

Distributionally Robust Regret Minimization

Eilyan Bitar*

Abstract

We consider decision-making problems involving the optimization of linear objective functions with uncertain coefficients. The probability distribution of the coefficients—which are assumed to be stochastic in nature—is unknown to the decision maker but is assumed to lie within a given ambiguity set, defined as a type-1 Wasserstein ball centered at a given nominal distribution. To account for this uncertainty, we minimize the worst-case expected *regret* over all distributions in the ambiguity set. Here, the (ex post) regret experienced by the decision maker is defined as the difference between the cost incurred by a chosen decision given a particular realization of the objective coefficients and the minimum achievable cost with perfect knowledge of the coefficients at the outset. For this class of ambiguity sets, the worst-case expected regret is shown to equal the expected regret under the nominal distribution plus a regularization term that has the effect of drawing optimal solutions toward the “center” of the feasible region as the radius of the ambiguity set increases. This novel form of regularization is also shown to arise when minimizing the worst-case conditional value-at-risk (CVaR) of regret. We show that, under certain conditions, distributionally robust regret minimization problems over type-1 Wasserstein balls can be recast as tractable finite-dimensional convex programs.

1 Introduction

We consider a class of uncertain linear optimization problems in which the coefficients of the objective function are unknown, and a decision must be made without prior knowledge of these coefficients. To handle such uncertainty in the problem data, a common approach is to assume that the unknown parameters belong to a given set, known as the uncertainty set, and to construct a solution that is optimal for the least-favorable parameter values in the given uncertainty set. This corresponds to the following class of robust linear optimization problems:

$$\inf_{x \in \mathcal{X}} \sup_{w \in \mathcal{W}} w^\top x. \quad (1)$$

In problem (1), the decision vector $x \in \mathbb{R}^n$ is constrained to lie within a feasible set $\mathcal{X} \subseteq \mathbb{R}^n$ and the vector of uncertain parameters $w \in \mathbb{R}^n$ is confined to an uncertainty set $\mathcal{W} \subseteq \mathbb{R}^n$.

Robust optimization problems, as defined in (1), are known to be computationally tractable under certain conditions. For example, if the feasible set \mathcal{X} is polyhedral and the uncertainty set \mathcal{W} is polyhedral or ellipsoidal, then the problem (1) can be reformulated as a polynomially sized linear program or a second-order cone program, respectively [6, 7]. However, despite its broad appeal, a potential drawback of the robust optimization framework is its inherent conservatism. By optimizing for the worst-case scenario within the uncertainty set, robust optimization methods can lead to overly cautious decisions that may underperform when facing less severe realizations of the uncertain parameters.

*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA. Email: eyb5@cornell.edu

To mitigate this conservatism, robust regret minimization has been explored as an alternative to the robust optimization approach. Originally proposed in [26], regret measures the difference between the cost realized by a chosen decision and the best possible cost that could have been achieved had the uncertain parameters been known at the outset. Mathematically, the regret incurred by a decision $x \in \mathcal{X}$ given a particular realization of the uncertain parameters $w \in \mathcal{W}$ is defined as

$$R(x, w) := w^\top x - \inf_{y \in \mathcal{X}} w^\top y. \quad (2)$$

With this definition, the robust regret minimization problem can be expressed as

$$\inf_{x \in \mathcal{X}} \sup_{w \in \mathcal{W}} R(x, w). \quad (3)$$

We note that a related class of robust regret minimization problems, where uncertainty is confined to the constraints on x , is known to be computationally tractable. Specifically, if the feasible set \mathcal{X} and the uncertainty set \mathcal{W} are both polyhedral, and the uncertain parameters w only appear linearly in the right-hand side of the constraints on x , then the corresponding robust regret minimization problem can be reformulated as an equivalent linear program [8, 14]. In contrast, robust regret minimization problems with uncertain coefficients in the objective function, as defined in (3), are more difficult to solve. In particular, problem (3) has been proven to be strongly NP-hard when the uncertainty set is an axis-aligned hyperrectangle [4]. As a result, most of the algorithms that have been developed to solve problem (3) are either conservative in nature or computationally expensive.¹ There are, however, special cases of the minimax regret optimization problem (3) that can be solved in polynomial time, including specific resource allocation problems with rectangular uncertainty sets [3, 13], and certain inventory management problems with simplex uncertainty sets [24].

1.1 A Distributionally Robust Approach to Regret Minimization

In this paper, we approach the problem of regret minimization from the perspective of distributionally robust optimization. Unlike the robust formulation in (3), where the unknown parameters are allowed to take any value within a predefined uncertainty set, the distributionally robust approach treats the uncertain parameters as a random vector \mathbf{w} whose distribution P is unknown, but assumed to lie within a given ambiguity set \mathcal{P} —a predefined family of possible distributions. Using the given ambiguity set, decisions are constructed to minimize the worst-case expected regret across all possible distributions within the ambiguity set. Mathematically, this leads to the following distributionally robust regret minimization problem considered in this paper:

$$\inf_{x \in \mathcal{X}} \sup_{P \in \mathcal{P}} \mathbb{E}_P [R(x, \mathbf{w})]. \quad (4)$$

It is important to note that, in problem (4), regret is evaluated *ex post*, meaning it is assessed after the uncertain parameters have been realized. This differs from the *ex ante* approach to distributionally robust regret minimization, where the *ex ante* regret of a decision is defined as the difference between the *expected cost* incurred by that decision and the minimum expected cost, assuming perfect knowledge of the underlying probability distribution. Distributionally robust *ex ante* regret minimization has been extensively studied in the context of the newsvendor model

¹We refer the reader to [24] for a comprehensive overview of the various algorithms that have been developed to tackle this class of single-stage robust regret minimization problems.

[11, 23, 27, 29], two-stage stochastic programs [12], and multi-stage stochastic programs defined over discrete probability spaces [25].

The formulation in [22] is more closely aligned with the approach taken in this paper, as it employs an *ex post* measure of regret. Specifically, they define the ambiguity set as the collection of all distributions with common support (a hyperrectangle), mean, and higher-order marginal moments. With this ambiguity set, they consider a risk-sensitive variant of the optimization model in (4), where decisions are chosen to minimize the worst-case conditional value-at-risk (CVaR) of the *ex post* regret. However, since the moments of a distribution are typically estimated from data and, therefore, may themselves be uncertain, we consider an alternative class of (distance-based) ambiguity sets that do not impose such rigid restrictions a priori. Specifically, in this paper, the ambiguity set is taken to be the type-1 Wasserstein ball distributions centered at a given nominal distribution (e.g., an empirical distribution constructed from measured data), where the radius of this ball represents the decision maker’s confidence in the accuracy of the nominal distribution.

While the focus of this paper is on single-stage stochastic programs, we note that Wasserstein distance-based ambiguity sets have also been utilized to synthesize distributionally robust (*ex post*) regret optimal control policies for unconstrained stochastic linear-quadratic optimal control problems, where ambiguity in the process disturbance and measurement noise distributions is modeled using type-2 Wasserstein balls [1, 16, 18].

1.2 Main Results and Paper Outline

In Section 2, we introduce the specific class of distributionally robust (*ex post*) regret minimization problems (4) considered in this paper, where the ambiguity set \mathcal{P} is defined as a type-1 Wasserstein ball, and the feasible region \mathcal{X} is assumed to be nonempty, compact, and possibly nonconvex. In Section 2.1, we show that the worst-case expected regret of a decision $x \in \mathcal{X}$ can be expressed as the sum of the expected regret under the nominal distribution and a regularization term. This regularization term, which scales with the radius of the ambiguity set, penalizes the distance of the decision x to the farthest point in the feasible set. This has the effect of drawing optimal solutions to problem (4) toward the “center” of the feasible region as the Wasserstein radius increases.

Based on this reformulation, the Wasserstein distributionally robust regret minimization problem (4) is shown to “interpolate” between the nominal expected cost minimization problem, as the radius of the ambiguity set approaches zero, and the robust regret minimization problem (3) with a norm-bounded uncertainty set, as the radius of the ambiguity set grows large. In Section 3, this specific form of regularization is also shown to emerge when considering a risk-sensitive variant of problem (4), where decisions are selected to minimize the worst-case conditional value-at-risk of regret over all distributions in the type-1 Wasserstein ball.

Finally, while the regularization function is shown to be NP-hard to compute in general (it is equivalent to a norm maximization problem over the feasible set \mathcal{X}), in Section 2.2, we identify conditions on the structure of the feasible set and the choice of norm used in the definition of the Wasserstein distance that enable the equivalent reformulation of distributionally robust regret minimization problem (8) as a tractable convex program.

1.3 Notation

Given a norm $\|\cdot\|$ on \mathbb{R}^n , the dual norm is defined as $\|x\|_* := \sup\{x^\top y \mid \|y\| \leq 1\}$ for all $x \in \mathbb{R}^n$. The support function associated with a set $\mathcal{X} \subseteq \mathbb{R}^n$ is defined as $\sigma_{\mathcal{X}}(y) := \sup_{x \in \mathcal{X}} x^\top y$ for all $y \in \mathbb{R}^n$. We use boldface symbols to denote random variables, and non-boldface symbols to denote particular values in the range of a random variable and other deterministic quantities.

2 Distributionally Robust Expected Regret Minimization

In this section, we employ a duality theory for optimal transport problems to show that the distributionally robust expected regret minimization problem (4), with a type-1 Wasserstein distance-based ambiguity set, is equivalent to a regularized version of the expected cost minimization problem under the nominal distribution. Interestingly, the resulting form of regularization is shown to encourage solutions that are “centrally positioned” within the feasible set \mathcal{X} . In Section 2.2, we establish conditions on the structure of the feasible set and the choice of norm used in defining the Wasserstein distance that enable the reformulation of problem (4) as a finite-dimensional convex program. In Section 2.3, we draw a comparison between this framework and the more widely studied distributionally robust expected cost minimization approach.

2.1 Strong Duality for Distributionally Robust Regret Minimization Problems

Unless stated otherwise, we let $\|\cdot\|$ be an arbitrary norm on \mathbb{R}^n , and denote by $\mathcal{M}(\mathbb{R}^n)$ the collection of Borel probability measures P on \mathbb{R}^n that satisfy $\mathbb{E}_P[\|\mathbf{w}\|] < \infty$. The type-1 Wasserstein distance between probability measures is defined as follows.

Definition 1 (Type-1 Wasserstein Distance). The *type-1 Wasserstein distance* between two distributions $P_1, P_2 \in \mathcal{M}(\mathbb{R}^n)$ is defined as

$$W_1(P_1, P_2) := \inf_{\pi \in \Pi(P_1, P_2)} \int_{\mathbb{R}^{N_x} \times \mathbb{R}^{N_x}} \|z_1 - z_2\| \pi(dz_1, dz_2),$$

where $\Pi(P_1, P_2)$ denotes the set of all joint distributions in $\mathcal{M}(\mathbb{R}^n \times \mathbb{R}^n)$ with marginal distributions P_1 and P_2 .

Given a *nominal distribution* $P_0 \in \mathcal{M}(\mathbb{R}^n)$, we define the ambiguity set \mathcal{P} as the set of all distributions whose type-1 Wasserstein distance to P_0 is at most $r \geq 0$, i.e.,

$$\mathcal{P} := \{P \in \mathcal{M}(\mathbb{R}^n) \mid W_1(P, P_0) \leq r\}. \quad (5)$$

The *radius of the ambiguity set*, r , can be interpreted as encoding the decision maker’s level of confidence in how accurately the nominal distribution reflects the true underlying distribution. In particular, the ambiguity set becomes a singleton containing only the nominal distribution if the ambiguity set radius is decreased to zero.

Given this specific family of ambiguity sets, we now turn our attention to the inner maximization (worst-case expectation) in the distributionally robust regret minimization problem (4). Strong duality has been shown to hold for a broad family of Wasserstein worst-case expectation problems [9, 15, 21]. In particular, the following result establishes conditions under which strong duality is guaranteed to hold, along with a characterization of the corresponding dual optimization problem.

Theorem 1 (Strong Duality, [15]). Let $r > 0$ and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a Borel-measurable function such that $\mathbb{E}_{P_0}[|f(\mathbf{w})|] < \infty$. It holds that

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P[f(\mathbf{w})] = \inf_{\lambda \geq 0} \mathbb{E}_{P_0} \left[\sup_{z \in \mathbb{R}^n} \{f(z) + \lambda(r - \|z - \mathbf{w}\|)\} \right]. \quad (6)$$

We utilize Theorem 1 to derive an equivalent dual reformulation for the worst-case expected regret over the type-1 Wasserstein ball of distributions.

Theorem 2 (Worst-Case Expected Regret). Let $r \geq 0$ and assume that the feasible set $\mathcal{X} \subset \mathbb{R}^n$ is nonempty and compact. The worst-case expected regret incurred by a decision $x \in \mathcal{X}$ can be expressed as:

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P [R(x, \mathbf{w})] = \mathbb{E}_{P_0} [R(x, \mathbf{w})] + r \sup_{v \in \mathcal{X}} \|x - v\|_*. \quad (7)$$

Furthermore, a point $x \in \mathcal{X}$ is an optimal solution to the distributionally robust regret minimization problem (4) if and only if it is an optimal solution to the following problem:

$$\inf_{x \in \mathcal{X}} \left\{ \mathbb{E}_{P_0} [\mathbf{w}^\top x] + r \sup_{v \in \mathcal{X}} \|x - v\|_* \right\}. \quad (8)$$

Proof. When $r = 0$, the result is immediate, as the ambiguity set reduces to a singleton containing only the nominal distribution P_0 . For $r > 0$, we begin by proving the identity stated in Equation (7). First, observe that for each $x \in \mathcal{X}$, the regret function $R(x, w) = w^\top x - \inf_{y \in \mathcal{X}} w^\top y$ is convex in w over \mathbb{R}^n and, as a result, is Borel measurable. Also, since it is assumed that $P_0 \in \mathcal{M}(\mathbb{R}^n)$ and that the feasible set \mathcal{X} is bounded, it holds that $\mathbb{E}_{P_0} [|R(x, \mathbf{w})|] \leq \mathbb{E}_{P_0} [\|\mathbf{w}\|] \sup_{y \in \mathcal{X}} \|x - y\|_* < \infty$ for all $x \in \mathcal{X}$. Given these properties, Theorem 1 can be applied to express the worst-case expected regret as

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P [R(x, \mathbf{w})] = \inf_{\lambda \geq 0} \mathbb{E}_{P_0} \left[\sup_{z \in \mathbb{R}^n} \{R(x, z) + \lambda(r - \|z - \mathbf{w}\|)\} \right]. \quad (9)$$

Substituting $R(x, z) := z^\top x - \inf_{y \in \mathcal{X}} \{z^\top y\}$ into the right-hand side of Eq. (9) and rearranging terms, the worst-case expected regret can be rewritten as

$$\inf_{\lambda \geq 0} \left\{ \lambda r + \mathbb{E}_{P_0} \left[\sup_{z \in \mathbb{R}^n} \sup_{y \in \mathcal{X}} \{z^\top (x - y) - \lambda \|z - \mathbf{w}\|\} \right] \right\}. \quad (10)$$

Next, we exchange the order of the supremum taken with respect to $z \in \mathbb{R}^n$ with the supremum taken with respect to $y \in \mathcal{X}$ in (10). Then, using the fact that the conjugate function of $\lambda \|\cdot\|$ is given by

$$\sup_{z \in \mathbb{R}^n} \{z^\top \xi - \lambda \|z\|\} = \begin{cases} 0 & \text{if } \|\xi\|_* \leq \lambda, \\ \infty & \text{otherwise,} \end{cases}$$

for all $\xi \in \mathbb{R}^n$, we can simplify the optimization problem in (10) to

$$\inf_{\lambda \geq 0} \left\{ \lambda r + \mathbb{E}_{P_0} \left[\sup_{y \in \mathcal{X}} \left\{ \begin{array}{ll} \mathbf{w}^\top (x - y) & \text{if } \|x - y\|_* \leq \lambda, \\ \infty & \text{otherwise} \end{array} \right\} \right] \right\}. \quad (11)$$

Clearly, the value of the expectation in (11) is infinite if there exists any element $v \in \mathcal{X}$ such that $\|x - v\|_* > \lambda$. Hence, the optimization problem (11) can be rewritten as

$$\inf_{\lambda \geq 0} \left\{ \lambda r + \mathbb{E}_{P_0} \left[\sup_{y \in \mathcal{X}} \{\mathbf{w}^\top (x - y)\} \right] \mid \|x - v\|_* \leq \lambda \text{ for all } v \in \mathcal{X} \right\}. \quad (12)$$

From (12), the optimal dual variable is found to be $\lambda^* = \sup_{v \in \mathcal{X}} \|x - v\|_*$. Thus, the optimal value of problem (12) is given by

$$r \sup_{v \in \mathcal{X}} \|x - v\|_* + \mathbb{E}_{P_0} \left[\sup_{y \in \mathcal{X}} \{\mathbf{w}^\top (x - y)\} \right], \quad (13)$$

where the second term in (13) is $\mathbb{E}_{P_0}[R(x, \mathbf{w})]$, proving the equivalence in (7). Finally, the equivalence of the optimal solution sets for problems (4) and (8) is a direct consequence of (7), which implies that

$$\inf_{x \in \mathcal{X}} \sup_{P \in \mathcal{P}} \mathbb{E}_P [R(x, \mathbf{w})] = \inf_{x \in \mathcal{X}} \left\{ \mathbb{E}_{P_0}[\mathbf{w}^\top x] + r \sup_{v \in \mathcal{X}} \|x - v\|_* \right\} - \mathbb{E}_{P_0} \left[\inf_{y \in \mathcal{X}} \mathbf{w}^\top y \right], \quad (14)$$

thus completing the proof. \square

Remark 1 (Regularization). Theorem 2 shows that the worst-case expected regret incurred by a decision $x \in \mathcal{X}$ over all distributions in the type-1 Wasserstein ball can be expressed as the sum of two terms: the expected regret of the decision x under the nominal distribution and a regularization term that penalizes the distance of the decision x to the farthest point in the feasible set \mathcal{X} . Specifically, the regularization function $\sup_{v \in \mathcal{X}} \|x - v\|_*$ represents the radius of the smallest dual-norm ball centered at x that contains the feasible set \mathcal{X} . Thus, as the radius r of the ambiguity set \mathcal{P} increases, the set of optimal solutions to the distributionally robust regret minimization problem (4) will gravitate towards the “center” of the feasible set \mathcal{X} , defined by $\arg \min_{x \in \mathcal{X}} \sup_{v \in \mathcal{X}} \|x - v\|_*$.²

The regularization function $\sup_{v \in \mathcal{X}} \|x - v\|_*$ can also be shown to coincide with the Lipschitz modulus of the regret function $R(x, \cdot)$ on \mathbb{R}^n with respect to the norm $\|\cdot\|$. “Lipschitz regularization” has also been shown to arise in distributionally robust optimization problems equipped with type-1 Wasserstein ambiguity sets centered at discrete (empirical) nominal distributions [21, Theorem 6.3], [19, Theorem 10].

Remark 2 (Interpolation). Interestingly, the regularization function can also be interpreted as the worst-case regret incurred by a decision $x \in \mathcal{X}$, where the supremum is taken over all cost vectors w within a norm-bounded uncertainty set $\mathcal{W} = \{w \in \mathbb{R}^n \mid \|w\| \leq 1\}$. Specifically, we have that

$$\sup_{\|w\| \leq 1} R(x, w) = \sup_{\|w\| \leq 1} \sup_{y \in \mathcal{X}} \{w^\top (x - y)\} = \sup_{y \in \mathcal{X}} \|x - y\|_*,$$

where the last equality is obtained by exchanging the order of the supremum taken with respect to w and y , and using the fact that $\sup_{\|w\| \leq 1} w^\top (x - y) = \|x - y\|_*$. With this interpretation, problem (8) can be rewritten as

$$\inf_{x \in \mathcal{X}} \left\{ \mathbb{E}_{P_0}[\mathbf{w}^\top x] + r \sup_{\|w\| \leq 1} R(x, w) \right\}. \quad (15)$$

Thus, the Wasserstein distributionally robust regret minimization problem can be understood as “interpolating” between the nominal expected cost minimization problem $\inf_{x \in \mathcal{X}} \mathbb{E}_{P_0}[\mathbf{w}^\top x]$ (when $r \rightarrow 0$) and the robust regret minimization problem $\inf_{x \in \mathcal{X}} \sup_{\|w\| \leq 1} R(x, w)$ (when $r \rightarrow \infty$). The radius of the ambiguity set r , which reflects the decision maker’s confidence in the accuracy of the nominal distribution, acts as the regularization parameter that governs this tradeoff.

2.2 Complexity and Tractable Reformulation

The evaluation of the regularization function $\sup_{v \in \mathcal{X}} \|x - v\|_*$ in problem (8) amounts to solving a norm maximization problem over the feasible set \mathcal{X} . Such problems are known to be NP-hard when

²While this notion of a “central point” in the feasible set \mathcal{X} is similar to the Chebyshev center of the set \mathcal{X} , which is typically defined as the center of the smallest ball that contains the set \mathcal{X} , it differs in that it requires the center of the enclosing ball to lie within the set \mathcal{X} . In contrast, the Chebyshev center may not belong to \mathcal{X} , even if \mathcal{X} is a convex set.

$\|\cdot\|_*$ is a p -norm with $p \in [1, \infty)$ and \mathcal{X} is a compact, convex polytope in halfspace representation [20, Theorem 4.1]. However, if the feasible set can be expressed as the convex hull of a finite set of points given by $\mathcal{X} = \text{conv}\{v_1, \dots, v_m\}$, then the norm maximization problem takes on a particularly simple form:

$$\sup_{v \in \mathcal{X}} \|x - v\|_* = \max_{i \in \{1, \dots, m\}} \|x - v_i\|_*.$$

This follows from the fact that the supremum of a convex function over a convex and compact set is attained at an extreme point of that set. With this simplification, the distributionally robust regret minimization problem (8) can be reformulated as the following tractable convex program in the variables x and λ :

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{P_0}[\mathbf{w}^\top x] + r\lambda \\ & \text{subject to} && x \in \mathcal{X}, \lambda \in \mathbb{R} \\ & && \|x - v_i\|_* \leq \lambda, \quad i = 1, \dots, m. \end{aligned} \tag{16}$$

As another example, consider the case where $\|\cdot\|_*$ is the ∞ -norm. In this setting, the norm maximization problem $\sup_{v \in \mathcal{X}} \|x - v\|_\infty$ can be solved by maximizing the absolute value of each component of the vector $x - v$ separately, subject to the constraint $v \in \mathcal{X}$. Specifically, given a (possibly nonconvex) feasible set $\mathcal{X} \subseteq \mathbb{R}^n$, we have

$$\sup_{v \in \mathcal{X}} \|x - v\|_\infty = \max_{i \in \{1, \dots, n\}} \max \left\{ \sup_{v \in \mathcal{X}} e_i^\top (x - v), \sup_{v \in \mathcal{X}} e_i^\top (v - x) \right\}, \tag{17}$$

where the vectors $e_1, \dots, e_n \in \mathbb{R}^n$ denote the standard basis for \mathbb{R}^n . This suggests that when the Wasserstein distance is defined in terms of the 1-norm (the dual of the ∞ -norm), the distributionally robust regret minimization problem (8) can be reformulated as shown in the following corollary, which follows directly from Theorem 2 and the simplified form of the regularization function in Eq. (17).

Corollary 1. Let $r \geq 0$ and assume that the set $\mathcal{X} \subseteq \mathbb{R}^n$ is nonempty and compact. If $\|\cdot\|_*$ is the ∞ -norm, then problem (8) is equivalent to the following optimization problem in the variables x and λ :

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{P_0}[\mathbf{w}^\top x] + r\lambda \\ & \text{subject to} && x \in \mathcal{X}, \lambda \in \mathbb{R} \\ & && \sigma_{\mathcal{X}}(e_i) - e_i^\top x \leq \lambda, \quad i = 1, \dots, n \\ & && \sigma_{\mathcal{X}}(-e_i) + e_i^\top x \leq \lambda, \quad i = 1, \dots, n. \end{aligned} \tag{18}$$

Here, $\sigma_{\mathcal{X}}$ denotes the support function associated with the set \mathcal{X} .

While problem (18) is a valid reformulation of the distributionally robust regret minimization problem (8) for any nonempty, compact set \mathcal{X} , the computational tractability of problem (18) depends on the underlying geometry of the feasible set \mathcal{X} . For example, if \mathcal{X} is a convex set, then problem (18) reduces to a convex optimization problem, where the support function values $\sigma_{\mathcal{X}}(\pm e_i)$ for $i = 1, \dots, n$ can be computed by solving $2n$ convex optimization problems, either numerically or in closed form in certain cases. See [5, Section 3] for an overview of a variety of convex sets whose support functions can be determined in closed form or recast as minimization problems using strong duality.

2.3 Comparison to Distributionally Robust Optimization

In this section, we compare the distributionally robust regret minimization approach with the more widely studied distributionally robust expected cost minimization framework—commonly referred to as distributionally robust optimization (DRO)—where the decision-maker seeks to minimize the worst-case expected cost across all distributions within a specified ambiguity set by solving:

$$\inf_{x \in \mathcal{X}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[\mathbf{w}^\top x]. \quad (19)$$

While both methods aim to enhance the robustness of decisions to distributional uncertainty, they induce distinct forms of regularization, leading to fundamental differences in the behavior of their optimal solutions as the Wasserstein radius r varies. To illustrate this difference more precisely, note that the worst-case expected cost incurred by a decision $x \in \mathcal{X}$ can be expressed as the sum of the expected cost under the nominal distribution and a regularization term that penalizes the dual norm of x , scaled by the radius of the ambiguity set r :

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P[\mathbf{w}^\top x] = \mathbb{E}_{P_0}[\mathbf{w}^\top x] + r\|x\|_*. \quad (20)$$

The equivalence in (20) was originally shown to hold for type-1 Wasserstein ambiguity sets centered at empirical nominal distributions [19, 21]. By applying Theorem 1, the reformulation in (20) can also be shown to hold for any nominal distribution $P_0 \in \mathcal{M}(\mathbb{R}^n)$. We omit the proof here, as it closely follows the proof of Theorem 2.

Eq. (20) reveals that distributionally robust expected cost minimization results in a form of regularization that encourages decisions that are small in dual norm, driving the set of optimal solutions toward the origin as the ambiguity set radius r increases. In contrast, the regularizer induced by the distributionally robust expected regret minimization approach draws the set of optimal solutions toward the “center” of the feasible region as r increases. Figure 1 illustrates these contrasting behaviors with a simple two-dimensional example.

3 Worst-Case Conditional Value-at-Risk of Regret Minimization

Drawing inspiration from the decision criterion proposed in [22], we now consider a risk-sensitive variant of problem (4), where decisions are made to minimize the worst-case conditional value-at-risk (CVaR) of regret over the type-1 Wasserstein ball of distributions. Specifically, given a confidence level $\alpha \in [0, 1)$, the conditional value-at-risk of regret incurred by a decision $x \in \mathcal{X}$ under a distribution $P \in \mathcal{P}$ is defined as

$$\text{CVaR}_\alpha^P(R(x, \mathbf{w})) := \inf_{\tau \in \mathbb{R}} \left\{ \tau + \frac{1}{1 - \alpha} \mathbb{E}_P[\max\{R(x, \mathbf{w}) - \tau, 0\}] \right\}. \quad (21)$$

As a coherent risk measure [2], CVaR is guaranteed to be convexity preserving. This ensures convexity of the (worst-case) CVaR of regret with respect to the decision variable x , since the regret function $R(x, w)$ is linear in x for each $w \in \mathbb{R}^n$. Utilizing this class of risk measures, we formulate the worst-case CVaR of regret minimization problem as

$$\inf_{x \in \mathcal{X}} \sup_{P \in \mathcal{P}} \text{CVaR}_\alpha^P(R(x, \mathbf{w})). \quad (22)$$

As the confidence level $\alpha \uparrow 1$, the CVaR of regret, $\text{CVaR}_\alpha^P(R(x, \mathbf{w}))$, converges to the essential supremum of regret, defined as $\text{ess sup}_P(R(x, \mathbf{w})) := \inf_{\tau \in \mathbb{R}} \{\tau \mid P(R(x, \mathbf{w}) > \tau) = 0\}$. When $\alpha =$




Figure 1: Illustration of optimal solutions to the distributionally robust regret minimization problem (4) (blue stars) and the distributionally robust cost minimization problem (19) (red squares) as the ambiguity set radius r ranges from 10^{-2} to 10 over 50 logarithmically spaced points. In this example, the nominal expected value of \mathbf{w} is taken to be $\mathbb{E}_{P_0}[\mathbf{w}] = [-1/2, 2]^\top$, the dual norm $\|\cdot\|_*$ is the ∞ -norm, and the feasible set $\mathcal{X} = \{x \in \mathbb{R}^2 \mid \|x - 1\|_2 \leq 1\}$ is shown as the shaded gray region.

0, it coincides with the expected regret $\mathbb{E}_P[R(x, \mathbf{w})]$, and problem (22) reduces to the worst-case expected regret minimization problem (4). The following result extends Theorem 2 to encompass this more general family of risk-sensitive regret minimization problems.

Theorem 3 (Worst-Case CVaR of Regret). Let $r \geq 0$ and assume that the feasible set $\mathcal{X} \subset \mathbb{R}^n$ is nonempty and compact. Given a confidence level $\alpha \in [0, 1)$, the worst-case conditional value-at-risk of regret incurred by a decision $x \in \mathcal{X}$ can be expressed as:

$$\sup_{P \in \mathcal{P}} \text{CVaR}_\alpha^P(R(x, \mathbf{w})) = \text{CVaR}_\alpha^{P_0}(R(x, \mathbf{w})) + \left(\frac{r}{1 - \alpha}\right) \sup_{v \in \mathcal{X}} \|x - v\|_*. \quad (23)$$

Proof. When $r = 0$, the result is immediate, as the ambiguity set reduces to a singleton containing only the nominal distribution P_0 . For $r > 0$, recall that

$$\sup_{P \in \mathcal{P}} \text{CVaR}_\alpha^P(R(x, \mathbf{w})) = \sup_{P \in \mathcal{P}} \inf_{\tau \in \mathbb{R}} \left\{ \tau + \frac{1}{1 - \alpha} \mathbb{E}_P[\max\{R(x, \mathbf{w}) - \tau, 0\}] \right\}. \quad (24)$$

Next, we next invoke a version of the minimax theorem to exchange the order of the infimum and supremum in (24). To do so, note that the objective function in (24) is convex in the variable $\tau \in \mathbb{R}$ and linear in the probability measure P over the type-1 Wasserstein ball \mathcal{P} , which is a nonempty, convex, and weakly compact set [28, Theorem 1]. The objective function can also be shown to be weakly continuous in P on the set \mathcal{P} . Given these properties, the minimax theorem from [10, Theorem 4.5] can be applied to reformulate the worst-case CVaR of regret in (24) as

$$\sup_{P \in \mathcal{P}} \text{CVaR}_\alpha^P(R(x, \mathbf{w})) = \inf_{\tau \in \mathbb{R}} \left\{ \tau + \frac{1}{1 - \alpha} \sup_{P \in \mathcal{P}} \mathbb{E}_P[\max\{R(x, \mathbf{w}) - \tau, 0\}] \right\}. \quad (25)$$

To complete the proof, one can employ Theorem 1 to reformulate the worst-case expectation on the right-hand side of Eq. (25) in its dual form, and follow a line of reasoning similar to that used in the proof of Theorem 2 to show that

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P [\max\{R(x, \mathbf{w}) - \tau, 0\}] = \mathbb{E}_{P_0} [\max\{R(x, \mathbf{w}) - \tau, 0\}] + r \sup_{v \in \mathcal{X}} \|x - v\|_*.$$

Substituting the above expression into Eq. (25) and rearranging terms yields the desired result. \square

Theorem 3 shows that for $\alpha \in [0, 1)$, the worst-case CVaR of regret minimization problem (22) can be reformulated as

$$\inf_{x \in \mathcal{X}} \left\{ \text{CVaR}_\alpha^{P_0}(R(x, \mathbf{w})) + \left(\frac{r}{1 - \alpha} \right) \sup_{v \in \mathcal{X}} \|x - v\|_* \right\}. \quad (26)$$

This reformulation possesses the same form of regularization identified in Theorem 2, but with a rescaling of the regularization parameter by a factor of $1/(1 - \alpha)$, reflecting a dependence on the confidence level α .

Remark 3 (Computational tractability). As expected, the reformulation (26) reduces to the one provided in Theorem 2 when $\alpha = 0$. To solve problem (26) for $\alpha \in (0, 1)$, one must calculate the CVaR of regret under the nominal distribution P_0 . However, despite their convexity, computing CVaR measures with high accuracy is known to be challenging, even for relatively simple cost functions and distributions.³ This difficulty is avoided when P_0 is a discrete distribution supported on a finite set of moderate cardinality. Together with this simplifying assumption on the support of P_0 , the structural conditions provided in Section 2.2 can be utilized to reformulate problem (26) as a finite-dimensional convex program.

References

- [1] Feras Al Taha, Shuhao Yan, and Eilyan Bitar. A distributionally robust approach to regret optimal control using the wasserstein distance. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 2768–2775. IEEE, 2023.
- [2] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.
- [3] Igor Averbakh. Minmax regret linear resource allocation problems. *Operations Research Letters*, 32(2):174–180, 2004.
- [4] Igor Averbakh and Vasilij Lebedev. On the complexity of minmax regret linear programming. *European Journal of Operational Research*, 160(1):227–231, 2005.
- [5] Aharon Ben-Tal, Dick Den Hertog, and Jean-Philippe Vial. Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical programming*, 149(1):265–299, 2015.
- [6] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton university press, 2009.
- [7] Aharon Ben-Tal and Arkadi Nemirovski. Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13, 1999.

³For example, [17, Corollary 1] shows that it is #P-hard to compute the expected value of the non-negative part of a linear combination of uniformly distributed random variables.

- [8] Dimitris Bertsimas and Iain Dunning. Relative robust and adaptive optimization. *INFORMS Journal on Computing*, 32(2):408–427, 2020.
- [9] Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.
- [10] Jonathan M Borwein and Qiji J Zhu. A variational approach to lagrange multipliers. *Journal of Optimization Theory and Applications*, 171:727–756, 2016.
- [11] Zhi Chen and Weijun Xie. Regret in the newsvendor model with demand and yield randomness. *Production and Operations Management*, 30(11):4176–4197, 2021.
- [12] Youngchae Cho and Insoon Yang. Wasserstein distributionally robust regret minimization. *IEEE Control Systems Letters*, 2024.
- [13] Eduardo Conde. On the complexity of the continuous unbounded knapsack problem with uncertain coefficients. *Operations research letters*, 33(5):481–485, 2005.
- [14] Virginie Gabrel and Céline Murat. Robustness and duality in linear programming. *Journal of the Operational Research Society*, 61(8):1288–1296, 2010.
- [15] Rui Gao and Anton Kleywegt. Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research*, 48(2):603–655, 2023.
- [16] Joudi Hajar, Taylan Kargin, and Babak Hassibi. Wasserstein distributionally robust regret-optimal control under partial observability. In *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1–6. IEEE, 2023.
- [17] Grani A Hanasusanto, Daniel Kuhn, and Wolfram Wiesemann. A comment on “computational complexity of stochastic programming problems”. *Mathematical Programming*, 159:557–569, 2016.
- [18] Taylan Kargin, Joudi Hajar, Vikrant Malik, and Babak Hassibi. Wasserstein distributionally robust regret-optimal control over infinite-horizon. In *6th Annual Learning for Dynamics & Control Conference*, pages 1688–1701. PMLR, 2024.
- [19] Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations research & management science in the age of analytics*, pages 130–166. Informs, 2019.
- [20] Olvi L Mangasarian and T-H Shiau. A variable-complexity norm maximization problem. *SIAM Journal on Algebraic Discrete Methods*, 7(3):455–461, 1986.
- [21] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.
- [22] Karthik Natarajan, Dongjian Shi, and Kim-Chuan Toh. A probabilistic model for minmax regret in combinatorial optimization. *Operations Research*, 62(1):160–181, 2014.
- [23] Georgia Perakis and Guillaume Roels. Regret in the newsvendor model with partial information. *Operations research*, 56(1):188–203, 2008.

- [24] Mehran Poursoltani and Erick Delage. Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems. *Operations Research*, 70(5):2906–2930, 2022.
- [25] Mehran Poursoltani, Erick Delage, and Angelos Georghiou. Risk-averse regret minimization in multistage stochastic programs. *Operations Research*, 72(4):1727–1738, 2024.
- [26] Leonard J Savage. The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67, 1951.
- [27] Jinfeng Yue, Bintong Chen, and Min-Chiang Wang. Expected value of distribution information for the newsvendor problem. *Operations research*, 54(6):1128–1136, 2006.
- [28] Man-Chung Yue, Daniel Kuhn, and Wolfram Wiesemann. On linear optimization over wasserstein balls. *Mathematical Programming*, 195(1):1107–1122, 2022.
- [29] Zhisu Zhu, Jiawei Zhang, and Yinyu Ye. Newsvendor optimization with limited distribution information. *Optimization methods and software*, 28(3):640–667, 2013.