

Highlights

Data-Driven Minimax-Regret Portfolio Optimization under Tail-Risk Ambiguity

Zheqi Fan

- Audit relative tail-risk rankings, not Expected Shortfall levels, and project the training mix onto admissible engine mixtures.
- Implement one locked interpolation between the validated mix and worst-case admissible regret, not a Wasserstein ball on return laws.
- Bound true Expected Shortfall regret in finite samples, and treat an empty admissible set as certified rejection of the engine library.
- Solve the rule as a linear program by enumerating extreme points of a small engine polytope.

Data-Driven Minimax-Regret Portfolio Optimization under Tail-Risk Ambiguity

Zheqi Fan^{a,b,*}

^a*Division of Emerging Interdisciplinary Areas, Hong Kong University of Science and Technology, Hong Kong SAR*

^b*Thrust of Financial Technology, Hong Kong University of Science and Technology (Guangzhou), China*

Abstract

We study data-driven portfolio choice under Expected Shortfall when a short training sample can distort a plug-in optimizer and when competing tail-risk engines disagree on which feasible portfolios are safer. Checking only the level of Value-at-Risk or Expected Shortfall can miss this disagreement: a common additive shift leaves every ranking unchanged, yet a relative distortion can change the optimizer. We form a training-score mix of the engines, independently audit relative Expected Shortfall differences on a frozen audit library, and project that mix onto the resulting admissible set of engine mixtures. The implemented rule interpolates, with a parameter locked in advance, between that validated mix and worst-case admissible regret. We prove a high-probability bound on true Expected Shortfall regret that separates model tolerance, validation error, coverage error, and optimization error. An empty admissible set, under the stated validation model, yields a level- β rejection of the engine library at the prescribed tolerance, rather than a solver failure. For a small finite library the rule is a linear program obtained from the extreme points of the admissible set. We do not claim a polynomial dual reformulation in general, and we do not place a Wasserstein ball around return laws.

Keywords: Expected Shortfall, minimax regret, portfolio optimization, data-driven optimization, finite-sample guarantee

1. Introduction

A portfolio manager who reports Expected Shortfall must turn noisy return samples and several competing tail-risk engines into one implementable weight vector. Two failures occur together. A short training window can make a plug-in optimizer chase an atypical sample. Independently, engines that look plausible in training can disagree on *which* feasible portfolios are safer, even when their absolute Expected Shortfall levels look similar on a few benchmarks. The operating decision is therefore not only “which weights minimize estimated Expected Shortfall,” but “which engine mix remains admissible after an independent audit, and how much robustness to retain if

*Corresponding author.

Email address: `zheqi.fan@connect.ust.hk` (Zheqi Fan)

several mixes still pass.” The insight of this paper is a two-sample rule with an explicit rejection option: score the engines on the training fold; audit relative tail-risk rankings on a held-out sample; project the training mix onto the mixes that pass; then interpolate, with a parameter locked in advance, between that validated mix and the worst admissible regret. If no mix passes, reject the library rather than force a portfolio. The resulting high-probability bound on true Expected Shortfall regret is conservative under a distribution-free radius, yet decomposable and decision-relevant; the construction is not a regime-switching Wasserstein program on return laws.

A perturbation that looks small in a risk forecast can still reorder feasible portfolios and change the optimizer (Goldfarb and Iyengar, 2003; Zhu and Fukushima, 2009; Huang et al., 2010). Expected Shortfall (ES), equivalently Conditional Value-at-Risk under the convention used here, is a natural objective for scenario-based construction: it is coherent and admits a convex representation that becomes linear under finite scenarios (Artzner et al., 1999; Acerbi and Tasche, 2002; Rockafellar and Uryasev, 2002). Worst-case and relative-robust CVaR models protect decisions against distributional uncertainty (Zhu and Fukushima, 2009; Huang et al., 2010), while a parallel stream builds uncertainty sets from estimation error and studies finite-sample guarantees for data-driven optimization (Delage and Ye, 2010; Bertsimas et al., 2018; Mohajerin Esfahani and Kuhn, 2018; Van Parys et al., 2021). These literatures do not always separate two finite-sample difficulties that arise simultaneously in tail-risk portfolio governance.

Sampling and parameter uncertainty. Even under correct specification, a short training sample can be atypical, and estimation errors transfer into and are amplified by optimization (the optimizer’s curse; (Smith and Winkler, 2006)). Estimate-then-optimize (ETO) pipelines treat a point estimate as truth. A complementary perspective treats the data-driven solution itself as a *function of the sample* and evaluates out-of-sample performance by averaging over the sampling distribution. That viewpoint is central to operational statistics (Liyanage and Shanthikumar, 2005), to the Optimize-via-Estimate programme of Loke et al. (2026), and to related finite-sample analyses (Besbes and Mouchtaki, 2023). We take from that literature only the two-sample discipline: a $\mathcal{G}_{\text{tr}}$ -measurable training-score anchor, followed by independent validation. We do not invoke their sufficient-statistic prior-optimality construction, and we do not interpret training scores as Bayesian posteriors.

Risk-function discrepancy. Candidate ES engines can disagree on relative risk across feasible portfolios even when absolute levels look similar on a few benchmarks. Level-based VaR/ES checks need not control portfolio choice: common additive shifts leave rankings unchanged, whereas relative-risk distortions can alter the optimizer. A decision-relevant discrepancy that is invariant to common additive shifts connects validation of relative ES differences to portfolio regret over a continuous feasible set. Existing methods typically emphasize one dimension: robust and distributionally robust optimization hedge uncertainty sets, often over return laws (Delage and Ye, 2010; Bertsimas et al., 2018; Mohajerin Esfahani and Kuhn, 2018); a training-only score

can reweight engines after a short sample without confronting independent evidence that the aggregation fails validation; forecast validation and elicibility research compare scores (Fissler and Ziegel, 2016) yet often stop short of a downstream ES-regret certificate; and predictive-to-prescriptive pipelines (Bertsimas and Kallus, 2020; Ban and Rudin, 2019; Mišić and Perakis, 2020) motivate aligning learning with decisions without, by themselves, delivering the audit object used here.

Decision-focused learning and finite-sample prescriptive analytics. Smart Predict-then-Optimize (SPO) of Elmachoub and Grigas (2022) is a foundational motivation for evaluating decisions by *decision regret* rather than by pure estimation error: the SPO loss measures the suboptimality gap induced by optimizing a predicted objective, and the convex SPO+ surrogate makes this training-stage criterion tractable for linear objectives under contextual covariates. That perspective is complementary to, but distinct from, the present framework. SPO/SPO+ is primarily designed for linear (or similarly structured) objective predictions in contextual settings and does not, by itself, provide finite-sample distribution-free regret certificates when candidate *nonlinear* tail-risk engines are severely misspecified. We do not develop a “tail-risk SPO+” surrogate theorem. Instead, we keep SPO as a conceptual cornerstone for caring about decision regret, and shift the methodological focus from training-stage surrogate minimization to decision-relevant risk-function ambiguity sets with explicit high-probability out-of-sample Expected Shortfall regret bounds. A training-score anchor addresses sampling variability only on the training fold; the hybrid rule further requires independent decision-relevant validation of risk-function aggregations. The gap is therefore operational as well as statistical: the desk needs a finite-sample rule that accounts for training-sample uncertainty and still refuses to optimize on engine mixes that fail an independent, decision-relevant audit. The resulting policy is two-sample, allows explicit rejection, and carries an Expected Shortfall-regret certificate.

1.1. Contributions

This paper changes the engine-mix decision, not the return-law ball. The implemented policy is a locked interpolation between a validated training mix and worst-case admissible Expected Shortfall regret, with certified rejection when the mix set is empty. Relative to Fan (2026), which develops validation-calibrated risk-function ambiguity for pure minimax regret over admissible engine aggregations, the present paper implements the hybrid rule H_λ that interpolates between a training-score mix and worst-case admissible regret after the same style of holdout audit. The supporting contributions are as follows.

1. *Two-sample admissible mixes.* A $\mathcal{G}_{\text{tr}}$ -measurable training-score anchor q^0 is paired with independent validation of relative ES differences on a frozen audit library. The resulting polytope $\widehat{\mathcal{Q}}_\delta$ over convex aggregations of fixed engines yields the validated anchor q^{vpw} by Euclidean projection; an empty polytope is treated as certified rejection of the library.

2. *Decision-relevant discrepancy.* An additive-shift-invariant discrepancy $d_{\mathcal{X}}$ and a regret-transfer lemma connect risk-function audits to portfolio ES regret, so level-based VaR/ES scorecards are not used as decision certificates.
3. *Hybrid rule and finite-sample certificate.* A locked interpolation H_{λ} between $r_{q^{\text{vpw}}}$ and worst-case admissible regret comes with a high-probability bound on true ES regret r_0 , separating model tolerance, validation error, coverage error, and optimization error, together with a bound on the decision-level cost of the validation correction.
4. *Tractable finite-library form.* Aggregate-risk regret is convex in the mix weights, so the inner supremum is attained at extreme points of $\widehat{\mathcal{Q}}_{\delta}$ and the hybrid rule is an LP for a small engine library. We do not claim that enumeration is polynomial in the number of audit inequalities, and we do not claim outperformance of $1/N$.

1.2. Positioning

The ambiguity object is a convex combination of *implemented ES risk functions*; the certificate object is true ES regret r_0 . Fan (2026) uses a related audit to calibrate an ambiguity set and then solves pure minimax regret; here the implemented decision is the locked hybrid H_{λ} rather than that pure minimax map. This differs from SPO/SPO+ and related predictive-to-prescriptive pipelines (Elmachtoub and Grigas, 2022; Liu and Grigas, 2021; El Balghiti et al., 2019; Bertsimas and Kallus, 2020; Ban and Rudin, 2019), which keep a decision-regret motive but typically train feature-to-decision or linear-objective surrogates: those guarantees are not holdout ES-regret certificates under misspecified nonlinear engines, and our engine weights are prior scores constrained by a validation polytope rather than a learned contextual policy. It also differs from classical DRO on return laws (Delage and Ye, 2010; Mohajerin Esfahani and Kuhn, 2018; Van Parys et al., 2021; Lam, 2019), which enlarge a ball around P (moments, Wasserstein, or empirical divergence), and from worst-case / relative-robust CVaR (Zhu and Fukushima, 2009; Huang et al., 2010; Xidonas et al., 2017), which optimize levels or Huang-style engine-wise regret $\tilde{r}_q = r_q + I(q)$: we control relative ES differences of risk-function aggregations on an audit library and optimize oracle-consistent aggregate-risk regret r_q and H_{λ} . Table I summarizes these object-level distinctions for the main neighboring streams.

Core neighbor. Pun et al. (2023) is the closest recent data-driven CVaR portfolio paper in the return-law DRO family: a regime-switching Wasserstein ball on P with an inclusion-style guarantee and a DRO-CVaR rule. Relative to that line (Table I), we change the ambiguity variable to mix weights q , certify holdout r_0 rather than inclusion of P_0 , keep radii audit-driven rather than performance-tuned, and treat empty $\widehat{\mathcal{Q}}_{\delta}$ as a first-class rejection output instead of claiming that $1/N$ is dominated. A clustered Wasserstein-CVaR comparator appears only in the Online Supplement as a labeled return-law benchmark, not as our estimator.

Table I: Object-level positioning against core neighbors. Abbreviations: W = Wasserstein; DRO = distributionally robust optimization; CVaR = conditional value-at-risk (equivalent to ES under our convention).

Stream	Ambiguity object	Finite-sample / certificate target	Decision object
This paper	polytope on mix weights q over fixed ES maps	holdout bound on true ES regret r_0 ; empty $\hat{\mathcal{Q}}_\delta$ rejects	hybrid H_λ on aggregate-risk regret r_q
Pun et al. (2023)	regime-switching W ambiguity set on return law P	high-probability inclusion of P_0 above a sample-size bound; radius tuned in implementation	data-driven DRO CVaR portfolio rule
Zhu and Fukushima (2009) ; Huang et al. (2010)	discrete distributional uncertainty / multiple expert priors	robust or relative optimality versus those priors (not a holdout r_0 certificate)	worst-case CVaR level (Zhu and Fukushima, 2009) or relative-robust CVaR (Huang et al., 2010)
Mohajerin Esfahani and Kuhn (2018)	W ball around the empirical return law	finite-sample / asymptotic performance guarantees for the DRO value	minimize worst-case risk over the W ball
Van Parys et al. (2021)	relative-entropy (KL) ball around the empirical law	least-conservative predictor-prescriptor under out-of-sample disappointment limits	DRO with relative-entropy ambiguity is uniquely optimal in that meta-problem
Elmachtoub and Grigas (2022) ; Liu and Grigas (2021)	predicted linear (or structured) objective costs	excess SPO risk and related generalization / calibration bounds	Smart Predict-then-Optimize / SPO+ training surrogate

Neighboring holdout and finite-sample strands are likewise adjacent rather than interchangeable: end-to-end DRO pipelines may learn robustness from validation outcomes (Costa and Iyengar, 2023), whereas we forbid using validation to select δ ; path-validation work can certify CVaR feasibility with abstention (Long, 2025), whereas we certify relative ES differences and hybrid regret; transient newsvendor analyses (Besbes and Mouchtaki, 2023) give sample-complexity schedules for a single loss, whereas our certificate covers a finite library of structurally misspecified tail engines after an independent audit. In one sentence: lock a training-score mix, independently audit relative Expected Shortfall rankings, project onto what remains admissible, and interpolate toward worst-case admissible regret, or reject the library if nothing remains. Validation controls decision-relevant discrepancy within the accepted engine class; a large certificate means the audit is not yet sharp, not that the portfolio is safe. Extensions to other coherent risk measures are left as future work.

2. Problem setting and the hybrid rule

This section assembles the objects used later: the feasible set and true-ES regret, the two-sample split, the engine library, a $\mathcal{G}_{\text{tr}}$ -measurable training-score anchor, a decision-relevant audit of risk-function aggregations, and the hybrid objective H_λ . Certificates and the computational form of H_λ are deferred to Section 3.

2.1. Feasible set and loss

Let

$$\mathcal{X} = \{x \in \mathbb{R}^n : x \geq 0, \mathbf{1}^\top x = 1, Gx \leq g\} \quad (2.1)$$

be nonempty, compact, and convex. For a holding-period return vector R , define the buy-and-hold loss $L(x, R) = -x^\top R$. We do not use within-period rebalanced product losses, which destroy linearity in the initial weights.

Let P_0 be a deployment reference law. For tail level $\alpha \in (0, 1)$ and $p = 1 - \alpha$,

$$f_0(x) = \text{ES}_{\alpha, P_0}(L(x, R)), \quad r_0(x) = f_0(x) - \min_{y \in \mathcal{X}} f_0(y). \quad (2.2)$$

2.2. Training and validation separation

Let Y^{tr} be the training sample and $Y^{\text{val}} = \{R_1^v, \dots, R_T^v\}$ an independent validation sample. Let $\mathcal{G}_{\text{tr}}$ contain the fitted engines, training scores, preference weights, audit library, tolerances, and feasible set.

Assumption 2.1 (Conditional i.i.d. validation). *Conditionally on $\mathcal{G}_{\text{tr}}$, the validation returns are i.i.d. from P_0 .*

Serial dependence and deployment drift are discussed in Section 5.

2.3. Candidate ES engines

Fix M engines with scenarios (R_{ms}, π_{ms}) and Rockafellar–Uryasev representations

$$f_m(x) = \min_{\zeta_m} \left\{ \zeta_m + \frac{1}{p} \sum_{s=1}^{S_m} \pi_{ms} \left[-x^\top R_{ms} - \zeta_m \right]_+ \right\}. \quad (2.3)$$

For $q \in \Delta_M = \{q \geq 0 : \mathbf{1}^\top q = 1\}$,

$$f_q(x) = \sum_m q_m f_m(x) \quad (2.4)$$

aggregates risk functions and is not necessarily ES under a mixture law.

The default implemented library, all fitted on Y^{tr} only, is

1. historical empirical scenarios;
2. a Gaussian engine with Ledoit–Wolf shrinkage covariance ([Ledoit and Wolf, 2004](#));
3. a multivariate Student- t engine with pre-specified degrees of freedom;
4. a two-state Gaussian mixture obtained by K -means on the training cloud.

Numerical cells may drop engines to $M \leq 4$ as locked in the protocol. No engine is tuned on the validation fold.

The engines enter the optimizer only through a convex aggregation f_q . The next subsection constructs a training-fold starting weight q^0 ; independent validation then decides which aggregations remain admissible.

2.4. Training-score anchor

Let $a_m(Y^{\text{tr}}) > 0$ be a training-fold fit or out-of-sample proxy score for engine m (for example a Laplace/BIC score or a cross-validated likelihood score), and let $\omega \in \Delta_M$ be a pre-specified preference weight. The default implementation uses

$$\log a_m(Y^{\text{tr}}) = \log g_m(Y^{\text{tr}} \mid \hat{\theta}_m) - \frac{\dim(\Theta_m)}{2} \log N_{\text{tr}}, \quad (2.5)$$

where $g_m(\cdot \mid \hat{\theta}_m)$ is an engine likelihood, or a ridge-Gaussian proxy for the historical engine, and $\dim(\Theta_m)$ counts free mean/covariance parameters. The historical, Ledoit–Wolf, Student- t , and K -means mixture engines do not share a common low-dimensional sufficient statistic, so the scores are merely comparable training diagnostics.

Definition 2.1 (Training-score anchor).

$$q_m^0 = \frac{\omega_m a_m(Y^{\text{tr}})}{\sum_{\ell=1}^M \omega_\ell a_\ell(Y^{\text{tr}})}. \quad (2.6)$$

Proposition 2.1 (Well-definedness). *If $\sum_m \omega_m a_m(Y^{\text{tr}}) > 0$, then $q^0 \in \Delta_M$.*

The vector q^0 is $\mathcal{G}_{\text{tr}}$ -measurable by construction and is *not* a Bayesian posterior. Subsequent projection onto $\widehat{\mathcal{Q}}_\delta$ and the regret certificates hold for every such measurable anchor; they do not require a prior-optimality, uniform-optimality, or sufficient-statistic interpretation. A training score does not, by itself, certify that f_{q^0} reproduces decision-relevant relative tail risk on an independent sample. That is the role of the audit constructed next.

2.5. Decision-relevant validation

Definition 2.2 (Decision discrepancy). *For continuous $f, g : \mathcal{X} \rightarrow \mathbb{R}$,*

$$d_{\mathcal{X}}(f, g) = \sup_{x, y \in \mathcal{X}} |[f(x) - f(y)] - [g(x) - g(y)]| = \max_{x \in \mathcal{X}} (f - g)(x) - \min_{x \in \mathcal{X}} (f - g)(x). \quad (2.7)$$

Thus $d_{\mathcal{X}}(f, g) = 0$ if and only if $f - g$ is constant on $\mathcal{X}$.

Lemma 2.1 (Regret transfer). *Let $r_f(x) = f(x) - \min_{y \in \mathcal{X}} f(y)$. For every $x \in \mathcal{X}$,*

$$|r_f(x) - r_g(x)| \leq d_{\mathcal{X}}(f, g). \quad (2.8)$$

Proof. Let $x_f^* \in \arg \min_{y \in \mathcal{X}} f(y)$ and $x_g^* \in \arg \min_{y \in \mathcal{X}} g(y)$, and write $\Delta = f - g$. Then

$$r_f(x) - r_g(x) = (f(x) - f(x_f^*)) - (g(x) - g(x_g^*)) \quad (2.9)$$

$$= \Delta(x) - \Delta(x_f^*) + (g(x_g^*) - g(x_f^*)). \quad (2.10)$$

Optimality of x_g^* for g forces $g(x_g^*) - g(x_f^*) \leq 0$, so

$$r_f(x) - r_g(x) \leq \Delta(x) - \Delta(x_f^*) \leq \max_{\mathcal{X}} \Delta - \min_{\mathcal{X}} \Delta = \text{osc}_{\mathcal{X}}(\Delta) = d_{\mathcal{X}}(f, g). \quad (2.11)$$

Exchanging the roles of f and g yields the opposite inequality. If $f - g$ is constant on $\mathcal{X}$ then $d_{\mathcal{X}}(f, g) = 0$ and $r_f \equiv r_g$: additive level shifts cancel in regret and cannot change the optimizer. $\square$

Let $\mathcal{B} = \{b_1, \dots, b_J\} \subset \mathcal{X}$ be a pre-specified audit library, and write

$$\theta_j = f_0(b_j), \quad \widehat{\theta}_j = \widehat{\text{ES}}_\alpha(b_j; Y^{\text{val}}), \quad A_{jm} = f_m(b_j). \quad (2.12)$$

Define the audit discrepancy

$$d_{\mathcal{B}}(q; \theta) = \text{rng}(Aq - \theta). \quad (2.13)$$

The remainder of this subsection turns that finite audit into an admissible polytope for q , a certified-rejection rule when the polytope is empty, and the validated projection of q^0 .

Assumption 2.2 (Bounded deployment and scenario returns). *There exists a constant $\bar{R}_0 < \infty$ such that $\|R\|_\infty \leq \bar{R}_0$ almost surely under P_0 . Moreover, all scenarios generated by the implemented risk engines are pre-truncated to satisfy*

$$\max_{m \in [M], s \in [S_m]} \|R_{ms}\|_\infty \leq \bar{R}_{\text{eng}}. \quad (2.14)$$

We set

$$\bar{R} := \max\{\bar{R}_0, \bar{R}_{\text{eng}}\}. \quad (2.15)$$

Let $e_T = e_T(\alpha, \beta, \bar{R}, T, J)$ be a simultaneous radius such that

$$\mathbb{P}\left(\max_j |\hat{\theta}_j - \theta_j| \leq e_T \mid \mathcal{G}_{\text{tr}}\right) \geq 1 - \beta. \quad (2.16)$$

The validation losses associated with different audit portfolios may be dependent because they are evaluated on the same validation returns. Nevertheless, conditional on $\mathcal{G}_{\text{tr}}$, a union bound over the fixed audit library yields simultaneous coverage by allocating failure probability β/J to each audit portfolio. A DKW-based translation to ES yields one sufficient (conservative) construction (Massart, 1990); Lemma Appendix A.1 records the official formula (A.2)–(A.3). That DKW-to-ES radius is the paper’s official finite-sample guarantee. Massart’s inequality (Massart, 1990) supplies an explicit two-sided CDF band; a union bound over the frozen audit library then yields a simultaneous ES radius that is fully transparent for certification. Sharper finite-sample CVaR/ES concentration bounds are available (Prashanth et al., 2020) and can improve leading constants once the empirical quantile and the tail residual are controlled jointly. We retain the DKW route for the official guarantee precisely because the audit set is finite and $\mathcal{G}_{\text{tr}}$ -measurable, dependence across audits is absorbed by the union bound, and the resulting formula needs no auxiliary variance estimation inside the certificate. A variance-adaptive empirical-ES inequality would have to control the empirical quantile and the tail residual process jointly: the positive-part residuals $(L_{jt} - \hat{\nu}_j)_+$ are not an i.i.d. sample independent of $\hat{\nu}_j$, so a direct empirical-Bernstein / Maurer–Pontil bound (Maurer and Pontil, 2009) on those residuals is not valid. We leave a tight variance-aware ES radius as future work.

Gaussian, Student- t , and mixture engines are theoretically unbounded. The implemented library therefore applies coordinate-wise winsorization at $\pm \bar{R}_{\text{eng}}$ before any ES evaluation, so every scenario cloud used by f_m lies in the same ℓ_∞ ball as the clipped deployment DGP. All official cells take $\bar{R}_0 = \bar{R}_{\text{eng}} = 0.20$.

Definition 2.3 (Admissible aggregations).

$$\mathcal{Q}_\delta = \{q \in \Delta_M : d_{\mathcal{B}}(q; \theta) \leq \delta\}, \quad (2.17)$$

$$\hat{\mathcal{Q}}_\delta = \{q \in \Delta_M : d_{\mathcal{B}}(q; \hat{\theta}) \leq \delta + 2e_T\}. \quad (2.18)$$

Lemma 2.2 (Range perturbation). *If $\max_j |\hat{\theta}_j - \theta_j| \leq e_T$, then for every q ,*

$$|\text{rng}(Aq - \hat{\theta}) - \text{rng}(Aq - \theta)| \leq 2e_T. \quad (2.19)$$

Proof. For each coordinate, $|(Aq - \hat{\theta})_j - (Aq - \theta)_j| = |\theta_j - \hat{\theta}_j| \leq e_T$. Hence

$$\max_j (Aq - \hat{\theta})_j \leq \max_j (Aq - \theta)_j + e_T, \quad \min_j (Aq - \hat{\theta})_j \geq \min_j (Aq - \theta)_j - e_T, \quad (2.20)$$

so

$$\text{rng}(Aq - \hat{\theta}) \leq \text{rng}(Aq - \theta) + 2e_T. \quad (2.21)$$

The opposite inequality follows by exchanging θ and $\hat{\theta}$. $\square$

Theorem 2.4 (Set containment). *On $\{\max_j |\hat{\theta}_j - \theta_j| \leq e_T\}$,*

$$\mathcal{Q}_\delta \subseteq \hat{\mathcal{Q}}_\delta \subseteq \mathcal{Q}_{\delta+4e_T}. \quad (2.22)$$

Proof. On the event of Lemma 2.2, if $q \in \mathcal{Q}_\delta$ then $\text{rng}(Aq - \theta) \leq \delta$, hence

$$\text{rng}(Aq - \hat{\theta}) \leq \delta + 2e_T, \quad (2.23)$$

so $q \in \hat{\mathcal{Q}}_\delta$. If $q \in \hat{\mathcal{Q}}_\delta$ then $\text{rng}(Aq - \hat{\theta}) \leq \delta + 2e_T$, hence

$$\text{rng}(Aq - \theta) \leq \delta + 4e_T, \quad (2.24)$$

so $q \in \mathcal{Q}_{\delta+4e_T}$. $\square$

Emptiness of $\hat{\mathcal{Q}}_\delta$ is a first-class diagnostic, not a solver failure. The next theorem makes that statement quantitative.

Theorem 2.5 (Audit adequacy and certified rejection). *Define the oracle and validation minimum audit discrepancies*

$$\delta_B^* := \min_{q \in \Delta_M} d_{\mathcal{B}}(q; \theta), \quad \hat{\delta}_B := \min_{q \in \Delta_M} d_{\mathcal{B}}(q; \hat{\theta}). \quad (2.25)$$

On the simultaneous validation event

$$\mathcal{E}_T = \{\max_{j \in [J]} |\hat{\theta}_j - \theta_j| \leq e_T\}, \quad (2.26)$$

$$|\hat{\delta}_B - \delta_B^*| \leq 2e_T. \quad (2.27)$$

Consequently:

(i) *If $\delta_B^* \leq \delta$, then $\hat{\mathcal{Q}}_\delta \neq \emptyset$.*

(ii) If $\widehat{\mathcal{Q}}_\delta = \emptyset$, then $\delta_B^* > \delta$ (certified rejection of the engine library).

(iii) If $\delta_B^* > \delta + 4e_T$, then $\widehat{\mathcal{Q}}_\delta = \emptyset$ with probability at least $1 - \beta$.

Proof. Lemma 2.2 gives

$$|d_{\mathcal{B}}(q; \widehat{\theta}) - d_{\mathcal{B}}(q; \theta)| \leq 2e_T \quad (2.28)$$

for every $q \in \Delta_M$. Taking the infimum over q on both sides yields $|\widehat{\delta}_B - \delta_B^*| \leq 2e_T$. (i) If $\delta_B^* \leq \delta$, some q^* satisfies $d_{\mathcal{B}}(q^*; \theta) \leq \delta$, hence $d_{\mathcal{B}}(q^*; \widehat{\theta}) \leq \delta + 2e_T$ and $q^* \in \widehat{\mathcal{Q}}_\delta$. (ii) If $\widehat{\mathcal{Q}}_\delta = \emptyset$, then $\widehat{\delta}_B > \delta + 2e_T$, so $\delta_B^* \geq \widehat{\delta}_B - 2e_T > \delta$. (iii) On $\mathcal{E}_T$, $\widehat{\delta}_B \geq \delta_B^* - 2e_T$. The hypothesis $\delta_B^* > \delta + 4e_T$ therefore forces $\widehat{\delta}_B > \delta + 2e_T$, i.e. $\widehat{\mathcal{Q}}_\delta = \emptyset$. The event $\mathcal{E}_T$ has probability at least $1 - \beta$. $\square$

Definition 2.6 (Validated training-score aggregation). *If $\widehat{\mathcal{Q}}_\delta \neq \emptyset$,*

$$q^{\text{vpw}} \in \arg \min_{q \in \widehat{\mathcal{Q}}_\delta} \frac{1}{2} \|q - q^0\|_2^2. \quad (2.29)$$

We retain the superscript in q^{vpw} as shorthand for the validated (prior/score-weighted) anchor. This is a convex quadratic program. A Kullback–Leibler projection is an optional appendix variant. If $q^0 \in \widehat{\mathcal{Q}}_\delta$, then the validation displacement is zero. What remains is to choose a portfolio from the admissible aggregations: the hybrid rule interpolates between the validated anchor and worst-case admissible regret.

2.6. Hybrid portfolio rule

For $q \in \Delta_M$, define aggregate-risk regret

$$v(q) = \min_{y \in \mathcal{X}} f_q(y), \quad r_q(x) = f_q(x) - v(q). \quad (2.30)$$

With $v_m = \min_y f_m(y)$,

$$\sum_{m=1}^M q_m [f_m(x) - v_m] = r_q(x) + I(q), \quad I(q) = v(q) - \sum_{m=1}^M q_m v_m \geq 0. \quad (2.31)$$

The term $I(q)$ is oracle incompatibility and does not depend on the implemented x ; aggregate-risk regret is the correct object for optimizing f_q .

Remark 2.1 (Huang relative-robust CVaR). *Write*

$$\widetilde{r}_q(x) = r_q(x) + I(q), \quad (2.32)$$

Worst-case ES (Zhu and Fukushima, 2009) hedges a worst engine level $\max_m f_m(x)$. Relative-robust CVaR (Huang et al., 2010) instead compares each engine to its own oracle through

$$\max_m (f_m(x) - v_m), \quad (2.33)$$

which is $\tilde{r}$ at simplex vertices. Both constructions are coherent when the decision maker refuses to mix engines, or when the hedge is chosen to be engine-wise. They are not the regret of the aggregate risk f_q : the gap $I(q) = v(q) - \sum_m q_m v_m$ is nonnegative and independent of the implemented x , so optimizing $\tilde{r}_q$ (or its vertex form) targets a strictly larger object than r_q . Our hybrid rule and certificate are built around aggregate-risk regret r_q and H_λ , with r_0 as the evaluation target. Numerical baseline B7 implements the finite-library maximum-level ES and engine-wise regret competitors on the same library precisely to make that object-level split visible: $I(q)$ lifts the Huang-style objective away from the certificate target even when the two rules look similar on paper.

Remark 2.2 (Connection to the SPO loss). Conceptually, $r_q(x)$ is related to the SPO loss of [Elmachtoub and Grigas \(2022\)](#): both quantify the suboptimality gap from implementing a decision under a surrogate or predicted objective relative to an underlying evaluation. The parallel stops at the measurement principle. SPO is formulated for linear contextual cost vectors and training-stage surrogate minimization, whereas our regret is defined over coherent tail-risk functions (Expected Shortfall) and is later controlled through independent validation rather than through an SPO+ training surrogate.

Definition 2.7 (Hybrid objective). For $\lambda \in [0, 1]$ and $\hat{\mathcal{Q}}_\delta \neq \emptyset$,

$$H_\lambda(x) = (1 - \lambda) r_{q^{\text{vpw}}}(x) + \lambda \sup_{q \in \hat{\mathcal{Q}}_\delta} r_q(x). \quad (2.34)$$

Let

$$\hat{x}_\lambda \in \arg \min_{x \in \mathcal{X}} H_\lambda(x), \quad \hat{V}_\lambda = \min_{x \in \mathcal{X}} H_\lambda(x). \quad (2.35)$$

Endpoints: $\lambda = 0$ is the validated training-score portfolio; $\lambda = 1$ is validation-only minimax regret; interior values are hybrid. We call λ a *training-to-robustness interpolation parameter*. Official numerical cells lock $\lambda = 0.5$; the remaining grid is a sensitivity path, not a decision rule.

If $\hat{\mathcal{Q}}_\delta = \emptyset$, the method returns a pre-specified rejection/fallback decision; the certificates of Section 3 do not apply.

3. Certificates and computation

The construction of Section 2 is now turned into a high-probability bound on true ES regret and into an implementable linear program. We first lift the finite audit to the whole simplex by a Lipschitz cover, then transfer the bound to r_q and H_λ , and finally record the decision-level cost of projecting the training anchor. The computational subsection explains why the inner supremum is an extreme-point program rather than a dual of the linear form $c(x)^\top q$.

3.1. Lipschitz cover and regret certificates

Define the coverage radius

$$h(\mathcal{B}, \mathcal{X}) = \sup_{x \in \mathcal{X}} \min_{b \in \mathcal{B}} \|x - b\|_1. \quad (3.1)$$

Lemma 3.1 (ES Lipschitz). *Under Assumption 2.2, the deployment risk f_0 and every implemented engine risk f_m (hence every aggregate f_q) satisfy*

$$|f(x) - f(y)| \leq \bar{R} \|x - y\|_1. \quad (3.2)$$

Proof. Under P_0 , holding-period loss satisfies

$$L(x, R) - L(y, R) = (y - x)^\top R \leq \bar{R}_0 \|x - y\|_1 \quad (3.3)$$

almost surely, and likewise with x and y exchanged, so

$$\|L(x, \cdot) - L(y, \cdot)\|_\infty \leq \bar{R}_0 \|x - y\|_1. \quad (3.4)$$

Each implemented engine is pre-truncated: $\|R_{ms}\|_\infty \leq \bar{R}_{\text{eng}}$ for every scenario, hence the same ℓ_∞ comparison yields Lipschitz constant $\bar{R}_{\text{eng}}$ for every engine loss. Set $\bar{R} = \max\{\bar{R}_0, \bar{R}_{\text{eng}}\}$. Expected Shortfall is monotone and cash-invariant: if $L \leq L' + c$ almost surely then $\text{ES}_\alpha(L) \leq \text{ES}_\alpha(L') + c$. Applying this to $L(x, \cdot) \leq L(y, \cdot) + \bar{R} \|x - y\|_1$ and the symmetric inequality gives

$$|f_0(x) - f_0(y)| \leq \bar{R} \|x - y\|_1. \quad (3.5)$$

The identical argument applies to each truncated engine map f_m , and therefore to every convex combination $f_q = \sum_m q_m f_m$:

$$|f_q(x) - f_q(y)| \leq \sum_m q_m |f_m(x) - f_m(y)| \leq \bar{R} \|x - y\|_1. \quad (3.6)$$

(Equivalently, in the Rockafellar–Uryasev epigraph, a uniform ℓ_∞ shift of the loss shifts every positive-part expectation by at most the same constant, so the infimum moves by at most that constant.) $\square$

Theorem 3.1 (Audit-to-domain bound). *On the containment event, for every $q \in \hat{\mathcal{Q}}_\delta$,*

$$d_{\mathcal{X}}(f_q, f_0) \leq \Gamma_T := \delta + 4e_T + 4\bar{R}h(\mathcal{B}, \mathcal{X}). \quad (3.7)$$

Proof. Fix $q \in \hat{\mathcal{Q}}_\delta$ on the containment event $\{\max_j |\hat{\theta}_j - \theta_j| \leq e_T\}$, and write $\Delta := f_q - f_0$. Theorem 2.4 yields $q \in \mathcal{Q}_{\delta+4e_T}$, so the finite audit controls the oscillation of Δ on $\mathcal{B}$:

$$\text{osc}_{\mathcal{B}}(\Delta) = d_{\mathcal{B}}(q; \theta) \leq \delta + 4e_T. \quad (3.8)$$

It remains to lift that control from $\mathcal{B}$ to $\mathcal{X}$. For arbitrary $x \in \mathcal{X}$ choose a nearest audit point $b(x) \in \mathcal{B}$ in the ℓ_1 metric, so that $\|x - b(x)\|_1 \leq h(\mathcal{B}, \mathcal{X})$ by definition of the covering radius. Lemma 3.1 applied to f_q and to f_0 gives

$$|\Delta(x) - \Delta(b(x))| = |(f_q(x) - f_0(x)) - (f_q(b(x)) - f_0(b(x)))| \quad (3.9)$$

$$\leq |f_q(x) - f_q(b(x))| + |f_0(x) - f_0(b(x))| \quad (3.10)$$

$$\leq 2\bar{R}h(\mathcal{B}, \mathcal{X}). \quad (3.11)$$

Consequently every value of Δ on $\mathcal{X}$ lies within $2\bar{R}h(\mathcal{B}, \mathcal{X})$ of some value on $\mathcal{B}$, and therefore

$$\max_{x \in \mathcal{X}} \Delta(x) \leq \max_{b \in \mathcal{B}} \Delta(b) + 2\bar{R}h(\mathcal{B}, \mathcal{X}), \quad (3.12)$$

$$\min_{x \in \mathcal{X}} \Delta(x) \geq \min_{b \in \mathcal{B}} \Delta(b) - 2\bar{R}h(\mathcal{B}, \mathcal{X}). \quad (3.13)$$

Subtracting these two displays produces

$$\text{osc}_{\mathcal{X}}(\Delta) \leq \text{osc}_{\mathcal{B}}(\Delta) + 4\bar{R}h(\mathcal{B}, \mathcal{X}) \leq \delta + 4e_T + 4\bar{R}h(\mathcal{B}, \mathcal{X}), \quad (3.14)$$

which is (3.7). The three summands have distinct sources: model tolerance δ , simultaneous validation error $4e_T$, and Lipschitz mesh error $4\bar{R}h(\mathcal{B}, \mathcal{X})$. $\square$

Error budget. Equation (3.7) is an explicit three-term budget: model tolerance δ , simultaneous DKW-to-ES width $4e_T$, and Lipschitz mesh lift $4\bar{R}h(\mathcal{B}, \mathcal{X})$. Experiment D isolates the mesh term; Experiment E (and its refinement cell) tracks the same budget as T grows: coverage stays one while the DKW term shrinks with sample size.

The ambiguity controlled by Γ_T is over engine mixes q , not over the return law P . Wasserstein DRO on distributions (Mohajerin Esfahani and Kuhn, 2018; Van Parys et al., 2021) and moment ambiguity sets (Delage and Ye, 2010) enlarge a ball around an estimated P (or its moments) and optimize a worst-case risk inside that ball. Here the engines f_m are fixed implemented maps; validation only restricts which convex combinations remain admissible, and Γ_T transfers that finite-audit control into a uniform bound on aggregate-risk regret. The two literatures therefore certify different objects even when both employ Expected Shortfall as the risk functional.

The constant $4\bar{R}h(\mathcal{B}, \mathcal{X})$ does not vanish for a *fixed* vertex library: if $\mathcal{B}$ is exactly the simplex vertices then $h = 2(1 - 1/n)$. That is a property of the audit mesh, not of the regret transfer. If the mesh is refined so that $h(\mathcal{B}, \mathcal{X}) \rightarrow 0$ and $e_T \rightarrow 0$, then $\Gamma_T \rightarrow \delta$; for a pre-specified $\delta = 0$ the Lipschitz certificate vanishes.

Lemma 3.2 (Enriched three-asset covering). *If $n = 3$ and $\mathcal{B}$ contains the three vertices, the three edge midpoints, and the barycenter $1/3$, then*

$$h(\mathcal{B}, \mathcal{X}) \leq 2/3. \quad (3.15)$$

For $n \geq 2$, the same construction yields the coarser certified bound

$$h(\mathcal{B}, \mathcal{X}) \leq \min(2(1 - 1/n), 1). \quad (3.16)$$

The proof is geometric and is deferred to Appendix [Appendix A](#).

Remark 3.1 (Convex interpolation remainder). *Expected Shortfall is convex, so a supporting-plane remainder can in principle replace part of the Lipschitz diameter when first-order tail sensitivities are close on $\mathcal{B}$. If $\mathcal{X} = \text{conv}(\mathcal{B})$ and $\Delta = f_q - f_0$, Jensen gaps*

$$\eta_f(x) = \sum_j \lambda_j f(b_j) - f(x) \quad (3.17)$$

yield

$$\text{osc}_{\mathcal{X}}(\Delta) \leq \text{osc}_{\mathcal{B}}(\Delta) + 2\bar{\eta} \quad (3.18)$$

for a uniform gap $\bar{\eta}$. A further subgradient comparison would introduce constants δ_G and κ that are not internally certified by the audit sample. We therefore do not treat those constants as a formal remainder, and the official certificate remains the Lipschitz cover of Theorem [3.1](#).

Lemma 3.3 (Uniform regret control). *On the same event, for all $x \in \mathcal{X}$ and $q \in \hat{\mathcal{Q}}_\delta$,*

$$|r_q(x) - r_0(x)| \leq \Gamma_T. \quad (3.19)$$

Moreover,

$$\left| \sup_{q \in \hat{\mathcal{Q}}_\delta} r_q(x) - r_0(x) \right| \leq \Gamma_T. \quad (3.20)$$

Proof. The first claim is Lemma [2.1](#) together with Theorem [3.1](#):

$$|r_q(x) - r_0(x)| \leq d_{\mathcal{X}}(f_q, f_0) \leq \Gamma_T. \quad (3.21)$$

For the supremum, the same bound gives

$$\sup_{q \in \hat{\mathcal{Q}}_\delta} r_q(x) \leq r_0(x) + \Gamma_T. \quad (3.22)$$

If $\hat{\mathcal{Q}}_\delta = \emptyset$, the hybrid rule follows the rejection branch and the lemma is not invoked. If the set is nonempty then $q^{\text{vpw}} \in \hat{\mathcal{Q}}_\delta$, so

$$\sup_{q \in \hat{\mathcal{Q}}_\delta} r_q(x) \geq r_{q^{\text{vpw}}}(x) \geq r_0(x) - \Gamma_T. \quad (3.23)$$

Hence $\left| \sup_q r_q(x) - r_0(x) \right| \leq \Gamma_T$. □

Theorem 3.2 (Hybrid regret approximation). *Suppose $\widehat{\mathcal{Q}}_\delta \neq \emptyset$. On the simultaneous validation event, for every $x \in \mathcal{X}$ and $\lambda \in [0, 1]$,*

$$|H_\lambda(x) - r_0(x)| \leq \Gamma_T. \quad (3.24)$$

Proof. Lemma 3.3 places both building blocks of the hybrid objective in a Γ_T -neighbourhood of true regret: for every $x \in \mathcal{X}$,

$$r_{q^{\text{vpw}}}(x), \sup_{q \in \widehat{\mathcal{Q}}_\delta} r_q(x) \in [r_0(x) - \Gamma_T, r_0(x) + \Gamma_T]. \quad (3.25)$$

The hybrid score is the convex combination

$$H_\lambda(x) = (1 - \lambda) r_{q^{\text{vpw}}}(x) + \lambda \sup_{q \in \widehat{\mathcal{Q}}_\delta} r_q(x). \quad (3.26)$$

Any convex combination of two points in an interval remains in that interval, so $H_\lambda(x) \in [r_0(x) - \Gamma_T, r_0(x) + \Gamma_T]$, which is exactly (3.24). The argument is uniform in $\lambda \in [0, 1]$: the same Γ_T controls the validated anchor ($\lambda = 0$), pure minimax ($\lambda = 1$), and every locked interpolation in between. $\square$

Theorem 3.3 (A posteriori certificate). *If*

$$H_\lambda(\widehat{x}_\lambda) \leq \widehat{V}_\lambda + \xi \quad (3.27)$$

and $\widehat{\mathcal{Q}}_\delta \neq \emptyset$, then on the good event

$$r_0(\widehat{x}_\lambda) \leq \widehat{V}_\lambda + \xi + \Gamma_T. \quad (3.28)$$

Proof. On the good event $\mathcal{E}_T$, Theorem 3.2 yields the pointwise comparison

$$r_0(\widehat{x}_\lambda) \leq H_\lambda(\widehat{x}_\lambda) + \Gamma_T. \quad (3.29)$$

The numerical solver is assumed only to return a ξ -optimal hybrid decision,

$$H_\lambda(\widehat{x}_\lambda) \leq \widehat{V}_\lambda + \xi = \min_{y \in \mathcal{X}} H_\lambda(y) + \xi. \quad (3.30)$$

Chaining the two displays gives

$$r_0(\widehat{x}_\lambda) \leq \widehat{V}_\lambda + \xi + \Gamma_T, \quad (3.31)$$

which is the a posteriori certificate (3.28). No additional concentration is required at this step: all statistical uncertainty has already been absorbed into Γ_T through Theorem 3.1. $\square$

Corollary 3.4 (A priori certificate). *Under the same conditions,*

$$r_0(\hat{x}_\lambda) \leq 2\Gamma_T + \xi. \quad (3.32)$$

Proof. Let $x_0^* \in \arg \min_{y \in \mathcal{X}} r_0(y)$, so $r_0(x_0^*) = 0$. Then

$$\hat{V}_\lambda = \min_x H_\lambda(x) \leq H_\lambda(x_0^*). \quad (3.33)$$

Theorem 3.2 gives

$$H_\lambda(x_0^*) \leq r_0(x_0^*) + \Gamma_T = \Gamma_T. \quad (3.34)$$

Combine with Theorem 3.3. $\square$

Theorem 3.5 (Validation-correction cost). *Let*

$$\Delta_{\text{val}} = \|q^{\text{vpw}} - q^0\|_1 \quad (3.35)$$

and

$$D = \max_{1 \leq m \leq M} \left(\max_{x \in \mathcal{X}} f_m(x) - \min_{x \in \mathcal{X}} f_m(x) \right). \quad (3.36)$$

Then for all $x \in \mathcal{X}$,

$$|r_{q^{\text{vpw}}}(x) - r_{q^0}(x)| \leq D\Delta_{\text{val}}. \quad (3.37)$$

Proof. Write $\Delta q = q^{\text{vpw}} - q^0$. For any constants c_m ,

$$f_{q^{\text{vpw}}}(x) - f_{q^0}(x) = \sum_{m=1}^M \Delta q_m (f_m(x) - c_m) + \sum_{m=1}^M \Delta q_m c_m. \quad (3.38)$$

Choose $c_m = \min_{y \in \mathcal{X}} f_m(y)$. Then $0 \leq f_m(x) - c_m \leq D$ and the second sum does not depend on x . Hence

$$|(f_{q^{\text{vpw}}} - f_{q^0})(x) - (f_{q^{\text{vpw}}} - f_{q^0})(y)| \leq D \|\Delta q\|_1 = D\Delta_{\text{val}}, \quad (3.39)$$

so $d_{\mathcal{X}}(f_{q^{\text{vpw}}}, f_{q^0}) \leq D\Delta_{\text{val}}$. Lemma 2.1 transfers the bound to regrets. $\square$

The quantity $D\Delta_{\text{val}}$ is the decision-level price of enforcing validation. If $q^0 \in \hat{\mathcal{Q}}_\delta$, then $\Delta_{\text{val}} = 0$.

Corollary 3.6 (Anchor-adaptive validation cost). *If $\hat{\mathcal{Q}}_\delta \neq \emptyset$, then for every $x \in \mathcal{X}$*

$$|r_{q^{\text{vpw}}}(x) - r_{q^0}(x)| \leq D\sqrt{M} \text{dist}_2(q^0, \hat{\mathcal{Q}}_\delta), \quad (3.40)$$

where

$$\text{dist}_2(q^0, \hat{\mathcal{Q}}_\delta) := \min_{q \in \hat{\mathcal{Q}}_\delta} \|q - q^0\|_2 \quad (3.41)$$

is the Euclidean distance from the training-score anchor to the admissible set.

Proof. Definition 2.6 is Euclidean projection, so

$$\|q^{\text{vpw}} - q^0\|_2 = \text{dist}_2(q^0, \hat{\mathcal{Q}}_\delta). \quad (3.42)$$

Norm comparison on $\mathbb{R}^M$ gives

$$\Delta_{\text{val}} = \|q^{\text{vpw}} - q^0\|_1 \leq \sqrt{M} \|q^{\text{vpw}} - q^0\|_2. \quad (3.43)$$

Theorem 3.5 then yields the claim. $\square$

The bounds control H_λ once a minimizer is in hand. The remaining question is computational: how to evaluate the inner supremum of aggregate-risk regret over the polytope.

3.2. Computation

The set $\hat{\mathcal{Q}}_\delta$ admits the compact range formulation with variables (q, ℓ, u) :

$$\ell \leq (Aq - \hat{\theta})_j \leq u, \quad u - \ell \leq \delta + 2e_T, \quad q \in \Delta_M. \quad (3.44)$$

Definition 2.6 is then a convex QP.

Proposition 3.1 (Extreme-point reduction of aggregate-risk regret). *Suppose $\hat{\mathcal{Q}}_\delta \neq \emptyset$. The map $q \mapsto v(q) = \min_{y \in \mathcal{X}} f_q(y)$ is concave on Δ_M (pointwise infimum of maps linear in q). Consequently, for each fixed $x \in \mathcal{X}$, the aggregate-risk regret $r_q(x) = f_q(x) - v(q)$ is convex in q . A convex function on a nonempty polytope attains its maximum at an extreme point, so*

$$\sup_{q \in \hat{\mathcal{Q}}_\delta} r_q(x) = \max_{v^k \in \text{ext}(\hat{\mathcal{Q}}_\delta)} [f_{v^k}(x) - \bar{v}_k], \quad \bar{v}_k := \min_{y \in \mathcal{X}} f_{v^k}(y). \quad (3.45)$$

Once $\text{ext}(\hat{\mathcal{Q}}_\delta)$ and the offline oracles $\bar{v}_k, \bar{v}_{\text{vpw}}$ are computed, the hybrid objective H_λ is an epigraph-linear function of (x, t) after each finite-scenario ES is expanded in Rockafellar–Uryasev form, and the hybrid rule is a linear program. We do not claim that $|\text{ext}(\hat{\mathcal{Q}}_\delta)|$ is polynomial in the number of audit inequalities.

Proof. Linearity of $q \mapsto f_q(x)$ is immediate from $f_q = \sum_m q_m f_m$. Concavity of $v(\cdot)$ follows because

$$v(q) = \min_{y \in \mathcal{X}} \sum_m q_m f_m(y) \quad (3.46)$$

is an infimum of affine functions of q . Hence $r_q(x) = f_q(x) - v(q)$ is convex in q . On a nonempty compact convex polytope the maximum of a continuous convex function is attained at an extreme point, which is (3.45). Substituting the finite list $\{v^k\}$ into the epigraph form of H_λ and linearizing each ES block (Rockafellar and Uryasev, 2002) yields an LP in (x, t) and the RU auxiliaries. $\square$

Dualizing the linear form

$$c(x)^\top q = \sum_m q_m (f_m(x) - v_m) \quad (3.47)$$

would instead recover

$$\sup_q \tilde{r}_q(x) = \sup_q (r_q(x) + I(q)), \quad (3.48)$$

which is engine-wise regret, not the hybrid target of Proposition 3.1.

The official solver therefore has two stages.

- (i) Offline. Compute $\text{ext}(\hat{\mathcal{Q}}_\delta)$ and, for each extreme point v^k , solve

$$\bar{v}_k = \min_y f_{v^k}(y). \quad (3.49)$$

In the same pass, solve the validated-anchor oracle

$$\bar{v}_{\text{vpw}} := \min_{y \in \mathcal{X}} f_{q^{\text{vpw}}}(y). \quad (3.50)$$

- (ii) Online. Introduce an epigraph variable t for

$$\max_k (f_{v^k}(x) - \bar{v}_k) \quad (3.51)$$

and expand each finite-scenario ES with a Rockafellar–Uryasev block. With both families of oracles precomputed, the objective

$$(1 - \lambda) [f_{q^{\text{vpw}}}(x) - \bar{v}_{\text{vpw}}] + \lambda t \quad (3.52)$$

has only known scalar offsets, so the program is a linear program in (x, t) and the RU auxiliaries.

Each offline oracle and the online epigraph block are Rockafellar–Uryasev linearizations of Expected Shortfall (Rockafellar and Uryasev, 2002), as already used in Proposition 3.1. Absolute and relative robust optimization with discrete ambiguity sets routinely reduce to vertex programs of this type; the survey of Aissi et al. (2009) places minimax-regret formulations in that computational family. The contribution here is not a new LP template, but the identification of the convex objective r_q (rather than the linear surrogate $\tilde{r}_q$) together with the validation polytope that supplies the vertices. For the locked library size $M \leq 4$ the enumeration is exact and fast.

Proposition 3.2 (Engine-wise dual as a surrogate). *Write*

$$c(x) = (f_1(x) - v_1, \dots, f_M(x) - v_M)^\top. \quad (3.53)$$

The linear program

$$\begin{aligned}
& \max_{q, \ell, u} && c(x)^\top q \\
& \text{s.t.} && q \in \Delta_M, \\
& && Aq - u\mathbf{1} \leq \widehat{\theta}, \\
& && -Aq + \ell\mathbf{1} \leq -\widehat{\theta}, \\
& && u - \ell \leq \delta + 2e_T
\end{aligned} \tag{3.54}$$

equals

$$\sup_{q \in \widehat{\mathcal{Q}}_\delta} \widetilde{r}_q(x) = \sup_{q \in \widehat{\mathcal{Q}}_\delta} [r_q(x) + I(q)]. \tag{3.55}$$

If $I(q) = 0$ for every $q \in \widehat{\mathcal{Q}}_\delta$, then it also equals

$$\sup_{q \in \widehat{\mathcal{Q}}_\delta} r_q(x). \tag{3.56}$$

A common minimizer

$$x^\dagger \in \bigcap_{m=1}^M \arg \min_{x \in \mathcal{X}} f_m(x) \tag{3.57}$$

is sufficient for $I(q) \equiv 0$ on the simplex. The converse is not asserted: for a particular x , equality of the two suprema may occur even when $I(q) > 0$ at other admissible weights.

Proof. By definition

$$c(x)^\top q = \sum_m q_m (f_m(x) - v_m) = r_q(x) + I(q) = \widetilde{r}_q(x). \tag{3.58}$$

Program (3.54) is therefore

$$\sup_{q \in \widehat{\mathcal{Q}}_\delta} \widetilde{r}_q(x) = \sup_{q \in \widehat{\mathcal{Q}}_\delta} [r_q(x) + I(q)]. \tag{3.59}$$

If $I(q) = 0$ for every $q \in \widehat{\mathcal{Q}}_\delta$, the two integrands coincide on the feasible set, so the program also equals $\sup_{q \in \widehat{\mathcal{Q}}_\delta} r_q(x)$. A common minimizer

$$x^\dagger \in \bigcap_{m=1}^M \arg \min_{x \in \mathcal{X}} f_m(x) \tag{3.60}$$

gives $v(q) = \sum_m q_m v_m$ for every $q \in \Delta_M$, hence $I \equiv 0$ on the simplex and a fortiori on $\widehat{\mathcal{Q}}_\delta$. The converse is not asserted: vanishing of I on $\widehat{\mathcal{Q}}_\delta$ is sufficient, not necessary, for the two suprema at a fixed x to coincide, since $I(q) > 0$ at some admissible weights need not raise the attained supremum of $r_q + I$. $\square$

If $\widehat{\mathcal{Q}}_\delta = \emptyset$, the method returns the pre-specified rejection/fallback of Theorem 2.5. Section 4 checks the mechanisms just stated: shift invariance, inactive validation under a conservative radius, certificate width, and the distinction from return-law DRO.

4. Numerical experiments

The numerical section checks mechanisms from Section 3 rather than ranking methods under a correctly specified Gaussian law. Contribution item 4 already declines to claim outperformance of $1/N$: when engines are well specified and the locked DKW radius is inert, equal weights are a historically strong baseline (DeMiguel et al., 2009), and the official certificate Γ_T remains conservative by construction. What the locked cells must show is that (i) additive level errors need not change decisions (A; Online Supplement table/figure); (ii) when the official radius is large, validation is inert and the hybrid tracks engine-wise regret (B); (iii) the same radius theory predicts when rejection activates (C; Online Supplement); (iv) audit mesh enters Γ_T through h and the union bound (D); (v) the a posteriori certificate covers true regret with an explicit width anatomy (E); (vi) when $1/N$ is a weak ES rule on an asymmetric cash simplex, the hybrid preserves a cash-heavy portfolio (B').

The protocol is locked before inspecting outputs; all optimization uses SciPy HiGHS; the official hybrid rule enumerates extreme points of $\widehat{\mathcal{Q}}_\delta$ as in Section 3.2. True ES regret is exact on a finite clipped DGP cloud ($\bar{R}_0 = 0.20$). Parentheses in tables are standard errors across replications. Full protocols, DGPs, and robustness cells are in the Online Supplement.¹

Default locked cells: $\alpha = 0.95$, $\beta = 0.05$, $\delta = 0.01$, official hybrid $\lambda = 0.5$, $n \in \{3, 6\}$, $N_{\text{tr}}, T \in \{100, 250, 500\}$, engine library of size at most $M = 4$. Experiment E uses 200 independent replications for $n = 3$ and 30 for $n = 6$.

Baselines: B1 empirical ES on the validation sample; B2 parametric ETO; B3 unvalidated training-score anchor q^0 ; B4 validation-only minimax ($\lambda = 1$); B5 validated training-score anchor ($\lambda = 0$); B6 hybrid at the locked $\lambda = 0.5$; B7 finite-library maximum-level ES and engine-wise regret baselines (Huang-inspired) on the same engines; B8 Wasserstein-CVaR; B9 $1/N$. An exploratory clustered Wasserstein-CVaR cash-tilt comparison is reported in the Online Supplement alongside Experiment B'; Pun et al. (2023) is cited as a neighboring return-law benchmark.

4.1. Experiment A: additive shifts are not decision errors

Two-asset exact engines shifted by a common loss constant $c \in \{0.01, 0.04, 0.08\}$ keep absolute ES error equal to $|c|$, yet $d_{\mathcal{X}} = 0$ to machine precision and $t^* = 0.25$ unchanged (Online Supplement).

¹Online Supplement: locked protocol; Experiment A table/figure; Experiment B regret figure; Experiment B' cash figure; Experiment C tight-radius table; Experiment D audit table; Experiment E width and refinement path; λ -grid coverage; B8 Wasserstein path; relative-robust versus aggregate-risk note.

This supports the shift-invariant discrepancy $d_{\mathcal{X}}$ and Lemma 2.1: level scorecards can report a large “error” that does not change the decision.

4.2. Experiment B: sampling error under a correct Gaussian law

The population is a clipped Gaussian finite support with a correctly specified Gaussian engine. Table II reports mean true ES regret over 100 draws ($n = 3$). B1 is an information-rich validation plug-in, not a training-only competitor to B3–B6. B3 and B5 coincide: the DKW-to-ES radius leaves $\hat{\mathcal{Q}}_{\delta} = \Delta_M$ and $\Delta_{\text{val}} = 0$ (Theorem 2.4). In that inert-projection regime the locked hybrid (B6, $\lambda = 0.5$) tracks engine-wise regret / minimax (B4, B7 engine-wise) and improves on ETO (B2) and maximum-level ES (B7 level), consistent with optimizing r_q (Remark 2.1). Equal-weight $1/N$ (B9) remains strong under this law (DeMiguel et al., 2009); that is not a claim the hybrid is designed to overturn. The companion regret figure is in the Online Supplement.

Table II: Experiment B: mean true ES regret (SE), Gaussian DGP, $n = 3$, 100 replications. B6 is the locked $\lambda = 0.5$ rule. B3=B5 because the locked DKW radius leaves $\hat{\mathcal{Q}}_{\delta} = \Delta_M$.

Method	$N_{\text{tr}} = 100$	250	500
B1 validation ES	0.00085 (0.00012)	0.00072 (0.00008)	0.00112 (0.00013)
B2 parametric ETO	0.00432 (0.00038)	0.00438 (0.00035)	0.00314 (0.00029)
B3 unvalidated q^0	0.00299 (0.00027)	0.00309 (0.00031)	0.00230 (0.00029)
B4 minimax $\lambda = 1$	0.00219 (0.00016)	0.00174 (0.00014)	0.00161 (0.00016)
B5 validated q^{vpw}	0.00299 (0.00027)	0.00309 (0.00031)	0.00230 (0.00029)
B6 hybrid $\lambda = 0.5$	0.00216 (0.00015)	0.00158 (0.00014)	0.00150 (0.00016)
B7 engine-wise regret	0.00219 (0.00016)	0.00174 (0.00014)	0.00161 (0.00016)
B7 maximum-level ES	0.00408 (0.00031)	0.00366 (0.00033)	0.00314 (0.00026)
B9 $1/N$	0.00138 (0.00012)	0.00153 (0.00013)	0.00146 (0.00012)

On the long-only simplex, positive training-fold Wasserstein radii for B8 collapse toward the same portfolio as $1/N$ because the dual-norm penalty dominates.²

4.3. Experiment B': when equal weight is weak

Experiment B is silent when $1/N$ is a poor ES rule. On a locked four-asset equity/bond/cash simplex (Table III), mean true ES regret for $1/N$ is about 0.0416 at $\varepsilon = 0$ and 0.0412 at $\varepsilon = 0.01$. At $\varepsilon = 0$, single-ball and clustered Wasserstein-CVaR rules concentrate on cash (mean weight 0.96) with $r_0 \approx 0.00011$; the locked hybrid tilts to cash (mean weight 0.90) with $r_0 \approx 0.00059$. At $\varepsilon = 0.01$, the W-rules lose cash (mean weight 0.30) and inflate regret to about 0.030–0.032, while the hybrid keeps cash near 0.88 and $r_0 \approx 0.00068$. Wasserstein hedges return-law balls; the hybrid hedges admissible ES aggregations. Pun et al. (2023) is cited as a neighboring return-law benchmark. The cash-weight figure is in the Online Supplement.

²Radius-split tables, the weight path, and the Experiment B regret figure are in the Online Supplement.

Table III: Experiment B': asymmetric stock/bond/cash simplex, 40 replications. Mean true ES regret (SE) and mean cash weight.

ε	$1/N$	single W	clustered W	hybrid	W cash	hybrid cash
0	0.0416 (0.0007)	0.00011 (0.00002)	0.00011 (0.00002)	0.00059 (0.00007)	0.96	0.90
0.01	0.0412 (0.0007)	0.0304 (0.0017)	0.0325 (0.0016)	0.00068 (0.00009)	0.30	0.88

4.4. Experiment C: validation projection under locked and tight radii

Under the locked DKW-to-ES radius, projecting a misspecified prior onto $\hat{\mathcal{Q}}_\delta$ leaves $\hat{\mathcal{Q}}_\delta = \Delta_M$ in every official replication (empty-set rate 0, $\Delta_{\text{val}} = 0$), as predicted by Theorems 2.4–2.5 when e_T is large relative to engine disagreement on $\mathcal{B}$. A pre-specified tight-radius diagnostic ($e_T = 0.005$, $\delta = 0$), not inserted into Γ_T , activates rejection and displacement (empty-set rate 0.525, mean $\Delta_{\text{val}} = 0.64$; Online Supplement).

4.5. Experiment D: audit geometry and coverage

At $n = 3$, vertices / conventional / hybrid audits share the certified covering radius $4/3$ whenever vertices are included, so extra interior audits do not tighten the Lipschitz floor in Γ_T ; they enlarge the union-bound factor in e_T and inflate Γ_T (vertices $\Gamma_T = 6.028$, conventional 6.174, hybrid 6.511), while mean true regrets coincide because $\hat{\mathcal{Q}}_\delta = \Delta_M$. This matches the mesh term in Theorem 3.1. Diagnostic versus certified h values are tabulated in the Online Supplement.

4.6. Experiment E: certificate coverage (mandatory)

For each replication we draw training and independent validation samples from a known regime-mixture DGP, form q^0 and $\hat{\mathcal{Q}}_\delta$, solve the hybrid rule, and test (3.28) with $\xi = 0$. Table IV uses $n = 3$, the hybrid library with $J = 8$, and 200 replications. Coverage is 1 at every (T, λ) : the a posteriori certificate of Theorem 3.3 holds in every locked draw. Γ_T is conservative by design (from 6.51 at $T = 100$ to 3.51 at $T = 500$), while true regrets stay near 0.002. The experiment therefore supports *validity* of the guarantee, not economic sharpness of its width. Across λ , the locked hybrid (B6) sits between the validated anchor (B5) and pure minimax (B4), as the interpolation H_λ intends.

The official width decomposes as in (3.7): at $T \in \{100, 250, 500\}$ the DKW term is $4e_T \in \{5.435, 3.437, 2.430\}$ while the mesh floor $4\bar{R}h = 1.067$ is flat for this vertex-rich library, matching Theorem 3.1. A sharper ES radius (Prashanth et al., 2020) would be needed before the certificate is economically informative; the present cell documents the anatomy of the official bound rather than claiming that anatomy is optimal. The $n = 6$ cell (30 replications) repeats the same validity pattern: coverage 1, with conservative $\Gamma_T \in [3.77, 6.78]$. The non-vanishing width is the mesh-plus-DKW issue of Section 3, not a coverage failure. Width decomposition and anatomy displays are in the Online Supplement.

Table IV: Experiment E, $n = 3$, hybrid library $J = 8$, 200 replications. Coverage of $r_0 \leq \hat{V} + \Gamma_T$. Mean r_0 reported with SE. Coverage equals one; Γ_T is a valid but conservative envelope.

T	method	λ	coverage	mean Γ_T	mean r_0
100	B5	0.00	1.00	6.51	0.00299 (0.00022)
100	B6	0.50	1.00	6.51	0.00202 (0.00012)
100	B4	1.00	1.00	6.51	0.00188 (0.00012)
250	B5	0.00	1.00	4.51	0.00301 (0.00022)
250	B6	0.50	1.00	4.51	0.00203 (0.00014)
250	B4	1.00	1.00	4.51	0.00191 (0.00012)
500	B5	0.00	1.00	3.51	0.00316 (0.00021)
500	B6	0.50	1.00	3.51	0.00228 (0.00014)
500	B4	1.00	1.00	3.51	0.00219 (0.00013)

A pre-specified refinement cell keeps the official DKW-to-ES radius and extends T through 2000: mean Γ_T falls from 6.03 to 2.18 and coverage stays 1.³

5. Conclusion

This paper proposes a two-sample hybrid rule for Expected Shortfall portfolio choice under a finite library of competing tail-risk engines. Engines are scored on the training fold, relative Expected Shortfall rankings are audited on an independent validation sample, the training mix is projected onto the admissible set, and a locked interpolation weight balances that validated mix against worst-case admissible regret. An empty admissible set is treated as certified rejection of the library rather than a solver failure. For a small finite library the rule is a linear program obtained from extreme points of the admissible set, and the accompanying certificate separates model tolerance, validation error, coverage error, and optimization error. The locked numerical cells are consistent with that theory: additive level errors need not change decisions; inert DKW radii leave projection inactive and make the hybrid track engine-wise regret; audit mesh and DKW width dominate certificate width; and the a posteriori bound covers true regret at the advertised level.

Future work may develop sharper Expected Shortfall concentration bounds, dependence-robust validation protocols, adaptive audit designs that tighten the mesh term without sacrificing certification, and scalable algorithms that avoid explicit extreme-point enumeration for larger engine libraries.

³Full λ -grid coverage, width decomposition/figure, refinement path, Experiment A/C/D tables, B8 path, and asymmetric cash displays: Online Supplement.

Appendix A. Proofs of auxiliary results

Throughout, $\mathcal{X}$ is nonempty, compact, and convex, and each risk function is finite and continuous on $\mathcal{X}$. Write

$$\text{osc}_S(\Delta) = \max_{z \in S} \Delta(z) - \min_{z \in S} \Delta(z), \quad \text{rng}(v) = \max_j v_j - \min_j v_j. \quad (\text{A.1})$$

Minimizers $x_f^* \in \arg \min_{y \in \mathcal{X}} f(y)$ exist by compactness.

Appendix A.1. Simultaneous DKW-to-ES radius

Lemma Appendix A.1 (Simultaneous DKW-to-ES radius). *Fix a $\mathcal{G}_{\text{tr}}$ -measurable audit library of size J . The official radius is*

$$e_T(\alpha, \beta, \bar{R}, T, J) = \frac{2\bar{R}}{1-\alpha} \sqrt{\frac{\log(2J/\beta)}{2T}}. \quad (\text{A.2})$$

It satisfies

$$\Pr\left(\max_{j \in [J]} |\hat{\theta}_j - \theta_j| \leq e_T(\alpha, \beta, \bar{R}, T, J) \mid \mathcal{G}_{\text{tr}}\right) \geq 1 - \beta. \quad (\text{A.3})$$

Proof. Work conditionally on $\mathcal{G}_{\text{tr}}$, so that the audit library $\mathcal{B} = \{b_1, \dots, b_J\}$ is fixed. For each $j \in [J]$ write $L_{jt} = L(b_j, Y_t^{\text{val}})$ for the validation losses; Assumption 2.2 gives $L_{jt} \in [-\bar{R}, \bar{R}]$ almost surely. Let F_j be the true CDF of $L(b_j, Y)$ under P_0 and let $\hat{F}_j$ be the empirical CDF from the T validation draws.

Step 1 (DKW band). Massart's inequality (Massart, 1990) yields, for every $\varepsilon > 0$,

$$\Pr\left(\|\hat{F}_j - F_j\|_\infty > \varepsilon \mid \mathcal{G}_{\text{tr}}\right) \leq 2e^{-2T\varepsilon^2}. \quad (\text{A.4})$$

Choosing $\varepsilon_T = \sqrt{\log(2J/\beta)/(2T)}$ makes the right-hand side at most β/J .

Step 2 (from Kolmogorov distance to ES). Fix j and drop the index. The Rockafellar–Uryasev representation reads

$$\text{ES}_\alpha(L) = \inf_{\nu \in \mathbb{R}} \phi(\nu), \quad \phi(\nu) = \nu + \frac{1}{1-\alpha} \mathbb{E}[L - \nu]_+, \quad (\text{A.5})$$

and likewise $\widehat{\text{ES}}_\alpha(L) = \inf_{\nu} \widehat{\phi}(\nu)$ with the empirical expectation. For any fixed ν ,

$$\mathbb{E}[L - \nu]_+ = \int_{-\bar{R}}^{\bar{R}} [t - \nu]_+ dF(t) = \int_{\nu}^{\bar{R}} (1 - F(t)) dt, \quad (\text{A.6})$$

and the same identity holds with $\widehat{F}$. Therefore

$$|\mathbb{E}[L - \nu]_+ - \widehat{\mathbb{E}}[L - \nu]_+| = \left| \int_{\nu}^{\bar{R}} (F(t) - \widehat{F}(t)) dt \right| \quad (\text{A.7})$$

$$\leq \int_{-\bar{R}}^{\bar{R}} \|\widehat{F} - F\|_{\infty} dt = 2\bar{R} \|\widehat{F} - F\|_{\infty}. \quad (\text{A.8})$$

Hence $|\widehat{\phi}(\nu) - \phi(\nu)| \leq 2\bar{R} \|\widehat{F} - F\|_{\infty} / (1 - \alpha)$ uniformly in ν , which implies

$$|\widehat{\text{ES}}_{\alpha}(L) - \text{ES}_{\alpha}(L)| \leq \frac{2\bar{R}}{1 - \alpha} \|\widehat{F} - F\|_{\infty}. \quad (\text{A.9})$$

On the DKW event of Step 1 this is at most $e_T = (2\bar{R}/(1 - \alpha))\varepsilon_T$.

Step 3 (union bound). The J validation losses are dependent because they are evaluated on the same returns, but the audits themselves are $\mathcal{G}_{\text{tr}}$ -measurable. A union bound therefore yields

$$\Pr\left(\max_j |\widehat{\theta}_j - \theta_j| > e_T \mid \mathcal{G}_{\text{tr}}\right) \leq \sum_{j=1}^J \frac{\beta}{J} = \beta, \quad (\text{A.10})$$

which is (A.3). $\square$

Appendix A.2. Enriched covering radius

Proof of Lemma 3.2. Identify $\mathcal{X} = \{x \in \mathbb{R}_+^n : \sum_i x_i = 1\}$. For the standard basis vector e_i ,

$$\|x - e_i\|_1 = 2(1 - x_i), \quad (\text{A.11})$$

so the vertex-only covering radius satisfies $h(\{e_1, \dots, e_n\}, \mathcal{X}) = 2(1 - 1/n)$, attained at the barycenter.

Now suppose $n = 3$ and $\mathcal{B}$ contains the three vertices, the three edge midpoints $m_{ij} = (e_i + e_j)/2$, and the barycenter $\bar{e} = (1/3, 1/3, 1/3)$. These seven points partition $\mathcal{X}$ into six congruent right triangles. By symmetry it is enough to bound the covering radius of one triangle, say $T = \text{conv}\{e_1, m_{12}, m_{13}\}$. For $x \in T$ one has $x_1 \geq 1/2$, hence $\|x - e_1\|_1 = 2(1 - x_1) \leq 1$. A direct computation at the interior peak $x^* = (2/3, 1/6, 1/6)$ gives

$$\|x^* - e_1\|_1 = \|x^* - m_{12}\|_1 = \|x^* - \bar{e}\|_1 = 2/3, \quad (\text{A.12})$$

and every other point of T is at most this far from $\{e_1, m_{12}, m_{13}, \bar{e}\}$. Thus $h(\mathcal{B}, \mathcal{X}) \leq 2/3$.

For general $n \geq 2$, the same coordinate argument shows that whenever some $x_i \geq 1/2$ one has $\|x - e_i\|_1 \leq 1$. If all coordinates are strictly less than $1/2$, then $n \geq 3$ and x lies in the convex hull of the barycenter and the edge midpoints of the face where the large coordinates live; the resulting ℓ_1 radius is at most 1. Intersecting with the vertex-only bound yields

$$h(\mathcal{B}, \mathcal{X}) \leq \min(2(1 - 1/n), 1).$$

□

References

- Acerbi, C., Tasche, D., 2002. On the coherence of expected shortfall. *Journal of Banking and Finance* 26, 1487–1503. doi:[10.1016/S0378-4266\(02\)00283-2](https://doi.org/10.1016/S0378-4266(02)00283-2).
- Aissi, H., Bazgan, C., Vanderpooten, D., 2009. Min-max and min-max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research* 197, 427–438. doi:[10.1016/j.ejor.2008.09.012](https://doi.org/10.1016/j.ejor.2008.09.012).
- Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1999. Coherent measures of risk. *Mathematical Finance* 9, 203–228. doi:[10.1111/1467-9965.00068](https://doi.org/10.1111/1467-9965.00068).
- Ban, G.Y., Rudin, C., 2019. The big data newsvendor: Practical insights from machine learning. *Operations Research* 67, 90–108. doi:[10.1287/opre.2018.1757](https://doi.org/10.1287/opre.2018.1757).
- Bertsimas, D., Gupta, V., Kallus, N., 2018. Data-driven robust optimization. *Mathematical Programming* 167, 235–292. doi:[10.1007/s10107-017-1125-8](https://doi.org/10.1007/s10107-017-1125-8).
- Bertsimas, D., Kallus, N., 2020. From predictive to prescriptive analytics. *Management Science* 66, 1025–1044. doi:[10.1287/mnsc.2018.3253](https://doi.org/10.1287/mnsc.2018.3253).
- Besbes, O., Mouchtaki, O., 2023. How big should your data really be? data-driven newsvendor: Learning one sample at a time. *Management Science* 69, 5848–5865. doi:[10.1287/mnsc.2023.4725](https://doi.org/10.1287/mnsc.2023.4725).
- Costa, G., Iyengar, G.N., 2023. Distributionally robust end-to-end portfolio construction. *Quantitative Finance* 23, 1465–1482. doi:[10.1080/14697688.2023.2236148](https://doi.org/10.1080/14697688.2023.2236148).
- Delage, E., Ye, Y., 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* 58, 595–612. doi:[10.1287/opre.1090.0741](https://doi.org/10.1287/opre.1090.0741).
- DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the $1/n$ portfolio strategy? *The Review of Financial Studies* 22, 1915–1953. doi:[10.1093/rfs/hhm075](https://doi.org/10.1093/rfs/hhm075).
- El Balghiti, O., Elmachtoub, A.N., Grigas, P., Tewari, A., 2019. Generalization bounds in the predict-then-optimize framework, in: *Advances in Neural Information Processing Systems*.
- Elmachtoub, A.N., Grigas, P., 2022. Smart “predict, then optimize”. *Management Science* 68, 9–26. doi:[10.1287/mnsc.2020.3922](https://doi.org/10.1287/mnsc.2020.3922).
- Fan, Z., 2026. Validation-calibrated risk-function ambiguity for minimax-regret portfolio optimization.
- Fissler, T., Ziegel, J.F., 2016. Higher order elicibility and Osband’s principle. *The Annals of Statistics* 44, 1680–1707. doi:[10.1214/16-AOS1439](https://doi.org/10.1214/16-AOS1439).

- Goldfarb, D., Iyengar, G., 2003. Robust portfolio selection problems. *Mathematics of Operations Research* 28, 1–38. doi:[10.1287/moor.28.1.1.14260](https://doi.org/10.1287/moor.28.1.1.14260).
- Huang, D., Zhu, S., Fabozzi, F.J., Fukushima, M., 2010. Portfolio selection under distributional uncertainty: A relative robust CVaR approach. *European Journal of Operational Research* 203, 185–194. doi:[10.1016/j.ejor.2009.07.010](https://doi.org/10.1016/j.ejor.2009.07.010).
- Lam, H., 2019. Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *Operations Research* 67, 1090–1105. doi:[10.1287/opre.2018.1786](https://doi.org/10.1287/opre.2018.1786).
- Ledoit, O., Wolf, M., 2004. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* 88, 365–411. doi:[10.1016/S0047-259X\(03\)00096-4](https://doi.org/10.1016/S0047-259X(03)00096-4).
- Liu, H., Grigas, P., 2021. Risk bounds and calibration for a smart predict-then-optimize method, in: *Advances in Neural Information Processing Systems*.
- Liyanage, L.H., Shanthikumar, J.G., 2005. A practical inventory control policy using operational statistics. *Operations Research Letters* 33, 341–348. doi:[10.1016/j.orl.2004.08.003](https://doi.org/10.1016/j.orl.2004.08.003).
- Loke, G.G., Zhu, T., Zuo, R., 2026. Optimize-via-estimate: How to make finite-sample corrections for sampling uncertainty in data-driven optimization.
- Long, D., 2025. Shift-aware Gaussian-supremum validation for Wasserstein-DRO CVaR portfolios.
- Massart, P., 1990. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *The Annals of Probability* 18, 1269–1283. doi:[10.1214/aop/1176990746](https://doi.org/10.1214/aop/1176990746).
- Maurer, A., Pontil, M., 2009. Empirical Bernstein bounds and sample-variance penalization, in: *Proceedings of the 22nd Annual Conference on Learning Theory*.
- Mišić, V.V., Perakis, G., 2020. Data analytics in operations management: A review. *Manufacturing & Service Operations Management* 22, 158–169. doi:[10.1287/msom.2019.0805](https://doi.org/10.1287/msom.2019.0805).
- Mohajerin Esfahani, P., Kuhn, D., 2018. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming* 171, 115–166. doi:[10.1007/s10107-017-1172-1](https://doi.org/10.1007/s10107-017-1172-1).
- Prashanth, L.A., Jagannathan, K., Kolla, R.K., 2020. Concentration bounds for CVaR estimation: The cases of light-tailed and heavy-tailed distributions, in: *Proceedings of the 37th International Conference on Machine Learning*, PMLR. pp. 5577–5586.
- Pun, C.S., Wang, T., Yan, Z., 2023. Data-driven distributionally robust CVaR portfolio optimization under a regime-switching ambiguity set. *Manufacturing & Service Operations Management* 25, 1779–1795. doi:[10.1287/msom.2023.1229](https://doi.org/10.1287/msom.2023.1229).
- Rockafellar, R.T., Uryasev, S., 2002. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance* 26, 1443–1471. doi:[10.1016/S0378-4266\(02\)00271-6](https://doi.org/10.1016/S0378-4266(02)00271-6).
- Smith, J.E., Winkler, R.L., 2006. The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science* 52, 311–322. doi:[10.1287/mnsc.1050.0451](https://doi.org/10.1287/mnsc.1050.0451).

- Van Parys, B.P.G., Mohajerin Esfahani, P., Kuhn, D., 2021. From data to decisions: Distributionally robust optimization is optimal. *Management Science* 67, 3387–3402. doi:[10.1287/mnsc.2020.3678](https://doi.org/10.1287/mnsc.2020.3678).
- Xidonas, P., Mavrotas, G., Hassapis, C., Zopounidis, C., 2017. Robust multiobjective portfolio optimization: A minimax regret approach. *European Journal of Operational Research* 262, 299–305. doi:[10.1016/j.ejor.2017.03.041](https://doi.org/10.1016/j.ejor.2017.03.041).
- Zhu, S., Fukushima, M., 2009. Worst-case conditional value-at-risk with application to robust portfolio management. *Operations Research* 57, 1155–1168. doi:[10.1287/opre.1080.0684](https://doi.org/10.1287/opre.1080.0684).