Conformal Robust Optimization and Satisficing for Prescriptive Analytics with Black-Box Predictors

Lingjie Zhao
Tsinghua University

Hansheng Jiang
University of Toronto

Wei Qi
Tsinghua University

Abstract

High-capacity black-box predictors are increasingly used to inform business operational decisions, yet optimizing directly over point forecasts can yield fragile solutions. We develop a conformal framework that converts any predictor into calibrated, context-dependent uncertainty sets for robust decision-making. Our framework supports two paradigms: Conformal Robust Optimization (CRO), which constructs data-driven uncertainty sets calibrated to a coverage level α, and Conformal Robust Satisficing (CRS), which minimizes fragility relative to a target τ over the full support. We prove that CRO and CRS are equivalent under mild conditions, provide finite-sample coverage guarantees concentrating at $O(n^{-1/2})$, and derive suboptimality bounds quantifying the value of prediction accuracy. Experiments on a fractional knapsack problem demonstrate that CRO improves objectives relative to baselines while maintaining calibrated coverage, and CRS improves feasibility over distributionally robust satisficing benchmarks.

1 INTRODUCTION

Modern enterprises increasingly deploy black-box predictive models, including deep neural networks, transformers, foundation models, to forecast uncertain quantities driving operational decisions. While these models deliver remarkable accuracy, plugging point forecasts directly into an optimizer often produces fragile or infeasible decisions when predictions err. This motivates a fundamental question: *how can we harness state-of-the-art black-box predictors without sacrificing the robustness that decision-makers require?*

Classical robust optimization (RO) hedges against uncertainty via worst-case reasoning over a prescribed set (Ben-Tal and Nemirovski, 1999; Bertsimas and Sim, 2004). Data-driven extensions construct sets from historical data (Bertsimas et al., 2018), but typically ignore contextual features. Distributionally robust optimization (DRO) centers ambiguity sets on empirical distributions (Delage and Ye, 2010), yet rarely exploits instance-specific predictions. Recent predict-then-calibrate approaches integrate predictors into RO, but are largely restricted to white-box models or specific problem classes (Zhu et al., 2022; Sun et al., 2023; Chan et al., 2024). A further practical hurdle is that the *size* of uncertainty sets lacks operational meaning, forcing practitioners to resort to cross-validation for tuning (Hao and Zhang, 2025).

We propose a unified framework based on *conformal prediction* (Shafer and Vovk, 2008; Angelopoulos et al., 2023), a distribution-free tool that generates prediction sets with finite-sample guarantees. Our contributions are:

(1) CRO with any black-box predictor. We construct uncertainty sets from conformal scores, replacing opaque size parameters with an interpretable coverage level $\alpha \in (0, 1)$.

(2) CRS for parameter uncertainty. We introduce Conformal Robust Satisficing, a target-oriented framework that minimizes conformal fragility—the first satisficing model integrating black-box predictors for parameter uncertainty.

(3) Theoretical guarantees. We establish finite-sample coverage with $O(n^{-1/2})$ concentration, a suboptimality bound for CRO, and prove CRO–CRS equivalence under mild convexity conditions.

(4) Numerical validation. Experiments on a fractional knapsack problem confirm the theoretical re-

---

Workshop "Towards Trustworthy Predictions: Theory and Applications of Calibration for Modern AI" at AISTATS 2026, Tangier, Morocco. Copyright 2026 by the author(s).

Algorithm 1 Conformal Uncertainty Set Construction

1: **Input**: $\mathcal{D}_{\text{train}}$, $\mathcal{D}_{\text{cal}}$, target coverage α.
2: Train black-box model $\hat{f}$ on $\mathcal{D}_{\text{train}}$.
3: Design score $s(z, d)$ (e.g., $\|d - \hat{f}(z)\|_p$).
4: Compute $S_i \leftarrow s(z_i, d_i)$ for all $(z_i, d_i) \in \mathcal{D}_{\text{cal}}$.
5: Set $\eta_\alpha = \text{Quantile}\Big(\{S_i\}, \frac{\lceil (n+1)\alpha \rceil}{n+1}\Big)$.
6: **Output**: $\mathcal{U}_\alpha(z) = \{d : s(z, d) \leq \eta_\alpha\}$.

sults and demonstrate substantial gains over context-agnostic and neighbor-based baselines.

2 FRAMEWORK: PREDICT-CALIBRATE-SOLVE

2.1 Setup and Uncertainty Set Construction

Let $\mathcal{D} = \{(z_i, d_i)\}_{i=1}^N$ denote historical data where z_i are contextual features and $d_i \in \mathbb{R}^J$ are uncertain parameters. We partition $\mathcal{D}$ into training set $\mathcal{D}_{\text{train}}$ and calibration set $\mathcal{D}_{\text{cal}}$. Our pipeline has three modules:

Predictor. A black-box model $\hat{f} : \mathcal{Z} \to \mathbb{R}^J$ trained on $\mathcal{D}_{\text{train}}$ produces $\hat{d} = \hat{f}(z)$. No restrictions on architecture or training are imposed.

Calibrator. A score function $s(z, d)$ quantifies prediction discrepancy. The uncertainty set is the level set $\mathcal{U}(z) = \{d : s(z, d) \leq \eta\}$, where η is a data-driven threshold from $\mathcal{D}_{\text{cal}}$.

Solver. The robust optimizer finds x feasible or optimal for all $d \in \mathcal{U}(z)$.

We instantiate the calibrator via conformal prediction (Algorithm 1), delivering model-agnostic, distribution-free guarantees under exchangeability.

The score function governs set geometry: norm-based scores $s(z, d) = \|d - \hat{f}(z)\|_p$ yield norm balls; coordinate-wise scores produce box sets; Mahalanobis-type scores produce ellipsoids. Adaptive scores incorporating learned residual scales $\hat{h}(z)$ allow the set to expand for hard instances and contract for easy ones, without affecting coverage validity.

2.2 Theoretical Foundation

Assumption 1 (Exchangeability)**.** *The calibration data and test point $(Z_1, D_1), \ldots, (Z_{n+1}, D_{n+1})$ are exchangeable.*

Theorem 2 (Finite-Sample Coverage)**.** *Under Assumption 1, the conformal set satisfies $\alpha \leq \mathbb{P}(D_{\text{test}} \in \mathcal{U}_\alpha(Z_{\text{test}})) \leq \alpha + \frac{1}{n+1}$.*

This guarantee is distribution-free and does not require the predictor to be correctly specified. The predictor affects only the *efficiency* (set size), not the *validity* of coverage.

Theorem 3 (Sample Complexity)**.** *Under i.i.d. data, the calibration-conditional coverage $p(\mathcal{D}_n) := \mathbb{P}(D_{\text{test}} \in \mathcal{U}_\alpha(Z_{\text{test}}) \mid \mathcal{D}_n)$ satisfies: for any $\Delta > 0$, $\mathbb{P}(p(\mathcal{D}_n) \leq \alpha - \Delta) \leq e^{-2n\Delta^2}$. Hence $p(\mathcal{D}_n) \geq \alpha - O(n^{-1/2})$ with high probability.*

3 CONFORMAL ROBUST OPTIMIZATION AND SATISFICING

3.1 Conformal Robust Optimization (CRO)

Given an objective $a(x, d)$, decision set $\mathcal{X}$, and conformal uncertainty set $\mathcal{U}_\alpha(z)$, CRO solves

$$\min_{x \in \mathcal{X}} \max_{d \in \mathcal{U}_\alpha(z)} a(x, d). \tag{1}$$

The coverage level α replaces traditional size parameters with a probabilistic guarantee (e.g., 95% reliability), allowing decision-makers to directly encode risk preferences.

Proposition 4 (Suboptimality Bound)**.** *If $a(x, d)$ is L-Lipschitz in d and $\mathbb{P}(d \in \mathcal{U}_\alpha(z)) \geq \alpha$, then $\mathbb{P}(0 \leq \delta(z, d) \leq L \cdot \text{diam}(\mathcal{U}_\alpha(z))) \geq \alpha$, where δ is the suboptimality gap.*

For linear programs with a single uncertain constraint $\min_{x \in \mathcal{X}}\{c^\top x : d^\top x \leq b\}$, we obtain a sharper structural bound.

Proposition 5. *If $d_{\text{true}} \in \mathcal{U}_\alpha(z)$ and the score satisfies the triangle inequality, then $\delta(z, d_{\text{true}}) \leq 2\eta_\alpha \cdot \lambda^* \cdot \|x^*\|_{s,*}$, where λ^* is the optimal dual multiplier and $\|x^*\|_{s,*}$ is the Lipschitz constant of x^* w.r.t. the score metric.*

The bound decomposes the "price of robustness" into three multiplicative factors. Two are of particular interest for framework design: (i) *Calibration efficiency*, captured by η_α, quantifies the size of the conformal uncertainty set and explicates the value of prediction within CRO. An accurate predictor (low bias) shifts the score distribution toward zero, directly reducing η_α; a precise predictor (low variance) yields a sharper, lighter-tailed score distribution, ensuring that realized coverage concentrates tightly around α. (ii) *Geometric alignment*, measured by $\|x^*\|_{s,*}$, captures the sensitivity of the optimization to deviations defined by the score metric s. Tailoring $s(z, d)$ (e.g., via residual scaling) to penalize deviations in directions of high objective sensitivity can reduce this term and tighten the bound.

Remark 6. *For general convex problems where the constraint $a(x, d) \leq 0$ is convex in x and differentiable in d, a similar decomposition holds under a first-order approximation. By perturbation analysis (Danskin's Theorem), the suboptimality gap is approximated locally as $\delta(z, d_{true}) \lesssim 2\eta_\alpha \lambda^* \|\nabla_d a(x^*, \tilde{d})\|_{s,*}$, where $\tilde{d}$ is the worst-case realization and $\|\nabla_d a\|_{s,*}$ replaces $\|x^*\|_{s,*}$. This confirms that the structural insights regarding prediction accuracy (η_α) and score design extend to broader problem classes.*

3.2 Conformal Robust Satisficing (CRS)

While CRO minimizes worst-case cost at a prescribed coverage level, many applications are more naturally expressed as *satisficing*: the decision-maker specifies a performance target τ and seeks the least fragile decision that meets it (Schwartz et al., 2011). Distributionally Robust Satisficing (DRS) (Long et al., 2023; Sim et al., 2024) defines fragility in terms of ambiguity over probability distributions (e.g., Wasserstein balls). In contrast, for instance-specific decisions driven by black-box forecasts, the dominant uncertainty is the realized parameter value around the forecast. We therefore introduce CRS, which measures fragility directly in the parameter space using conformal scores.

Let $\mathfrak{s}(d, \hat{d})$ be the conformal deviation metric induced by the score function, conditioned on the point forecast $\hat{d} = \hat{f}(z)$. CRS solves:

$$\begin{aligned} \kappa(\tau) = \min\ & k \\ \text{s.t. } & a(x, d) - \tau \leq k \cdot \mathfrak{s}(d, \hat{d}),\ \forall d \in \mathfrak{D}, \\ & x \in \mathcal{X},\ k \geq 0, \end{aligned} \tag{2}$$

where $\mathfrak{D} \subseteq \mathbb{R}^J$ is the support set of d. The fragility k controls the maximum ratio of target violation to parameter deviation, admitting a geometric interpretation as $\rho(v) = \sup_{d \in \mathfrak{D} \setminus \{\hat{d}\}} \frac{\max\{0, a(x,d) - \tau\}}{\mathfrak{s}(d, \hat{d})}$. Larger violations are permitted only when d is sufficiently far from $\hat{d}$ under the conformal metric; if the target is violated at $\hat{d}$ (where $\mathfrak{s} = 0$), fragility is infinite.

The nontrivial range for τ is $[\underline{\tau}, \bar{\tau}]$, where $\underline{\tau} = \min_{x \in \mathcal{X}} a(x, \hat{d})$ is the predict-then-optimize value and $\bar{\tau} = \inf_{x \in \mathcal{X}} \sup_{d \in \mathfrak{D}} a(x, d)$ is the worst-case robust value. Below $\underline{\tau}$, feasibility fails; above $\bar{\tau}$, the optimal fragility is zero.

Proposition 7 (Tail Guarantee)**.** *If $\mathbb{P}(s(d, \hat{d}) > y) \leq \Psi(y)$ for a decreasing Ψ, then for any feasible (x, k) and margin $\delta > 0$: $\mathbb{P}(a(x, d) > \tau + \delta) \leq \Psi(\delta/k)$.*

Proposition 7 shows that minimizing k compresses the tail of the cost distribution, maximizing the decay rate of violation probabilities. This tail behavior can be directly estimated from the calibration set without strong distributional assumptions.

CRS vs. DRS. The Wasserstein-based DRS (Long et al., 2023) can be equivalently written as

$$\kappa(\tau) = \min\ k, \quad \text{s.t. } \frac{1}{S} \sum_{s \in [S]} \sup_{d \in \mathfrak{D}} \{a(x, d) - k\rho(d, \hat{d}_s)\} \leq \tau, \tag{3}$$

where $\{\hat{d}_s\}_{s \in [S]}$ are historical samples and $\rho(\cdot, \cdot)$ is the transport cost. Comparing (3) with (2): DRS constrains an *average* of robustified losses anchored at historical points, whereas CRS anchors the envelope at the context-specific prediction $\hat{d} = \hat{f}(z)$ and allows feature-dependent scores to reflect heterogeneous predictive uncertainty. This makes CRS naturally suited to prediction-centric, context-dependent decision-making.

3.3 CRO–CRS Equivalence

Although CRO and CRS have different objectives, i.e., minimizing worst-case cost versus minimizing fragility, we show that they are theoretically equivalent under regularity conditions. Let $\mathfrak{D}(\theta) = \{d : \mathfrak{s}(d, \hat{d}) \leq \theta\}$, where CRO($\theta$) denotes the CRO model with this set. Under strong duality, CRO(θ) can be reformulated by introducing intermediate variables τ and k:

$$\begin{aligned} \min\ & \tau + k \cdot \theta \\ \text{s.t. } & \sup_{d \in \mathfrak{D}} \{a(x, d) - k \cdot \mathfrak{s}(d, \hat{d})\} \leq \tau, \\ & x \in \mathcal{X},\ k \geq 0. \end{aligned} \tag{4}$$

The constraint in (4) is structurally identical to that of CRS (2). This observation underpins the equivalence.

Proposition 8 (CRO to CRS)**.** *For any CRO(θ), there exists τ such that the optimal solutions of CRO(θ) are a subset of those of CRS(τ).*

The reverse direction requires regularity conditions.

Theorem 9 (CRS–CRO Equivalence)**.** *Suppose $\mathcal{X}$ is convex and $a(x, d)$ is convex in x. For any target $\tau \in [\underline{\tau}, \bar{\tau}]$, there exists θ^* such that CRS(τ) and CRO(θ^*) share the same optimal solutions, with $\theta^* = -1/\beta(\tau)$ where $\beta(\tau) \in \partial\kappa(\tau)$.*

Mechanics of Equivalence. Substituting the optimal fragility $\kappa^*(\tau)$ into (4) yields an equivalent two-stage formulation of CRO: CRO(θ) $\equiv \min_\tau\{\tau + \theta \cdot \kappa^*(\tau)\}$. The objective decomposes into a nominal cost threshold τ and a *robust cost* $\theta \cdot \kappa^*(\tau)$. Recall from Proposition 7 that κ^* governs the tail decay of violation probabilities, so the robust cost term quantifies the penalty for tail risk. The parameter θ thus acts as the "shadow price" of fragility, representing the decision-maker's marginal rate of substitution between cost efficiency and robustness. A smaller θ drives the

Table 1: CRO performance comparison.

(a) Average objective (higher magnitude = better).

α	CRO-Bo	CRO-E	CRO-Bu	KNN	KMeans	Ellips.
0.8	1149	1123	1154	608	615	0
0.9	1146	1121	1154	568	669	0
0.95	1143	1118	1155	536	490	0

Note: Ellips. achieves zero utility due to excessive conservatism.

(b) Out-of-sample coverage.

α	CRO-Bo	CRO-E	CRO-Bu	KNN	KMeans	Ellips.
0.8	0.79	0.80	0.77	0.76	0.63	0.80
0.9	0.89	0.89	0.89	0.85	0.74	0.89
0.95	0.95	0.94	0.94	0.89	0.81	0.94

solution toward a tighter target τ and a smaller uncertainty set; a larger θ signals high risk aversion, justifying a relaxed target with lower fragility. The system reaches equilibrium when the marginal reduction in robust cost from relaxing τ is exactly offset by the marginal increase in the nominal target.

3.4 Extensions

We extend the framework along both theoretical and practical dimensions. First, while Section 2.2 establishes marginal coverage guarantees, decision-making settings often call for stronger conditional validity at the instance level; we design a non-parametric algorithm to achieve this property. We also aggregate K predictors by defining the conformal score as the minimum over K basic scores, yielding a union of uncertainty sets; due to the induced decomposition, the downstream optimization remains feasible and scales linearly in K. Details are deferred to the appendix.

4 NUMERICAL EXPERIMENTS

We validate our framework on a robust fractional knapsack problem where item utilities c are uncertain and estimated from contextual features z (Ho-Nguyen and Kılınç-Karzan, 2022). The decision variable $x \in \mathbb{R}^n$ represents the fraction of each item to include, with a budget constraint. The objective is to maximize utility under worst-case realizations of c within the conformal uncertainty set.

Define the disutility function $a(x, c) = -c^\top x$, where feasible region $\mathcal{X} = \{x \in [0,1]^n : p^\top x \leq B\}$.

CRO Performance. We compare three CRO variants (Box, Ellipsoid, Budget scores) against context-agnostic Ellipsoid RO, KNN RO, and KMeans RO baselines. Table 1 reports average objectives under true parameters and out-of-sample coverage across 1,000 test instances.

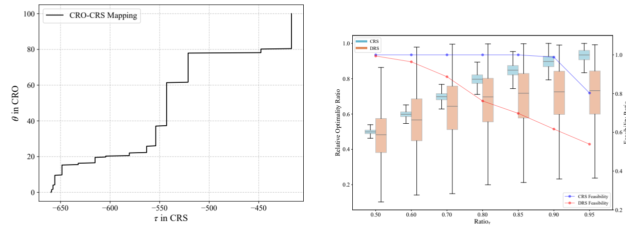

(a) CRO–CRS mapping. (b) CRS vs. DRS.

Figure 1: (a) Parameter mapping between CRS target τ and CRO size θ, validating equivalence. (b) CRS outperforms DRS in feasibility and optimality.

CRO achieves substantially higher realized utilities while closely tracking targeted coverage. The context-agnostic Ellipsoid baseline is excessively conservative (zero utility). Clustering-based methods under-cover and under-perform, confirming the value of black-box predictors for instance-specific uncertainty.

CRO–CRS Equivalence. We verify the theoretical equivalence on this problem. The prerequisites are satisfied: $a(x, c) = -c^\top x$ is affine, $\mathcal{X}$ is convex, and strong duality holds. Figure 1(a) visualizes the parameter mapping between τ and θ, corroborating Theorem 9.

CRS vs. DRS. We compare CRS against Wasserstein-based DRS (Long et al., 2023). Figure 1(b) shows CRS achieves higher feasibility and optimality ratios across target levels. DRS feasibility degrades as targets tighten due to reliance on raw historical samples, whereas CRS leverages context-specific predictions.

More Numerical Results. Additional numerical experiments on a facility location problem and a real data case study on perishable inventory management are in the appendices.

5 CONCLUSION

We introduced CRO and CRS, a unified conformal framework for prescriptive analytics with black-box predictors. CRO converts context-dependent forecasts into calibrated uncertainty sets with an interpretable coverage parameter. CRS offers a target-oriented counterpart minimizing conformal fragility. Under mild conditions, the two are equivalent with an explicit parameter mapping. Experiments confirm substantial gains over context-agnostic baselines. Future work includes extensions to conditional coverage and distributionally robust settings.

References

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

Ben-Tal, A. and Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13.

Bertsimas, D., Gupta, V., and Kallus, N. (2018). Data-driven robust optimization. *Mathematical Programming*, 167:235–292.

Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations research*, 52(1):35–53.

Chan, T., Delage, E., and Lin, B. (2024). Conformal inverse optimization for adherence-aware prescriptive analytics. *Available at SSRN 4970102*.

Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612.

Hao, H. and Zhang, P. (2025). Robust paths: Geometry and computation. *arXiv preprint arXiv:2508.20039*.

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

Long, D. Z., Sim, M., and Zhou, M. (2023). Robust satisficing. *Operations Research*, 71(1):61–82.

Schwartz, B., Ben-Haim, Y., and Dacso, C. (2011). What makes a good decision? robust satisficing as a normative standard of rational decision making. *Journal for the Theory of Social Behaviour*, 41(2):209–227.

Shafer, G. and Vovk, V. (2008). A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3).

Sim, M., Tang, Q., Zhou, M., and Zhu, T. (2024). The analytics of robust satisficing: predict, optimize, satisfice, then fortify. *Operations Research*, 73(5):2708–2728.

Sun, C., Liu, L., and Li, X. (2023). Predict-then-calibrate: A new perspective of robust contextual lp. *Advances in Neural Information Processing Systems*, 36:17713–17741.

Zhu, T., Xie, J., and Sim, M. (2022). Joint estimation and robustness optimization. *Management Science*, 68(3):1659–1677.