg}(\widehat{x}; \theta)$. The inclusion $\{\theta \in U\} \subseteq \{Y \leq M\}$ implies $\mathbb{P}(Y \leq M) \geq 1 - \alpha$, hence $M \geq \eta^*$. The regret sublevel at η^* has the required coverage by right continuity of the distribution function and attains this bound.

For a nondecreasing transformation φ, the set $\{r : \varphi(r) \leq t\}$ is an initial interval, so its preimage under nominal regret is an open or closed regret sublevel. Theorem 1 applies to closed sublevels. For a positive cutoff c, the open sublevel is dense in the closed one: convexity gives $f((1-s)\theta + s\widehat{\theta}) \leq (1-s)c < c$ for $s \in (0, 1)$ and $f(\theta) \leq c$. Since each scenario loss is continuous, its supremum is the same over either sublevel. The self-centering conclusion therefore also holds for open sublevels. □

A.2 Reflected Scores and Conformal Guarantees

Proof of Theorem 2. Set $g = \Psi(x_\theta) - \Psi(\widehat{x})$ and $d = \widehat{x} - x_\theta$. Oracle optimality yields

$$\operatorname{Reg}(\widehat{x}; \theta) + \lambda\Delta_{\widehat{\theta}}(x_\theta) = (\lambda - 1)g + \langle\lambda\widehat{\theta} - \theta, d\rangle.$$

To express this quantity as a scaled FY difference, write $q\{L_\Psi(u; x_\theta) - L_\Psi(u; \widehat{x})\} = q\{g + \langle u, d\rangle\}$. Matching the coefficients of g and d uniquely determines $q = \lambda - 1$ and $u = r_\lambda$. Rearranging gives (9), whose three terms are nonnegative FY gaps. The Bregman representation follows from (25). □

Proof of Corollary 1. Let w_c be the selected minimizer of $c^\top w$ on $\mathcal{X}$ and $z(c) = c^\top w_c$. For $\Psi = I_{\mathcal{X}}$,

$$S_2(-c, -\widehat{c}) = \max_{w \in \mathcal{X}} (c - 2\widehat{c})^\top w + 2\widehat{c}^\top w_c - z(c),$$

which is the SPO+ loss with the same optimizer selection. The identity uses only attained linear optimization, so convexity of $\mathcal{X}$ is unnecessary. □

The symmetric choice $\lambda = 2$. The factor 2 in the SPO+ function was originally chosen so that the population minimizer coincides with the conditional mean rather than a positive rescaling of it (Elmachtoub & Grigas, 2022). Corollary 1 gives this choice a precise geometric interpretation. Within the reflection family $r_\lambda = \hat{\theta} + (\widehat{\theta} - \theta)/(\lambda - 1)$, $\lambda = 2$ is the unique choice for which $\widehat{\theta}$ is the midpoint of θ and r_λ and the map $\theta \mapsto r_\lambda$ is self-inverse. Accordingly, the central-symmetry assumption used in the SPO+ consistency analysis is invariance under this same reflection.

Applying the reflection twice gives $r_\lambda(r_\lambda(\theta, \widehat{\theta}), \widehat{\theta}) = \widehat{\theta} + (\theta - \widehat{\theta})/(\lambda - 1)^2$. Hence the reflection is self-inverse and places $\widehat{\theta}$ midway between the two parameters exactly at $\lambda = 2$. If A is twice continuously differentiable near $\widehat{\theta}$, a Taylor expansion with $\delta = \theta - \widehat{\theta}$ gives

$$S_\lambda(\widehat{\theta} + \delta, \widehat{\theta}) = \frac{\lambda^2}{2(\lambda - 1)} \delta^\top \nabla^2 A(\widehat{\theta}) \delta + o(\|\delta\|^2).$$

The factor $\lambda^2/(\lambda - 1)$ has its unique minimum at two. If A is globally quadratic with positive curvature, then $S_\lambda = \lambda^2 D_A(\theta, \widehat{\theta})/(\lambda - 1)$ exactly. For finite thresholds, separate calibration on the same sample produces the same regret sublevel for every λ. For full sublevels in the effective parameter space, the robust decision remains nominal.

Proof of Theorem 3. Conditional on the information used to fix the predictor and score, the calibration and test pairs (z, θ) are exchangeable. Applying the same measurable score rule at each pair's context preserves exchangeability of the $n + 1$ scores. Independent continuous tie breakers give the test score a uniform rank in $\{1, \ldots, n + 1\}$. If this rank is at most $k = \lceil (n + 1)(1 - \alpha) \rceil$, the test score does not exceed the calibration threshold, giving coverage probability at least $k/(n + 1)$. If $k = n + 1$, the threshold is infinite and coverage is automatic. Averaging over the fitting data proves the marginal statement.

To obtain the simultaneous bounds, consider the coverage event with a finite threshold. The test parameter is then admissible, and each nonnegative term in (9) is bounded by the score, giving all bounds in (11) on the same event. When the threshold is finite almost surely, this event has probability at least $1 - \alpha$. □

Optimizer localization. Strong convexity and $\widehat{\theta} \in \partial\Psi(\widehat{x})$ give $\Delta_{\widehat{\theta}}(x) \geq \mu \|x - \widehat{x}\|^2/2$. Combining this with $\lambda \Delta_{\widehat{\theta}}(x_\theta) \leq S_\lambda \leq \eta$ proves the stated distance bound. More generally, if $\Delta_{\widehat{\theta}}(x) \geq \phi(\|x - \widehat{x}\|)$ for a nondecreasing growth function ϕ, the same argument gives $\|x_\theta - \widehat{x}\| \leq \phi^\dagger(\eta/\lambda)$, where $\phi^\dagger(t) = \sup\{r \geq 0 : \phi(r) \leq t\}$.

Proposition 5 (Comparison with exact-regret calibration)**.** *Suppose $A \in C^2(\mathbb{R}^d)$ and $mI \preceq \nabla^2 A(u) \preceq LI$ for all u, with $0 < m \leq L$. Write $\kappa_A = L/m$ and $K_\lambda = 1 + \kappa_A\{\lambda^2/(\lambda - 1) - 1\}$. Then*

$$\mathrm{Reg}(\widehat{x}; \theta) \leq S_\lambda(\theta, \widehat{\theta}) \leq K_\lambda \, \mathrm{Reg}(\widehat{x}; \theta). \tag{26}$$

At a fixed context, the population threshold $\eta_\lambda^{\mathrm{FY}} = q_{1-\alpha}(S_\lambda(\theta, \widehat{\theta}))$ therefore satisfies

$$\eta^* \leq \eta_\lambda^{\mathrm{FY}} \leq K_\lambda \eta^*. \tag{27}$$

The same comparison holds for paired calibration order statistics whenever the curvature bounds hold uniformly across contexts.

Proof. The Hessian bounds give $m\|u - v\|_2^2/2 \leq D_A(u, v) \leq L\|u - v\|_2^2/2$. Hence

$$D_A(\widehat{\theta}, \theta) \leq \kappa_A D_A(\theta, \widehat{\theta}), \qquad D_A(r_\lambda, \widehat{\theta}) \leq \frac{\kappa_A}{(\lambda - 1)^2} D_A(\theta, \widehat{\theta}).$$

Substitution into the three-gap identity gives the upper bound, since $\lambda+1/(\lambda-1) = \lambda^2/(\lambda-1)-1$. The lower bound is immediate from that identity. Monotonicity of quantiles and order statistics proves the threshold comparison. □

A larger score threshold need not increase the nominal regret certificate, whose change is quantified in Appendix A.5.

A.3 Convexity and the Robust Policy Guarantee

Proof of Proposition 2. For each θ, the loss $L_\Psi(\theta; x)$ is convex in x; its supremum is therefore convex. For $\bar{\theta} = \sum_j p_j \theta^j \in \operatorname{conv}(U)$, convexity of A gives

$$L_\Psi(\bar{\theta}; x) \le \sum_j p_j L_\Psi(\theta^j; x) \le \rho_U(x).$$

Taking the supremum over $\bar{\theta}$ and using $U \subseteq \operatorname{conv}(U)$ proves equality of the worst-case losses and their infima. □

Proof of Corollary 2. On the conformal coverage event, realized regret is bounded by worst-case regret over the calibrated set. For an exact robust optimizer, (14) bounds this by $\eta_{\alpha,\lambda}$. Theorem 3 gives the stated marginal probability.

Feasible algorithmic policies retain this guarantee if their worst-case regret has a certified global upper bound at most the threshold almost surely. A certified optimality gap at most ε almost surely instead gives threshold $\eta_{\alpha,\lambda} + \varepsilon$ by (14). □

Polyhedral FY sets. Let $\mathcal{X} = \operatorname{conv} V$ be a bounded polytope and write $h = \widehat{\theta}$. Among true-optimal vertices, we select one maximizing $h^\top v$, resolving remaining ties by a fixed vertex order. For a true-optimal vertex v, positive homogeneity of A gives

$$S_{\lambda,v}(\theta, h) = A(\lambda h - \theta) - (\lambda h - \theta)^\top v = A(\theta) + A(\lambda h - \theta) - \lambda h^\top v.$$

The prediction-preferred rule therefore minimizes the score among true optimizers; the final vertex order does not change its value. In any closed polyhedral parameter domain Θ, including the full space,

$$\mathcal{U}_{\lambda,\eta}(h) = \bigcup_{v \in V} \left\{ \theta \in \Theta : (u-v)^\top \theta \le 0, \quad (u-v)^\top (\lambda h - \theta) \le \eta \quad \forall u \in V \right\}. \tag{28}$$

Each branch in (28) identifies a true optimizer with score at most η, so the selected optimizer has no larger score. Conversely, every parameter in the FY set belongs to the branch of its selected vertex. The set is therefore a finite union of closed polyhedra and is itself closed.

Define the width $w_{\mathcal{X}}(a) = A(a) + A(-a)$. Subadditivity of A and $h^\top v \le A(h)$ imply

$$S_\lambda(\theta, h) \ge w_{\mathcal{X}}(\theta) - \lambda w_{\mathcal{X}}(h).$$

The width is positive on every nonzero direction of $E = \operatorname{span}(\mathcal{X} - \mathcal{X})$. Compactness of its unit sphere therefore gives $w_{\mathcal{X}}(a) \ge c_{\mathcal{X}} \|a\|_2$ on E for some $c_{\mathcal{X}} > 0$. After projection onto the effective space, the branches remain closed polyhedra, and the width bound makes their union bounded for finite η. Thus the FY set is compact in effective coordinates.

For the portfolio QP, compactness of the feasible set and strong convexity of the penalty give a unique continuous oracle and hence a continuous score. The effective-space bound in Appendix B.3 establishes compactness of the FY set. The general results in Section 3 retain their stated compactness and continuity assumptions.

A.4 Robust Optimality and Screening

For Legendre Ψ and compact $U \subset \operatorname{int} \operatorname{dom} A$, the FY loss equals the Bregman divergence generated by Ψ from x to $x_\theta = \nabla A(\theta)$. Problem (13) is then a right-Bregman Chebyshev-center problem for the oracle decisions. In this setting, Bauschke et al. (2010, Theorems 4.3–4.4) give the interior subdifferential and center conditions in Theorem 4. We use convex subdifferentials directly below to accommodate nonsmooth Ψ and indicator constraints.

Proof of Theorem 4. Write $\rho_U = \Psi + h_U$, where $h_U(x) = \max_{\theta \in U}\{A(\theta) - \langle \theta, x\rangle\}$. Compactness of U and continuity of A make h_U finite and continuous, with a nonempty compact active set. The maximum rule gives $\partial h_U(x) = -\operatorname{conv} W_U(x)$. The convex sum rule and Fermat's rule then yield

$$\partial \rho_U(x) = \partial \Psi(x) - \operatorname{conv} W_U(x), \qquad 0 \in \partial \rho_U(x) \iff \partial \Psi(x) \cap \operatorname{conv} W_U(x) \neq \varnothing.$$

If Ψ is differentiable at $\widehat{x}$, nominal optimality gives $\partial \Psi(\widehat{x}) = \{\widehat{\theta}\}$. □

Proof of Corollary 3. The intersection in the optimality condition provides a certificate, and Carathéodory's theorem bounds the number of active parameters needed to represent it. To relate this condition to the movement residual, set $K = \partial \rho_U(\widehat{x})$. Since $\partial \Psi(\widehat{x})$ is nonempty and closed and $\operatorname{conv} W_U(\widehat{x})$ is compact, K is closed and $\mathfrak{M}(\widehat{x}; U) = \min_{s \in K} \|s\|_*$. This distance is zero exactly when $0 \in K$.

Under local finiteness of Ψ, K is also compact, and the convex directional derivative equals its support function. Convex minimax gives

$$\min_{\|v\| \leq 1} \rho'_U(\widehat{x}; v) = \min_{\|v\| \leq 1} \max_{s \in K} \langle s, v\rangle = \max_{s \in K}(-\|s\|_*) = -\mathfrak{M}(\widehat{x}; U).$$

□

Proof of Corollary 4. For $\theta \in U$, either $x_\theta = \widehat{x}$ and nominal regret is zero, or the three-gap identity gives $\operatorname{Reg}(\widehat{x}; \theta) \leq \eta - \lambda\Gamma$. Thus $\rho_U(\widehat{x}) \leq (\eta - \lambda\Gamma)_+$, which is zero when $\eta \leq \lambda\Gamma$. □

Proposition 6 (Local curvature and decision movement)**.** *In one dimension, suppose A is C^3 near $\widehat{\theta}$, with $h = A''(\widehat{\theta}) > 0$ and $\tau = A'''(\widehat{\theta})$. For fixed $\lambda > 1$, assume the complete sublevel $U_\eta = \{\theta : S_\lambda(\theta, \widehat{\theta}) \leq \eta\}$ is a compact interval shrinking to $\widehat{\theta}$ as $\eta \downarrow 0$. Its unique robust center satisfies*

$$x^{\mathrm{rob}} - \widehat{x} = \frac{2 - \lambda}{3\lambda^2} \frac{\tau}{h} \eta + o(\eta). \tag{29}$$

Proof. Translate $\widehat{\theta}$ to zero and write $U_\eta = [t_-, t_+]$. Since the loss is convex in the parameter, its maximum occurs at an endpoint. We therefore consider the decision that equalizes the two endpoint losses:

$$x_0 = \frac{A(t_+) - A(t_-)}{t_+ - t_-}.$$

The mean value theorem gives $m \in (t_-, t_+)$ with $A'(m) = x_0$. Thus $m \in \partial \Psi(x_0) \cap \operatorname{conv}\{t_-, t_+\}$, so the active-hull condition establishes optimality. The subgradient bound is strict when using t_+ for $x < x_0$ and t_- for $x > x_0$, proving uniqueness and $x^{\mathrm{rob}} = x_0$.

Expanding $S_\lambda(t, 0)$ gives

$$S_\lambda(t, 0) = a_2 t^2 + a_3 t^3 + o(t^3), \qquad a_2 = \frac{\lambda^2 h}{2(\lambda - 1)}, \quad a_3 = \frac{\lambda^2(2\lambda - 3)\tau}{6(\lambda - 1)^2}.$$

With $r = \sqrt{\eta / a_2}$, the two boundary roots satisfy $t_\pm = \pm r - a_3 r^2/(2a_2) + o(r^2)$. Substituting into the secant slope yields

$$x^{\mathrm{rob}} - A'(0) = \left(\frac{\tau}{6} - \frac{h a_3}{2a_2}\right) r^2 + o(r^2),$$

which simplifies to (29). □

When $\tau \neq 0$, curvature asymmetry produces a first-order displacement for $\lambda \neq 2$; symmetric reflection cancels this leading term. Nonquadraticity alone, however, does not force movement. For example, take $\widehat{\theta} = 0$, $\lambda = 2$, and $A(\theta) = \|\theta\|_2^2/2 + \theta_1^4/4$. Then $S_2(\theta, 0) = 2(\|\theta\|_2^2 + \theta_1^4)$. Both its compact sublevels and the nominal regret $A(\theta)$ are symmetric, so the active hull contains zero and the robust center remains nominal.

A.5 The Value of Robustification

Proof of Proposition 3. The gain $\mathcal{G}_{\mathrm{rob}}(U)$ is zero exactly when $\widehat{x}$ minimizes ρ_U. By the preceding optimality condition, this is equivalent to $\mathfrak{M}(\widehat{x}; U) = 0$. If Ψ is μ-strongly convex, so is ρ_U; its quadratic lower bound implies coercivity and a unique minimizer x^{rob}. Applying strong convexity at this minimizer gives

$$\mathcal{G}_{\mathrm{rob}}(U) \geq \frac{\mu}{2}\|\widehat{x} - x^{\mathrm{rob}}\|^2.$$

For a minimum-norm $s \in \partial\rho_U(\widehat{x})$, strong convexity at $\widehat{x}$, followed by the dual-norm inequality, gives

$$\mathcal{G}_{\mathrm{rob}}(U) \leq \langle s, \widehat{x} - x^{\mathrm{rob}}\rangle - \frac{\mu}{2}\|\widehat{x} - x^{\mathrm{rob}}\|^2 \leq \max_{t\geq 0}\left\{\mathfrak{M}t - \frac{\mu}{2}t^2\right\} = \frac{\mathfrak{M}^2}{2\mu}.$$

□

Population comparison. Fix a set U with conditional coverage at least $1-\alpha$. Every decision x satisfies $\mathbb{P}\{L_\Psi(\theta; x) \leq \rho_U(x)\} \geq 1-\alpha$, and hence $q_{1-\alpha}(L_\Psi(\theta; x)) \leq \rho_U(x)$. Taking infima over decisions gives $V_\alpha \leq R^*(U)$. Adding and subtracting $R^*(U)$ then yields (21).

For the population-calibrated set $U_\lambda = \{\theta : S_\lambda(\theta, \widehat{\theta}) \leq \eta_\lambda^{\mathrm{FY}}\}$, define its nominal certificate cost as $C_\lambda = \rho_{U_\lambda}(\widehat{x}) - \eta^* \geq 0$. Then

$$R^*(U_\lambda) - \eta^* = C_\lambda - \mathcal{G}_{\mathrm{rob}}(U_\lambda). \tag{30}$$

Under Proposition 5, $C_\lambda \leq (K_\lambda - 1)\eta^*$. Thus a better certificate than direct regret calibration requires the optimization gain to exceed this cost.

Proposition 7 (Better calibration without decision movement)**.** *Let $\Psi(x) = x^\top H^{-1}x/2$ on $\mathbb{R}^d$, $\widehat{\theta} = 0$, and $\theta \sim \mathcal{N}(0, \sigma^2 I_d)$, with $d \geq 2$ and $\sigma > 0$. Let B_2 be the smallest Euclidean ball centered at zero with coverage $1-\alpha$, and let U_λ be the population-calibrated FY set. For $H = \mathrm{diag}(1, \varepsilon, \ldots, \varepsilon)$, $0 < \varepsilon \leq 1$, both sets have robust center $\widehat{x} = 0$, but*

$$\lim_{\varepsilon\downarrow 0}\frac{R^*(B_2)}{R^*(U_\lambda)} = \frac{\chi^2_{d,1-\alpha}}{\chi^2_{1,1-\alpha}}, \tag{31}$$

where $\chi^2_{k,p}$ denotes the p-quantile of a chi-squared distribution with k degrees of freedom. For fixed $\alpha \in (0,1)$, this limiting ratio grows linearly with dimension.

Proof. Here $A(\theta) = \theta^\top H\theta/2$ and $S_\lambda(\theta, 0) = \lambda^2 A(\theta)/(\lambda - 1)$. Thus $U_\lambda = \{\theta : A(\theta) \leq \eta^*\}$ and $R^*(U_\lambda) = \eta^*$. Symmetry of either set and strict convexity of Ψ make zero its unique robust center. The squared radius of B_2 is $\sigma^2\chi^2_{d,1-\alpha}$, so

$$R^*(B_2) = \frac{\sigma^2}{2}\chi^2_{d,1-\alpha}, \qquad \eta^* = \frac{\sigma^2}{2}\,q_{1-\alpha}\left(Z_1^2 + \varepsilon\sum_{j=2}^{d} Z_j^2\right),$$

where the Z_j are independent standard normals. For each fixed d, the latter quantile converges to $\chi^2_{1,1-\alpha}$ as $\varepsilon \downarrow 0$. Finally, $\chi^2_{d,1-\alpha}/d \to 1$ as $d \to \infty$. □

The residual ball treats all parameter directions equally, whereas regret calibration accounts for their different decision costs. The example thus isolates the benefit of calibration geometry when both robust decisions remain nominal.

A.6 Adversarial Barycenters and Scenario Generation

Proof of Theorem 5. For $\nu \in \mathcal{P}(U)$, write $m_\nu = \mathbb{E}_\nu[\theta]$ and $J(\nu) = \mathbb{E}_\nu[A(\theta)] - A(m_\nu)$. Expanding the FY loss gives

$$\rho_U(x) - J(\nu) = L_\Psi(m_\nu; x) + \mathbb{E}_\nu[\rho_U(x) - L_\Psi(\theta; x)] \geq 0. \tag{32}$$

This establishes weak duality. To show attainment, take a robust optimizer x^*. Theorem 4 provides $m^* \in \partial\Psi(x^*) \cap \mathrm{conv}\, W_U(x^*)$. Represent m^* as a convex combination of at most $\dim(\mathrm{aff}\, W_U(x^*)) + 1$ active parameters, and let ν^* be the corresponding discrete distribution. Both terms in (32) vanish, giving $J(\nu^*) = R^*(U)$. Each support point has regret $R^*(U)$. Finally, $m^* \in \partial\Psi(x^*)$ implies $x^* \in \partial A(m^*)$, and differentiability gives $x^* = x_{m^*}$. □

For a finite scenario set, the barycenter representation becomes

$$R_J = \max_{p \in \Delta_J} \left\{ \sum_j p_j A(\theta^j) - A\left(\sum_j p_j \theta^j \right) \right\}, \qquad \Delta_J = \{p \geq 0 : \textstyle\sum_j p_j = 1\}, \tag{33}$$

provided the restricted minimum is attained. For Legendre Ψ and scenarios in $\operatorname{int} \operatorname{dom} A$, this is the smallest-enclosing-Bregman-ball dual of Nock & Nielsen (2005, Eqs. (6)–(7)), applied to oracle decisions x_{θ^j}: conjugacy gives $\nabla\Psi(x_{\theta^j}) = \theta^j$ and reduces their dual objective to (33). Any feasible weights give a lower bound on R_J.

Proof of Proposition 4. Restricting scenarios to $U_J \subseteq U$ gives the lower bound $R_J \leq R^*(U)$. Feasibility of the master decision x_J for the full problem gives the upper bound $R^*(U) \leq \rho_U(x_J)$. Subtracting these bounds proves the stopping rule. The same argument applies to approximate solves:

$$\underline{R}_J \leq R^*(U) \leq \rho_U(x_J) \leq \overline{P}_J, \tag{34}$$

where $\underline{R}_J$ is a master lower bound and $\overline{P}_J$ is a global pricing upper bound for that decision. □

A feasible adversarial scenario bounds the pricing maximum from below and cannot replace $\overline{P}_J$ in (34).

B Experimental Details

We first describe the common protocol and baselines, then present the data, solution methods, and additional results for the experiments in Section 4.

B.1 Common Protocol and Baselines

Data splits and calibration. For the synthetic simplex and portfolio experiments, we generated 4,000 training, 1,000 tuning, and 2,000 calibration observations using master seed 20260903 and independent random streams for each split. Their test sets contain 5,000 and 1,000 contexts, respectively. The solar-grid experiment uses the chronological splits in Appendix B.4.

Predictors and residual geometry are fitted before calibration, and all methods share the realized parameters, constraints, and hindsight optimum. Each method calibrates its own score using (2). With 2,000 calibration observations, the selected ranks are 1,801 at $\alpha = 0.1$ and 1,901 at $\alpha = 0.05$.

Baseline construction. PTO maximizes $J(x; \widehat{\theta})$, FY minimizes worst-case regret over its reflected-score set, and the residual baselines maximize $\min_{\theta \in U} J(x; \theta)$ using the same predictor. Their scores are

$$s_2 = \|\theta - \widehat{\theta}\|_2, \qquad s_{\text{box}} = \max_j \frac{|\theta_j - \widehat{\theta}_j|}{\widehat{s}_j(z)}, \qquad s_{\text{ell}} = \|\widehat{\Sigma}^{-1/2}(\theta - \widehat{\theta})\|_2.$$

Each box scale estimates the 0.8 quantile of absolute training residuals using 40 depth-two boosting trees, with a floor of 10^{-6}. The matrix $\widehat{\Sigma}$ is the Ledoit–Wolf covariance of training residuals. KNN RO uses the mean and Ledoit–Wolf covariance of the outcomes at the 100 nearest training contexts in standardized context space, or all available observations if fewer than 100. In the controlled QP, all these quantities are fitted on historical training data.

For an ellipsoid $U = \{m + e : e^\top \Sigma^{-1} e \leq q^2\}$, the support function gives the convex counterpart

$$\max_{x \in \mathcal{X}} \left\{ m^\top x - \phi(x) - q\sqrt{x^\top \Sigma x} \right\}, \tag{35}$$

where $\phi = 0$ for the linear programs and $\phi(x) = \frac{1}{2} x^\top H x$ for the QP. The residual ℓ_2 case uses $m = \widehat{\theta}$ and $\Sigma = I$; the box replaces the norm penalty by $q \sum_j \widehat{s}_j(z)|x_j|$.

Discussion of other decision-aware methods. Our experiments examine calibration and robust refinement around fixed point predictions. Conformal Risk Control (Angelopoulos et al., 2024) controls expected bounded monotone losses, including the miscoverage indicator. It calibrates a prescribed family of outputs, leaving the choice of score and downstream policy to the application. Applying it directly to decision regret would require a suitable bounded, monotone loss family; the realized regret of a reoptimized decision need not be monotone as its uncertainty set grows.

Utility-Directed Conformal Prediction (Cortes-Gomez et al., 2025) combines class probabilities with set costs to construct low-cost prediction sets over finite labels. Risk-Averse Calibration (Kiyani et al., 2025) uses label probabilities to construct utility-superlevel sets and max–min actions targeting a utility quantile. Both allow post-processing of a fixed predictor. Adapting their finite-label algorithms to our experiments would require a probabilistic model for the outcomes and a continuous-domain implementation or an explicit discretization of the set-construction problems. These are additional modeling choices beyond the point forecasts used in our comparison.

The end-to-end approaches of Chenreddy & Delage (2024) and Yeh et al. (2025) learn uncertainty representations through downstream optimization losses, with the latter also differentiating through conformal calibration. These methods apply to continuous optimization, including portfolio allocation. Evaluating them would require a separate training and tuning protocol and would compare the combined effects of decision-aware learning and subsequent calibration and optimization. Our predictor-based comparisons keep training fixed to examine calibration and refinement after prediction.

Wang & Dobriban (2026) propose Risk-Optimal Conformal Prediction using probabilistic predictions and a worst-case expected-loss criterion that accounts for outcomes outside the prediction set. Our criterion minimizes worst-case regret within the calibrated set. On the unbounded parameter domains in our synthetic experiments, worst-case out-of-set regret can be infinite, so a comparison under their criterion would also require specifying an outcome-domain or loss restriction. This would change both the information supplied to the method and the risk formulation being evaluated.

Evaluation. We evaluate regret as $A(\theta) - J(x;\theta)$ using the same realized θ for every method. A reduction in mean regret therefore equals an increase of the same size in mean realized objective value. Q95 uses linear interpolation, and CVaR_{95} averages the largest 5% of empirical regret mass. Movement is measured by $\|x - \widehat{x}\|_2$, and a decision is counted as changed when its infinity-norm difference from PTO exceeds 10^{-6}. Empirical coverage is evaluated using each method's own parameter or score event.

The synthetic comparisons use a single training seed. We construct pointwise 95% intervals from 2,000 paired bootstrap resamples of the common test observations, conditional on the fitted predictor and calibration sample.

B.2 Contextual Simplex Allocation

Data and prediction. The decision is $x = (p, 1-p)$ with $p \in [0,1]$, and the objective is $\theta^\top x$. Contexts satisfy $z \sim N((1,1)^\top, 2.25 I_2)$, and returns are

$$\theta_1 = -5z_1 - 2z_2^2 - \epsilon_1, \qquad \theta_2 = -3z_1^2 - z_2 - \epsilon_2, \qquad \epsilon \sim N(0, I_2).$$

We fit a ridge predictor with quadratic features and regularization parameter 10, selected by five-fold cross-validation on the training set using the one-standard-error rule. This predictor is held fixed across comparisons. Figure 3 uses $\alpha = 0.05$, while Table 4 also reports the default $\alpha = 0.1$ and the other coverage levels.

Exact refinement. In the two-asset LP, regret depends only on $d = \theta_1 - \theta_2$. Relabel the nominally preferred asset as asset 1, so that $h = |\widehat{\theta}_1 - \widehat{\theta}_2| > 0$, and let p be its allocation. At $\lambda = 2$, when $2h < \eta$, the FY set projects to $[2h - \eta, 2h + \eta]$. Hence

$$\rho_U(p) = \max\{p(\eta - 2h), (1-p)(\eta + 2h)\}, \qquad p_{\mathrm{FY}} = \frac{1}{2} + \frac{h}{\eta}.$$

When $2h \geq \eta$, the set lies within the nominal decision region and $p_{\mathrm{FY}} = 1$. Thus the normalized return gap in Figure 3 is $2h/\eta$, and the allocation transferred away from PTO is $\frac{1}{2}(1 - 2h/\eta)_+$.

For comparison, direct nominal-regret calibration with threshold $q > 0$ gives the projected set $[-q, \infty)$. Any $p < 1$ then has infinite worst-case regret, whereas PTO has value q. PTO also remains minimax optimal at $q = 0$. These explicit solutions give the two curves in panel (a).

At $\alpha = 0.05$, direct regret and FY calibration give $q = 0.193961$ and $\eta = 1.620846$. FY changes the decision in 423 of the 5,000 contexts; in the remaining 4,577, both PTO and FY have zero worst-case regret over the FY set. The marked context has $\rho_U(\widehat{x}) = 1.278587$, $R^*(U) = 0.774287$, and $\mathcal{G}_{\mathrm{rob}}(U) = 0.504300$. The average gain is 0.023001 across all contexts and 0.271882 among those with decision movement.

Additional results and sensitivity. Table 3 gives the full comparison at $\alpha = 0.1$. Covariance shrinkage produces a scalar multiple of the identity, so the residual ellipsoid and residual ℓ_2 methods yield the same objective-robust policy. The regret-robust ellipsoid uses this same set with the minimax-regret objective, while the global ellipsoid is fitted to all training outcomes.

Table 4 varies one parameter at a time, holding the predictor and observations fixed and recalibrating each score. Higher target coverage or smaller λ can increase decision movement. When calibration gives a zero threshold because the score has mass at zero, FY retains PTO. The results also show that mean and tail regret need not improve together.

Table 3: Contextual simplex results at $\alpha = 0.1$. Movement is mean Euclidean distance from PTO; coverage uses each method's own set.

Method	Mean	Q95	CVaR95	Movement	Coverage
PTO	0.0502	0.1234	0.9945	0.0000	–
FY	0.0499	0.2174	0.9621	0.0077	89.8%
Residual ℓ_2	0.1149	0.7273	1.0428	0.1009	91.7%
Residual box	0.0502	0.1514	0.9930	0.0045	90.9%
Residual ellipsoid	0.1149	0.7273	1.0428	0.1009	91.7%
KNN RO	0.2410	1.3794	2.7948	0.1244	89.7%
Residual ellipsoid (regret)	0.1539	0.7888	1.0700	0.1259	91.7%
Global ellipsoid	4.2643	13.5813	18.6646	0.7331	89.8%

Table 4: Simplex FY sensitivity. The α grid fixes $\lambda = 2$, and the λ grid fixes $\alpha = 0.1$. Each configuration uses its own calibrated threshold.

α	λ	η	Coverage	Mean	CVaR95	Moved	Mean $\mathcal{G}_{\mathrm{rob}}$
0.05	2	1.6208	94.50%	0.0568	0.9334	8.46%	0.0230
0.10	2	0.3912	89.82%	0.0499	0.9621	2.06%	0.0014
0.20	2	0.0000	87.56%	0.0502	0.9945	0.00%	0.0000
0.30	2	0.0000	87.56%	0.0502	0.9945	0.00%	0.0000
0.40	2	0.0000	87.56%	0.0502	0.9945	0.00%	0.0000
0.10	1.25	1.2419	89.70%	0.0602	0.9355	10.36%	0.0215
0.10	1.5	0.8732	89.62%	0.0530	0.9398	5.98%	0.0090
0.10	3	0.0000	91.04%	0.0502	0.9945	0.00%	0.0000
0.10	5	0.0000	92.70%	0.0502	0.9945	0.00%	0.0000

B.3 Constrained Portfolio Optimization

Data and prediction. The main setting has $m = 20$, $H = I_{20}$, and cap $u = 0.2$:

$$J(x;\theta) = \theta^\top x - \tfrac{1}{2}\|x\|_2^2, \qquad \mathcal{X} = \{x \in \mathbb{R}^{20} : \mathbf{1}^\top x = 1,\ 0 \le x_j \le 0.2\}.$$

Calibration and test observations follow $\theta = c + Bz + \sigma\xi$, where $z \sim N(0, I_5)$, $\xi \sim N(0, I_{20})$, and $\sigma = 0.03$. The intercept is

$$c = (0.35, 0.35, 0.18, 0.16, 0.14, 0.12, -0.02, -0.30, \ldots, -0.30)^\top.$$

The columns of B are 0.003 times

$$\frac{e_3 - e_4}{\sqrt{2}},\ \frac{e_4 - e_5}{\sqrt{2}},\ \frac{e_5 - e_6}{\sqrt{2}},\ \frac{e_3 + e_4 + e_5 + e_6 - 4e_7}{\sqrt{20}},\ \frac{e_1 - e_2}{\sqrt{2}}.$$

Historical training and tuning returns include the additional shift $\kappa\sigma v$, with $v = (e_7 - e_3)/\sqrt{2}$ and $\kappa \in \{0, 0.5, 1, 2, 4\}$. Calibration and test observations follow the current distribution, so the shift affects predictions through the historical training data. We select the predictor by historical tuning MSE from ridge models with penalties $\{1, 10, 100, 1000\}$ and a neural network with two hidden layers. Ridge with penalty 1000 is selected in every shift regime; the main table uses $\kappa = 2$.

Master and pricing formulations. For general $H \succ 0$, write $\phi(x) = \frac{1}{2}x^\top Hx$ and use the capped simplex $\mathcal{X} = \{x : \mathbf{1}^\top x = 1,\ 0 \le x \le u\mathbf{1}\}$. The oracle and regret are

$$A(\theta) = \max_{y \in \mathcal{X}} \{\theta^\top y - \phi(y)\}, \qquad L_\Psi(\theta; x) = A(\theta) + \phi(x) - \theta^\top x.$$

Since $A(\theta + c\mathbf{1}) = A(\theta) + c$, common-return shifts leave both regret and the FY score unchanged. We center $\widehat{\theta}$ and work on $E = \{\theta : \mathbf{1}^\top \theta = 0\}$ for FY pricing.

For scenarios $\theta^1, \dots, \theta^J$ satisfying the FY score constraint, separating the common quadratic term gives the master QP

$$R_J = \min_{x \in \mathcal{X},\ t \in \mathbb{R}} \left\{ \tfrac{1}{2}x^\top Hx + t : t + (\theta^j)^\top x \ge A(\theta^j),\ j = 1, \dots, J \right\}. \tag{36}$$

The scenario values are cached, and each new scenario adds a linear constraint. The master multipliers supply the weights in (33).

Let $y = x_\theta$. With bound multipliers $a, b \ge 0$ and budget multiplier $\gamma \in \mathbb{R}$, the oracle KKT conditions are

$$y \in \mathcal{X}, \qquad \theta = Hy - a + b + \gamma\mathbf{1}, \qquad a_j y_j = b_j(u - y_j) = 0. \tag{37}$$

Multiplying stationarity by y gives $A(\theta) = \frac{1}{2}y^\top Hy + u\mathbf{1}^\top b + \gamma$. For the reflected oracle at $r = r_\lambda$, introduce dual variables $\widetilde{a}, \widetilde{b} \ge 0$, $\widetilde{\gamma} \in \mathbb{R}$, and a free vector v, with

$$Hv = r + \widetilde{a} - \widetilde{b} - \widetilde{\gamma}\mathbf{1}, \qquad D_r = \tfrac{1}{2}v^\top Hv + u\mathbf{1}^\top \widetilde{b} + \widetilde{\gamma}.$$

QP duality gives $A(r) = \min D_r$ over these variables. Substituting the true-oracle identity into the score yields the equivalent pricing model

$$\begin{aligned} \max \quad & \tfrac{1}{2}y^\top Hy + u\mathbf{1}^\top b + \gamma + \phi(x) - \theta^\top x \\ \text{s.t.} \quad & (37), \quad \theta \in E, \quad r = (\lambda\widehat{\theta} - \theta)/(\lambda - 1), \\ & Hv = r + \widetilde{a} - \widetilde{b} - \widetilde{\gamma}\mathbf{1}, \quad a, b, \widetilde{a}, \widetilde{b} \ge 0, \\ & (\lambda - 1)D_r + \tfrac{\lambda+1}{2}y^\top Hy - \lambda\widehat{\theta}^\top y + u\mathbf{1}^\top b + \gamma \le \eta. \end{aligned} \tag{38}$$

Every feasible point of (38) satisfies $D_r \ge A(r)$ and hence the original score constraint. Conversely, for every parameter satisfying that constraint, an optimal reflected dual solution gives a feasible point of the reformulation. We use this nonconvex formulation to obtain global pricing bounds for the certification study below, and the following boundary search to compute the reported policies.

Boundary search. On an oracle region C_I, the KKT system gives $x_\theta = F_I\theta + f_I$ and $A(\theta) = \frac{1}{2}\theta^\top F_I\theta + f_I^\top\theta + c_I$, with $F_I \succeq 0$. For a true/reflected region pair (I, K), write $y = x_\theta$ and $y_r = x_{r_\lambda}$. Differentiation gives

$$\begin{aligned} g &:= \nabla_\theta L_\Psi(\theta; x) = y - x, \\ s &:= \nabla_\theta S_\lambda = y - y_r + \lambda F_I(\theta - \widehat{\theta}), \\ Q_{IK} &:= \nabla^2_\theta S_\lambda = (\lambda + 1)F_I + F_K/(\lambda - 1). \end{aligned} \tag{39}$$

These derivatives guide the search within each region pair. When a proposal crosses a region boundary, we evaluate its score and regret using the oracles.

For a nonzero direction $d \in E$, define the boundary map $\mathcal{B}(d) = \widehat{\theta} + \tau(d)d$, where $\tau(d) = \sup\{t \ge 0 : S_\lambda(\widehat{\theta} + td, \widehat{\theta}) \le \eta\}$. Convexity of A and $F_I \succeq 0$ make the score nondecreasing along each ray. We bracket the boundary and solve the quadratic equation within each cell, using bisection when a proposed root leaves the bracket.

When Q_{IK} is positive definite on E, completing the square in the current score piece gives an ellipsoid. Maximizing the current quadratic regret over that ellipsoid, without the region inequalities, reduces after rescaling to

$$\max_{\|z\|_2 \le 1} \left\{ \tfrac{1}{2} z^\top M z + b^\top z \right\}, \qquad (\tau I - M) z = b, \quad \tau \ge \lambda_{\max}(M), \quad \|z\|_2 = 1.$$

Here $M \succeq 0$ and b are the transformed regret coefficients. An eigendecomposition and a scalar search for the multiplier give a proposal ϑ. We map it to $\mathcal{B}(\vartheta - \widehat{\theta})$ and evaluate it with the full oracles. If the proposal does not improve regret, we try the linear support point of the same ellipsoid, followed by a boundary ascent step. For $\delta = \theta - \widehat{\theta}$, differentiation through the radial map gives

$$g_{\mathrm{rad}} = g - \frac{g^\top \delta}{s^\top \delta} s.$$

We precondition it with H, project onto E, and backtrack after each radial retraction. This last step also applies when the local score curvature is singular, provided $s^\top \delta > 0$.

The search starts from $\widehat{\theta}$ and uses up to four high-regret scenarios from distinct directions as warm starts. At each expansion, it processes the queued point with the largest regret; every third expansion also adds a new boundary direction. Directions that change the nominal active set take priority, followed by training residuals, coordinate directions, and random directions. Candidates are accepted only if they satisfy the full score constraint and improve the pricing objective. Each call allows at most 32 expansions and a search depth of eight.

Algorithm 1 Scenario generation for the portfolio QP

Input: $\widehat{\theta}, H, u, \lambda, \eta$; tolerance ε; time budget T; initial search directions.
Output: A feasible decision x, master lower bound, and any available global pricing bound.

1. Compute $\widehat{x} = x_{\widehat{\theta}}$. If $\eta = 0$, return $\widehat{x}$; otherwise initialize the master and scenario pool with $\widehat{\theta}$.
2. Solve (36) for x and retain its lower bound.
3. Search for adverse parameters by the boundary search above. Validate their FY scores and add them to the scenario pool.
4. If a global pricing upper bound for x is available and its gap to the master lower bound is at most ε, return x with that certificate.
5. Add up to six pool scenarios with $L_\Psi(\theta; x) > R_J + \varepsilon$ to the master. Repeat from step 2 until the time budget is exhausted or four consecutive searches find no violation after at least $2m$ ray probes.
6. Use the remaining budget to resolve the master against known pool violations. Return the last feasible master decision.

The controlled QP uses $T = 1$ second, $\varepsilon = 10^{-6}$, and pricing calls of at most 0.15 seconds. Gurobi solves the convex QPs; the final master solve against violated pool scenarios uses tolerance 10^{-7} and at most 0.1 seconds. Convergence on the scenario pool does not certify optimality over the full FY set. The evaluation includes all 5,000 decisions across shift levels; 181 runs reached the time limit.

Exact two-asset diagnostic. To assess solver accuracy against an exact solution, we construct a two-asset QP using the simplex data and predictor from Appendix B.2, with $u = 0.8$, $\alpha = 0.1$, and $\lambda = 2$. We scale the Ledoit–Wolf covariance of training returns so that $(1, -1)H(1, -1)^\top$ equals the median absolute predicted return difference on that sample. Calibrating the FY score on the 2,000 observations gives $\eta = 0.4648713$.

For $x = (p, 1 - p)$ and $d = \theta_1 - \theta_2$, write the quadratic penalty as $\psi(p) = \frac{1}{2} a p^2 + bp + c$, where $a = H_{11} - 2H_{12} + H_{22} > 0$, $b = H_{12} - H_{22}$, and $c = H_{22}/2$. Then

$$p_d = \mathrm{clip}_{[1-u,u]}\left(\frac{d - b}{a} \right), \qquad A_1(d) = d p_d - \psi(p_d).$$

The true and reflected clipping boundaries divide the scalar score into at most five quadratic pieces. Solving their sublevel inequalities gives the full interval $[d_-, d_+]$. Convexity of $A_1(d) - dp + \psi(p)$

in d reduces pricing to its two endpoints. For $d_- < d_+$, equalizing their losses yields

$$p^{\mathrm{rob}} = \frac{A_1(d_+) - A_1(d_-)}{d_+ - d_-}.$$

The resulting allocation lies in $[1-u, u]$ and minimizes the maximum endpoint loss. For a singleton interval, the robust decision is the oracle decision at that parameter.

We compare the one-second heuristic with the exact solution on 27 prescribed cases covering interior, saturated, and boundary-crossing configurations, and on 1,000 evenly spaced contexts from the 5,000-context simplex test set. Exact solutions are used only for evaluation. Table 6(a) reports the objective error $\varepsilon_{\mathrm{opt}} = \rho_U(x^{\mathrm{alg}}) - R^*(U)$ and decision error $\varepsilon_x = \|x^{\mathrm{alg}} - x^{\mathrm{exact}}\|_\infty$; both are below 10^{-7} in all 1,027 cases. Replacing the heuristic test decisions by exact solutions changes both mean regret and CVaR_{95} by less than 10^{-8}.

Table 5: Exact two-asset QP diagnostics. The gain threshold is 10^{-7}. Displacement is the absolute change in the first asset's allocation.

FY set	Contexts	Positive gain	Mean $\mathcal{G}_{\mathrm{rob}}$	Mean displacement
Within one region	43	0	$< 10^{-15}$	$< 10^{-15}$
Crossing regions	957	919	0.005847	0.006233

The exact solutions also illustrate when refinement occurs. Among the 1,000 test contexts, all 43 FY sets contained in one optimizer region retain PTO, while 919 of the 957 sets crossing region boundaries yield a gain above 10^{-7} (Table 5). Crossing a region boundary permits refinement, and the active-hull condition determines whether it occurs.

Global certification in twenty assets. We solve (38) by spatial branch-and-bound in Gurobi 13.0.1, encoding true-oracle complementarity with type-1 special ordered sets. Pricing ranges over the complete FY set. To obtain finite parameter bounds, set $r_0 = \min\{1/m, u - 1/m\}$ and $\Delta_H = \frac{1}{2}\lambda_{\max}(H)u$. Oracle optimality and the feasible ball around $\mathbf{1}/m$ in the budget hyperplane give

$$S_\lambda(\theta, \widehat{\theta}) \geq 2r_0\|\theta\|_2 - \lambda\left(\max_j \widehat{\theta}_j - \min_j \widehat{\theta}_j + \Delta_H\right), \qquad \theta \in E.$$

The score constraint $S_\lambda(\theta, \widehat{\theta}) \leq \eta$ therefore bounds every parameter coordinate, and stationarity bounds an optimal choice of dual multipliers. For numerical accuracy, we normalize the score constraint and objective by $\max\{\eta, 10^{-3}\}$ and use an absolute pricing tolerance of 10^{-7} in the original units.

We apply global pricing to the reported heuristic decisions in all 1,000 test contexts at $(\alpha, \kappa) = (0.1, 2)$. We also obtain reference decisions certified within 10^{-6} by adding pricing scenarios to the master while keeping predictions and calibration fixed. Each run starts from the lower bound and feasible scenarios obtained by the heuristic solver. The solver uses one thread, with time limits of 180 seconds per instance and 60 seconds per pricing call.

For the worst-case regret objective, let $L \leq R^*(U)$ be the master lower bound and $\overline{P}(x) \geq \rho_U(x)$ the global pricing upper bound. Then

$$0 \leq \rho_U(x) - R^*(U) \leq \varepsilon_{\mathrm{cert}}(x) := \overline{P}(x) - L.$$

For $L > 0$, the relative objective error $[\rho_U(x) - R^*(U)]/R^*(U)$ is bounded above by $\varepsilon_{\mathrm{cert}}(x)/L$. A lower bound on the absolute error is obtained by subtracting a reference decision's certified upper bound from the regret at a feasible adversarial parameter for x. Taking the positive part of this difference and dividing by that same positive upper bound gives a lower bound on the relative error.

All 1,000 reference decisions attain the prescribed tolerance, as do 919 of the reported heuristic decisions (Table 6, panel (b)). For the heuristic decisions, the relative objective error is certified below 0.021% in at least 95% of instances; for the worst instance, it lies between 0.97% and 0.99%. Independent evaluation using the scalar-projection oracle confirms feasibility and regret values for all 2,045 adversarial solutions within tolerance 10^{-7}. Replacing the heuristic decisions by the certified

references changes mean regret and CVaR_{95} by less than 10^{-7} in the original units, leaving both unchanged at the precision reported in the main table.

The largest global upper bound on worst-case regret among the 1,000 reported heuristic decisions is 0.005873906, below the calibration threshold 0.018474363 in the original units. Thus every tested decision satisfies the sufficient condition for retaining the original regret threshold, independently of its optimality error.

Pricing sensitivity across settings. To assess pricing stability across settings, we use the first eight test contexts for every combination of the five shift levels and four α values in Table 6(c). Holding each heuristic decision fixed, we run an additional one-second search from 64 newly sampled Gaussian directions and apply local boundary search to the best eight candidates. This search is initialized independently of the scenario-generation pool and uses no realized test outcomes. We define δ_{price} as the positive part of the increase in the largest feasible regret found, relative to scenario generation. The increase is at most 10^{-6} in 144 of the 160 evaluations; panel (c) summarizes the results by coverage level.

Table 6: Solver accuracy. Panel (a) compares the implemented solver with the exact two-asset solution. Panel (b) reports certified upper bounds $\varepsilon_{\mathrm{cert}}$ on the objective gap $\rho_U(x) - R^*(U)$ in twenty assets at $(\alpha, \kappa) = (0.1, 2)$. Panel (c) reports increases detected by additional heuristic pricing across settings. All objective values use the original, unscaled units.

(a) Exact two-asset comparison				
Instances	n	Mean $\varepsilon_{\mathrm{opt}}$	Max $\varepsilon_{\mathrm{opt}}$	Max ε_x
Prescribed cases	27	1.19×10^{-9}	8.13×10^{-9}	7.72×10^{-9}
Test contexts	1,000	1.19×10^{-9}	1.21×10^{-8}	1.11×10^{-8}
(b) Global twenty-asset certificates				
Policies	n	Count $\varepsilon_{\mathrm{cert}} \le 10^{-6}$	Mean $\varepsilon_{\mathrm{cert}}$	Max $\varepsilon_{\mathrm{cert}}$
Heuristic policy	1,000	919	8.48×10^{-7}	5.67×10^{-5}
Certified reference	1,000	1,000	4.73×10^{-7}	9.99×10^{-7}
(c) Additional twenty-asset pricing				
α	n	Count $\delta_{\mathrm{price}} > 10^{-6}$	Mean δ_{price}	Max δ_{price}
0.1	40	9	6.93×10^{-7}	4.81×10^{-6}
0.3	40	2	1.95×10^{-7}	2.52×10^{-6}
0.4	40	5	1.48×10^{-6}	4.70×10^{-5}
0.5	40	0	6.42×10^{-8}	5.06×10^{-7}

Additional results and sensitivity. Table 7 reports paired FY-minus-baseline differences at $(\alpha, \kappa) = (0.1, 2)$. The intervals lie below zero for all five baselines, for both mean regret and CVaR_{95}. Figure 4 and Table 8 examine sensitivity to the historical shift and coverage level. Within each shift level, we vary α while holding the predictor, residual geometry, calibration scores, and test contexts fixed.

At $\kappa = 2$, FY makes its largest average adjustment and improves both regret criteria, although its advantage over PTO decreases as α increases. At $\kappa = 4$, average movement falls from 0.006385 to 0.000279 relative to the preceding shift level, and regret approaches that of PTO at the default $\alpha = 0.1$. Setting $\alpha = 0.3$ or 0.4 produces larger improvements.

FY score coverage at $\alpha = 0.1$ is 90.1%, 90.2%, 90.3%, 89.3%, and 89.5% as the historical shift increases. The ranges at $\alpha = 0.3$, 0.4, and 0.5 are 66.1–69.9%, 56.1–59.9%, and 46.2–47.8%; at $\alpha = 0.1$, residual ℓ_2 and box coverage ranges are 89.9–90.5% and 90.3–91.9%, respectively. These empirical rates evaluate the marginal score event in Theorem 3. The historical shift affects the predictor, while calibration and test data continue to share the current distribution.

Table 7: Portfolio QP: paired FY-minus-baseline regret differences with pointwise 95% bootstrap intervals, multiplied by 10^3. Negative differences favor FY.

Reference	Mean difference [95% interval]	CVaR95 difference [95% interval]
PTO	-0.283 $[-0.295, -0.271]$	-0.403 $[-0.468, -0.333]$
Residual ℓ_2	-2.394 $[-2.463, -2.327]$	-3.400 $[-3.623, -3.148]$
Residual box	-0.348 $[-0.369, -0.328]$	-0.564 $[-0.699, -0.443]$
Residual ellipsoid	-2.385 $[-2.453, -2.318]$	-3.386 $[-3.608, -3.135]$
KNN RO	-2.451 $[-2.524, -2.380]$	-3.527 $[-3.776, -3.267]$

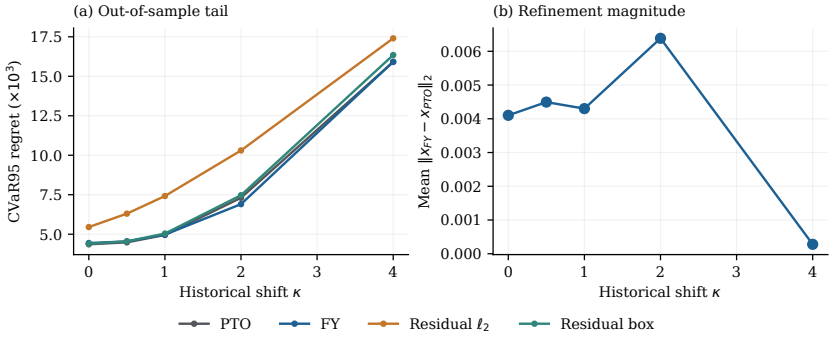


Figure 4: Controlled portfolio QP across historical shifts, with $\alpha = 0.1$ and $\lambda = 2$. (a) Test CVaR_{95}. (b) Mean Euclidean adjustment from PTO across the same 1,000 contexts per shift.

B.4 Real Solar-Grid Allocation

Data and decision problem. We use the public four-substation data from Zhou et al. (2026), comprising 168 monthly records from January 2011 to December 2024 and 182 solar installations. The release provides dates and counts without additional covariates. As in their stylized setting, upgrade costs and overload penalties equal one, with a budget of two. Let N_j denote the monthly count at station j, and let $x_j = 1 - a_j$ be its unupgraded fraction. The continuous relaxation is

$$\min_{x \in \mathcal{X}} q(N)^\top x, \qquad q_j(N) = \frac{1}{1 + \exp[-0.5(N_j - \tau_j)]}, \qquad \mathcal{X} = \{x \in [0,1]^4 : \mathbf{1}^\top x \geq 2\}. \quad (40)$$

The threshold τ_j is the mean count over the predictor's training period. With $\theta = -q(N)$ and $\Psi = I_{\mathcal{X}}$, we have $J(x;\theta) = \theta^\top x$ and $\mathrm{Reg}(x;\theta) = A(\theta) - \theta^\top x$. Regret is thus the excess realized smoothed penalty over the hindsight minimum. Expressing the decision through unupgraded fractions retains the parameter-dependent total penalty $\sum_j q_j$. Dropping this term when rewriting the objective as upgrade rewards would change the objective-robust baselines. Calibration and optimization use the parameter domain $\mathbb{R}^4$.

Chronological protocol. We evaluate twelve 24-month windows, starting each January from 2012 through 2023. Within each window, we use the first eight months to fit the predictor and residual geometry, the next eight for calibration, and the final eight for testing. This gives 96 test months covering May–December of 2013–2024. The predictor is the training mean of θ, and each method calibrates its own score. We set $\alpha = 0.2$, which selects the eighth order statistic among eight calibration scores; $\alpha = 0.1$ would require the augmented ninth value $+\infty$. This chronological evaluation measures decision performance and empirical coverage over time; the finite-sample guarantee in Theorem 3 requires exchangeability of the calibration and test pairs.

The box uses the scale learner from Appendix B.1 with a constant input. KNN RO uses all eight training observations, so it coincides with the residual ellipsoid under the mean predictor.

Exact FY reformulation. Write $h = \widehat{\theta}$ and $\eta = \eta_{\alpha,2}$. The feasible set has eleven vertices $V = \{v \in \{0,1\}^4 : \mathbf{1}^\top v \geq 2\}$. For each $v \in V$, let D_v have rows $(u - v)^\top$ for $u \in V \setminus \{v\}$. Its

Table 8: Controlled QP across shifts and coverage targets. Each entry is mean regret / CVaR_{95}, multiplied by 10^3. PTO is invariant to α.

κ	α	PTO	FY	Residual ℓ_2	Residual box
0	0.1	1.413 / 4.360	1.424 / 4.440	2.327 / 5.453	1.431 / 4.395
0	0.3	1.413 / 4.360	1.421 / 4.438	2.164 / 5.218	1.427 / 4.380
0	0.4	1.413 / 4.360	1.415 / 4.402	2.111 / 5.146	1.425 / 4.378
0	0.5	1.413 / 4.360	1.413 / 4.374	2.063 / 5.083	1.424 / 4.375
0.5	0.1	1.464 / 4.485	1.473 / 4.544	2.862 / 6.302	1.484 / 4.549
0.5	0.3	1.464 / 4.485	1.474 / 4.528	2.679 / 6.021	1.479 / 4.525
0.5	0.4	1.464 / 4.485	1.468 / 4.501	2.621 / 5.934	1.478 / 4.520
0.5	0.5	1.464 / 4.485	1.459 / 4.460	2.562 / 5.847	1.476 / 4.515
1	0.1	1.748 / 4.953	1.756 / 4.974	3.537 / 7.417	1.787 / 5.041
1	0.3	1.748 / 4.953	1.758 / 4.950	3.342 / 7.114	1.779 / 5.013
1	0.4	1.748 / 4.953	1.754 / 4.945	3.279 / 7.018	1.776 / 5.005
1	0.5	1.748 / 4.953	1.758 / 4.952	3.216 / 6.920	1.774 / 4.999
2	0.1	3.172 / 7.307	2.888 / 6.904	5.282 / 10.304	3.236 / 7.468
2	0.3	3.172 / 7.307	2.988 / 7.094	5.080 / 10.007	3.223 / 7.434
2	0.4	3.172 / 7.307	3.039 / 7.186	5.015 / 9.910	3.218 / 7.424
2	0.5	3.172 / 7.307	3.103 / 7.300	4.940 / 9.799	3.215 / 7.414
4	0.1	8.719 / 15.908	8.730 / 15.917	9.971 / 17.404	8.879 / 16.346
4	0.3	8.719 / 15.908	8.285 / 15.313	9.873 / 17.285	8.850 / 16.259
4	0.4	8.719 / 15.908	8.161 / 15.167	9.835 / 17.239	8.840 / 16.229
4	0.5	8.719 / 15.908	8.665 / 15.916	9.798 / 17.193	8.832 / 16.206

parameter cell is

$$C_v(h,\eta) = \{\theta : D_v\theta \le 0,\ -D_v\theta \le \eta\mathbf{1} - 2D_vh\}. \tag{41}$$

Under the common prediction-preferred rule, (28) gives the FY set as $\bigcup_{v\in V} C_v(h,\eta)$. A cell is empty if and only if $2[A(h) - h^\top v] > \eta$. For the remaining cells, put $B_v = [D_v; -D_v]$ and $b_v = [0; \eta\mathbf{1} - 2D_vh]$. Dualizing the linear pricing problem in each cell gives

$$\begin{aligned} &\min_{x,t,\{y_v\}} \quad t \\ &\text{subject to} \quad x \in \mathcal{X}, \quad y_v \ge 0, \quad B_v^\top y_v = v - x, \quad b_v^\top y_v \le t \quad \text{for every nonempty } C_v. \end{aligned} \tag{42}$$

We include every nonempty cell and screen the others using the stated condition. Solving the primal pricing problems independently on the retained cells certifies all FY decisions within 10^{-7} over the full set.

Results and allocation geometry. Table 2 reports regret and empirical coverage under the chronological evaluation. FY reduces tail regret relative to PTO at a small cost in mean regret. Its worst-case criterion does not weight scenarios by their probabilities, so it need not minimize mean regret or CVaR under the data distribution. Paired resampling of the twelve test-year blocks gives a 95% interval of $[-0.1521, -0.0288]$ for FY minus PTO in CVaR_{95}, in the original units. The corresponding FY-minus-ℓ_2 interval contains zero.

To explain why FY and the residual ball give similar allocations, let $b = \frac{1}{2}\mathbf{1}$ and $P = I_4 - \mathbf{1}\mathbf{1}^\top/4$. For a residual ball of radius r, the objective-robust problem is $\min_{x\in\mathcal{X}}\{-h^\top x + r\|x\|_2\}$. Since the fitted coordinates of h are negative, the solution uses the full upgrade budget, $\mathbf{1}^\top x = 2$. Writing $g = Ph$, the interior solution on this face is

$$x_{\ell_2} = b + \frac{g}{\sqrt{r^2 - \|g\|_2^2}}, \qquad r > \|g\|_2, \tag{43}$$

provided this vector lies in $[0,1]^4$. To obtain the expression, substitute $x = b + y$ with $\mathbf{1}^\top y = 0$, reducing the objective to $-g^\top y + r\sqrt{1 + \|y\|_2^2}$ plus a constant. Thus, when differences between station predictions are small relative to r, the allocation is close to b.

For FY, consider the symmetric case $h = -c\mathbf{1}$, $c > 0$, with $0 < \eta < 2c$. The screening rule retains exactly the six vertices satisfying $\mathbf{1}^\top v = 2$. Choose the orthonormal basis

$$Q = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & -1 \\ -1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}, \qquad Q^\top Q = I_3, \qquad Q^\top \mathbf{1} = 0.$$

The budget face can be written as $x = b + Qw$ with $\|w\|_1 \le 1$, and its six vertices correspond to $\pm e_1, \pm e_2, \pm e_3$. For $\theta = -c\mathbf{1} + Qy$ with $\|y\|_\infty \le \eta/2$, the true and reflected optimizers both lie on this face, giving $S_2(\theta, h) = 2\|y\|_\infty$. Conversely, comparison with these six vertices gives $S_2(\theta, h) \ge 2\|Q^\top \theta\|_\infty$ throughout the FY set. Together, these observations show that its projection onto the decision-relevant subspace is the cube

$$Q^\top \mathcal{U}_{2,\eta}(h) = \{y : \|y\|_\infty \le \eta/2\}.$$

Consequently, the two robust criteria on the budget face are

$$\begin{aligned} \rho_{\mathcal{U}_{2,\eta}(h)}(b + Qw) &= \sup_{\|y\|_\infty \le \eta/2} \{\|y\|_\infty - y^\top w\} = \frac{\eta}{2}(1 + \|w\|_1), \\ -h^\top (b + Qw) + r\|b + Qw\|_2 &= 2c + r\sqrt{1 + \|w\|_2^2}. \end{aligned} \tag{44}$$

Both criteria are uniquely minimized at $w = 0$. For any $x \in \mathcal{X}$ outside this face, the parameters $-c\mathbf{1} + Qy$ give FY worst-case regret at least $c(\mathbf{1}^\top x - 2) + (\eta/2)(1 + \|Q^\top x\|_1) > \eta/2$; the ℓ_2 criterion also increases with the unupgraded total. Hence b is the unique global solution of both problems.

The observed predictions have small stationwise contrasts: $\|Ph\|_2/r$ is at most 0.00962 across the twelve windows. This is partly explained by the training-mean construction. Centering at τ_j makes the linear term in the local expansion $\mathrm{sigmoid}(t) = 1/2 + t/4 + O(t^3)$ average to zero for $t = 0.5(N_j - \tau_j)$. The resulting FY and residual ℓ_2 decisions have a mean Euclidean distance of 0.00210 and a maximum of 0.00556. Replacing h by its coordinate mean while keeping each window's calibrated thresholds gives b for both methods in all twelve windows.