IEEE Control Systems Letters paper presented at
2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Wasserstein Distributionally Robust Regret Minimization

Youngchae Cho Insoon Yang

***Abstract*— We propose a decision-making model under uncertainty to minimize the *ex-ante* regret in a distributionally robust manner using the Wasserstein metric, where the regret is defined as the difference of the expected cost achieved and the best achievable expected cost for any given distribution of uncertainty. First, we formulate a minimization problem of the worst-case *ex-ante* regret over a Wasserstein ball. Subsequently, we derive its surrogate as the proposed model using the minimax inequality, whose objective function is above the worst-case *ex-ante* regret on the decision space. Our main contributions are to show that (i) the approximation error of our model is uniformly bounded, and it vanishes depending on the cost function, the uncertainty set, and the Wasserstein ball's radius; (ii) our model provides a finite-sample performance guarantee and is asymptotically optimal; and (iii) an optimal solution of our model for a class of two-stage linear programs can be obtained using the cutting-plane method. Simulation results demonstrate the effectiveness of our model.**

***Index Terms*— Optimization, distributional robustness.**

I. Introduction

Stochastic programming is a fundamental mathematical framework for decision-making under uncertainty. In the pursuit of cost reduction, risk-neutral decision-makers seek to minimize the expected cost. However, this is impossible in most cases as the underlying distribution of uncertainty is unknown. Nevertheless, access to sample data of uncertainty allows for making decisions with desired qualities.

Sample average approximation (SAA) is one of the simplest data-driven methods, where the underlying distribution is replaced by an empirical distribution [1], [2]. Although conceptually straightforward, SAA may be unreliable if only a few sample data are used. Robust optimization (RO) could be a safe alternative in such a case, typically minimizing the maximum possible cost for a (possibly data-driven) uncertainty set [3], [4]. However, RO is criticized for being too conservative, as it considers only worst-case scenarios.

Distributionally robust optimization (DRO) is yet another prominent data-driven approach, where worst-case analyses are conducted over a family of distributions referred to as an *ambiguity set* to mitigate the impact of potential misspecifications in the underlying distribution. The objective function of a DRO problem in the context of cost minimization is defined as the worst-case expected cost. The mathematical properties of a DRO problem heavily depend on the ambiguity set, which can be defined using, for example, moment conditions [5], total variation distances [6] and Wasserstein distances [7]–[9]. The size of an ambiguity set is adjustable, which makes DRO more robust than SAA and less conservative than RO.

Regret minimization also stands as a renowned paradigm of decision-making under uncertainty. To be precise in the subsequent discussions, we define and differentiate two regret concepts, *ex-ante* and *ex-post* regret. On the one hand, we employ the term "*ex-ante* regret" for a decision and a *distribution* of uncertainty to describe the difference of the expected cost incurred by the decision and the infimum of achievable expected costs (cf. [10]).[1] On the other hand, the "*ex-post* regret" is defined for a decision and a *realization* of uncertainty and indicates the difference of the cost incurred by the decision and the infimum of achievable costs.

Attributed to the seminal publication [11], regret minimization strategies combined with RO and DRO have demonstrated the capacity to reduce conservatism in their counterparts focused on minimizing costs, when the latter approaches could even lead to pessimistic decisions. For example, worst-case *ex-post* regret minimization models based on RO are studied for linear programs (LPs) [12] and two-stage LPs [13]. In [14], a worst-case *ex-ante* regret minimization model is investigated for machine learning, which reveals that regret is more robustly compared across an ambiguity set than cost. However, the solvable data-driven counterpart of the model cannot account for scenarios out of the sample set. In [10], a comprehensive regret minimization model is developed in a multi-stage setting with a variety of risk measures, which allows the benchmark decision to accurately forecast future information up to an arbitrary number of steps ahead. As such, it incorporates both *ex-ante* and *ex-post* regret minimization models based on DRO. However, the analyses are conducted only for discrete probability spaces. In [15], a linear-quadratic controller is designed for discrete-time linear dynamical systems to minimize the worst-case expected *ex-post* regret over a 2-Wasserstein ball. In addition, [16] and [17] study the newsvendor problem to minimize the worst-case *ex-ante* regret over moment-based ambiguity sets.

To the best of our knowledge, however, the literature lacks an *ex-ante* regret minimization model based on DRO that can be applied across a wide range of probability distributions (without assuming the data-generating distribution to be discrete as [10] or absolutely continuous with respect to the empirical distribution as [14]) and application domains (not in a domain-specific manner as [16] and [17]). Motivated

This work was supported by the Information and Communications Technology Planning and Evaluation (IITP) Grant funded by MSIT under Grant 2022-0-00124 and Grant 2022-0-00480.

The authors are with the Department of Electrical and Computer Engineering, ASRI, Seoul National University, Seoul 08826, South Korea {youngchaecho,insoonyang}@snu.ac.kr

[1] Whereas we define the *ex-ante* regret to assess the suboptimality of a decision in terms of the expected cost for a specified distribution, it can also be defined under various risk measures, as generalized in [10].

by this fact, we propose a novel decision-making model to minimize the *ex-ante* regret in a distributionally robust way.

To this end, we first formulate a minimization problem of the worst-case *ex-ante* regret over a Wasserstein ball, which is referred to as the *exact* model. We use the Wasserstein metric to consider a general class of data-generating distributions, which leads to ambiguity sets that are richer and yet less conservative compared to, for example, moment conditions or ϕ-divergences [9]. Subsequently, as no globally optimal algorithm is readily available for the exact model, we derive its surrogate as the proposed model using the minimax inequality. The objective function of our model is above the worst-case *ex-ante* regret on the decision space.

As our main contributions, we show the following:

- First, the approximation error of our model as a function on the decision space is bounded. The error becomes identically zero depending on the cost function, the uncertainty set, and the Wasserstein ball's radius.
- Second, our model provides an upper confidence bound on the *ex-ante* regret with respect to the true distribution of uncertainty and is asymptotically optimal.
- Third, an optimal point of our model within any given tolerance can be obtained by the cutting-plane method in finite iterations for a class of two-stage LPs.

The rest of this letter is organized as follows: In Section II, we formulate our regret minimization model and analyze its properties concerning approximation error. In Section III, we explain in detail the performance guarantees offered by our model with respect to the underlying distribution of uncertainty. In Section IV, we discuss the tractability of our model. In Section V, we provide simulation results.

Notation: We use $\mathbb{R}$, $\mathbb{R}_+$, $\xi \in \mathbb{R}^m$, Ξ, $\mathbb{P}^\star$, and $\mathbb{E}$ to represent the real number set, the non-negative real number set, the random vector of interest, a knowledge-based uncertainty set of ξ (we have $\xi \in \Xi$ deterministically), the unknown distribution of ξ, and the expectation operator, respectively. We assume that Ξ is closed. We let $\mathcal{P}(\Xi)$ and δ_ξ denote the set of Borel probability measures on Ξ and the Dirac delta distribution centered at any $\xi \in \Xi$, respectively. We assume that a set of N i.i.d. samples $\hat{\xi}_1, \ldots, \hat{\xi}_N$ of ξ is given, which is a random object governed by the N-fold product $\mathbb{P}^{\star N}$ of $\mathbb{P}^\star$. For any vectors $v_1, \ldots, v_N$ of the same dimension, we let $v_{1:N} := (v_1, \ldots, v_N)$. Moreover, we denote by $\mathbf{1}_m$ and $\mathcal{V}(\cdot)$ the m-dimensional vector of ones and the vertex set of any given closed convex polytope, respectively.

II. Proposed Model

Our goal is to minimize the worst-case *ex-ante* regret using Wasserstein DRO. This can be ideally achieved by solving the optimization problem

$$\inf_{x \in \mathcal{X}} \hat{R}^{\mathbf{a}}_{\varepsilon,p}(x), \tag{1}$$

where $x \in \mathbb{R}^n$ and $\mathcal{X}$ represent the decision vector and its non-empty feasible set, respectively. Here, $\hat{R}^{\mathbf{a}}_{\varepsilon,p}(x)$ denotes the worst-case *ex-ante* regret over a p-Wasserstein ball for any $p \in [1, \infty)$, i.e., $\hat{R}^{\mathbf{a}}_{\varepsilon,p}(x) := \sup_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} R^{\mathbf{a}}(x, \mathbb{P})$, where $R^{\mathbf{a}}(x, \mathbb{P}) := \mathbb{E}_{\xi \sim \mathbb{P}}[f(x, \xi)] - \inf_{y \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}}[f(y, \xi)]$ denotes the *ex-ante* regret of choosing $x \in \mathcal{X}$, with a cost function $f : \mathcal{X} \times \Xi \to \mathbb{R}$ of interest, when ξ follows $\mathbb{P} \in \mathcal{P}(\Xi)$. The Wasserstein ball $\hat{\mathcal{P}}_{\varepsilon,p} := \{\mathbb{P} \in \mathcal{P}(\Xi) : W_p(\mathbb{P}, \hat{\mathbb{P}}) \le \varepsilon\}$ is of radius $\varepsilon \in \mathbb{R}_+$ and centered at the empirical distribution $\hat{\mathbb{P}} := \frac{1}{N}\sum_{i=1}^N \delta_{\hat{\xi}_i}$, where $W_p(\mathbb{P}, \mathbb{P}') := \{\inf_{\pi \in \Pi(\mathbb{P}, \mathbb{P}')} \int_{\mathbb{R}^m \times \mathbb{R}^m} \|\xi - \xi'\|^p \pi(d\xi, d\xi')\}^{1/p}$ denotes the p-Wasserstein distance of $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\mathbb{R}^m)$ defined with any norm $\|\cdot\|$ on $\mathbb{R}^m$. Moreover, $\Pi(\mathbb{P}, \mathbb{P}')$ represents the set of all joint distributions of $(\xi, \xi') \in \mathbb{R}^m \times \mathbb{R}^m$ with marginals $\mathbb{P}$ and $\mathbb{P}'$ for ξ and ξ', respectively. Unless stated explicitly otherwise, $\varepsilon \in \mathbb{R}_+$, $p \in [1, \infty)$, and $\|\cdot\|$ are arbitrary.

It is worth contrasting $\hat{R}^{\mathbf{a}}_{\varepsilon,p}$ with the objective functions of the other two Wasserstein DRO approaches mentioned in Section I. To be specific, the worst-case expected cost minimization model (M1) is formulated as

$$\inf_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}}[f(x, \xi)].$$

Further, the worst-case expected *ex-post* regret minimization model (M2) is written as

$$\inf_{x \in \mathcal{X}} \hat{R}^{\mathbf{p}}_{\varepsilon,p}(x) = \inf_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}}[R^{\mathbf{p}}(x, \xi)], \tag{2}$$

where $R^{\mathbf{p}} : \mathcal{X} \times \Xi \to \mathbb{R}_+$, $R^{\mathbf{p}}(x, \xi) := f(x, \xi) - \inf_{y \in \mathcal{X}} f(y, \xi)$ represents the *ex-post* regret of a decision x with respect to a realization of ξ. We refer to the models M1 and M2 as needed, such as when discussing the properties of our model and numerically testing its performance.

Throughout this letter, we impose the following three assumptions on f:

Assumption 1: We have $\inf_{x \in \mathcal{X}, \xi \in \Xi} f(x, \xi) > -\infty$.

Assumption 2: We have $\sup_{x \in \mathcal{X}, \xi \in \Xi} f(x, \xi) < \infty$.

Assumption 3: For any $x \in \mathcal{X}$, $f(x, \cdot)$ is measurable.

We introduce Assumptions 1 and 2 to have the approximation error of our model as a function on $\mathcal{X}$ well-defined; note that if $\sup_{\xi \in \Xi} f(x', \xi)$ is infinite for some $x' \in \mathcal{X}$, so is $\hat{R}^{\mathbf{a}}_{\varepsilon,p}(x')$ for any $\varepsilon > 0$. Assumption 3 is required to reformulate (1) using the strong duality result [8, Th. 7] for Wasserstein DRO that is used in the following paragraph.

For notational simplicity, we define $g : \mathcal{X}^2 \times \Xi \to \mathbb{R}$, $g(x, y, \xi) := f(x, \xi) - f(y, \xi)$ and $\tilde{g} : \mathcal{X}^2 \times \Xi^N \to \mathbb{R}$, $\tilde{g}(x, y, \xi_{1:N}) := \frac{1}{N}\sum_{i=1}^N g(x, y, \xi_i)$. As $g(x, y, \cdot)$ is measurable for any $x, y \in \mathcal{X}$, [8, Th. 7] implies that

$$\hat{R}^{\mathbf{a}}_{\varepsilon,p}(x) = \sup_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}, y \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}}[g(x, y, \xi)] \tag{3}$$

$$= \sup_{y \in \mathcal{X}} \inf_{\lambda \in \mathbb{R}_+} \sup_{\xi_{1:N} \in \Xi^N} \varepsilon^p \lambda + \hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \lambda), \tag{4}$$

where $\hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \lambda) := \tilde{g}(x, y, \xi_{1:N}) - \frac{1}{N}\sum_{i=1}^N \lambda \|\xi_i - \hat{\xi}_i\|^p$. Further, $\lambda \in \mathbb{R}_+$ is the Lagrange multiplier associated with the Wasserstein distance constraint $\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}$ in (3); see the proof of [7, Th. 4.2]. Thus, (1) is rewritten as

$$\inf_{x \in \mathcal{X}} \sup_{y \in \mathcal{X}} \inf_{\lambda \in \mathbb{R}_+} \sup_{\xi_{1:N} \in \Xi^N} \varepsilon^p \lambda + \hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \lambda). \tag{5}$$

To our knowledge, however, no globally optimal algorithm for (5) is readily available in the literature except for trivial cases where $f(x, \xi)$ is additively separable, i.e., $f(x, \xi) =$

$f_1(x) + f_2(\xi)$ for some $f_1 : \mathcal{X} \to \mathbb{R}$ and $f_2 : \Xi \to \mathbb{R}$.[2] [3] Hence, we derive a surrogate of (5) as

$$\inf_{x \in \mathcal{X}} \hat{R}_{\varepsilon,p}(x), \tag{6}$$

which is the *proposed regret minimization model*, where our objective function $\hat{R}_{\varepsilon,p} : \mathcal{X} \to \mathbb{R}$ is defined as

$$\hat{R}_{\varepsilon,p}(x) := \inf_{\lambda \in \mathbb{R}_+} \sup_{y \in \mathcal{X}, \xi_{1:N} \in \Xi^N} \varepsilon^p \lambda + \hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \lambda).$$

The objective function is designed by applying the minimax inequality to (4). Thus, $\hat{R}^{\mathbf{a}}_{\varepsilon,p} \le \hat{R}_{\varepsilon,p}$ by construction. Moreover, (6) can be regarded as an RO problem, which suggests the possibility of applying existing RO algorithms [18]. The tractability of our model is elaborated on in Section IV.

In what follows, we discuss some properties of the proposed model as a surrogate of (1). First, we show that the objective function of our model is below the worst-case expected *ex-post* regret on the decision space.

Proposition 1: For all $x \in \mathcal{X}$, $\hat{R}_{\varepsilon,p}(x) \le \hat{R}^{\mathbf{p}}_{\varepsilon,p}(x)$.

Proof: As $R^{\mathbf{p}}(x, \cdot)$ is measurable for any $x \in \mathcal{X}$, according to [8, Th. 7], $\hat{R}^{\mathbf{p}}_{\varepsilon,p}(x)$ is equal to

$$\inf_{\lambda \in \mathbb{R}_+} \left\{ \varepsilon^p \lambda + \frac{1}{N} \sum_{i=1}^{N} \sup_{\xi \in \Xi, y \in \mathcal{X}} g(x, y, \xi) - \lambda \|\xi - \hat{\xi}_i\|^p \right\}.$$

Since $\frac{1}{N} \sum_{i=1}^{N} \sup_{y \in \mathcal{X}} g(x, y, \xi_i) \ge \sup_{y \in \mathcal{X}} \tilde{g}(x, y, \xi_{1:N})$ for any $(x, \xi_{1:N}) \in \mathcal{X} \times \Xi^N$, the statement holds. ■

As a corollary to Proposition 1, we additionally show that the approximation error of our model is bounded. To be precise, we define $\hat{\Delta}_{\varepsilon,p} : \mathcal{X} \to \mathbb{R}_+$, $\hat{\Delta}_{\varepsilon,p}(x) := \hat{R}_{\varepsilon,p}(x) - \hat{R}^{\mathbf{a}}_{\varepsilon,p}(x)$ to represent the *approximation error*.

Corollary 1: The approximation error is uniformly bounded as follows: $\hat{\Delta}_{\varepsilon,p} \le \inf_{x \in \mathcal{X}} \hat{R}^{\mathbf{p}}_{\varepsilon,p}(x)$.

Proof: For any $x \in \mathcal{X}$, we observe that

$$\begin{aligned}
\hat{\Delta}_{\varepsilon,p}(x) &\le \sup\nolimits_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}}[R^{\mathbf{p}}(x, \xi)] \\
&\qquad - \sup\nolimits_{\mathbb{P}' \in \hat{\mathcal{P}}_{\varepsilon,p}, y \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}'}[g(x, y, \xi)] \\
&= \sup\nolimits_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}}[f(x, \xi) - \inf\nolimits_{y' \in \mathcal{X}} f(y', \xi)] \\
&\qquad - \sup\nolimits_{\mathbb{P}' \in \hat{\mathcal{P}}_{\varepsilon,p}, y \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}'}[f(x, \xi) - f(y, \xi)] \\
&\le \sup\nolimits_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \inf\nolimits_{y \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathbb{P}}[f(y, \xi) - \inf\nolimits_{y' \in \mathcal{X}} f(y', \xi)] \\
&\le \inf\nolimits_{y \in \mathcal{X}} \sup\nolimits_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}}[R^{\mathbf{p}}(y, \xi)],
\end{aligned}$$

where the penultimate and last inequalities follow from the fact that we have $\sup_{\mathbb{P}' \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}'}[f(x, \xi) - f(y, \xi)] \ge \mathbb{E}_{\xi \sim \mathbb{P}}[f(x, \xi) - f(y, \xi)]$ for any $(\mathbb{P}, x, y) \in \hat{\mathcal{P}}_{\varepsilon,N} \times \mathcal{X}^2$ and the minimax inequality, respectively. As $\inf_{x \in \mathcal{X}} \hat{R}^{\mathbf{p}}_{\varepsilon,p}(x)$ is finite under Assumptions 1 and 2, the statement holds. ■

Subsequently, we show that the approximation error can vanish regardless of $x \in \mathcal{X}$, depending on $\tilde{g}$, Ξ, and ε.

Proposition 2: Suppose that either (i) $\sup_{y \in \mathcal{X}} \tilde{g}(x, y, \cdot)$ is upper semicontinuous for any $x \in \mathcal{X}$ and $\varepsilon = 0$ or (ii) $D(\Xi) := \sup_{\xi, \xi' \in \Xi} \|\xi - \xi'\| < \infty$ and $\varepsilon \ge D(\Xi)$. Then, we have $\hat{\Delta}_{\varepsilon,p} = 0$.

Proof: First, we assume that (i) holds. Let $\overline{g} := \sup_{x, y \in \mathcal{X}, \xi_{1:N} \in \Xi^N} \tilde{g}(x, y, \xi_{1:N})$, which is finite due to Assumptions 1 and 2. For any $x \in \mathcal{X}$ and $\gamma > 0$, we have

$$\begin{aligned}
\hat{R}_{0,p}(x) &\le \sup_{y \in \mathcal{X}, \xi_{1:N} \in \Xi^N} \hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \overline{g}/\gamma) \\
&= \sup_{y \in \mathcal{X}, \xi_{1:N} \in \Xi^N} \hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \overline{g}/\gamma) \\
&\qquad \text{s.t. } \frac{1}{N} \sum\nolimits_{i=1}^{N} \|\xi_i - \hat{\xi}_i\|^p \le \gamma,
\end{aligned} \tag{7}$$

where the inequality and equality follow from the definition of $\hat{R}_{\varepsilon,p}$ and the fact that $\hat{R}_{0,p} \ge 0$, respectively. Since $\overline{g} \ge 0$, we have $\hat{R}_{0,p}(x) \le G(x, \gamma)$ for any $x \in \mathcal{X}$ and $\gamma > 0$, where $G(x, \gamma) := \sup_{y \in \mathcal{X}, \xi_{1:N} \in \Xi^N} \{\tilde{g}(x, y, \xi_{1:N}) : (7)\}$. If $\hat{\Delta}_{0,p}(x') > 0$ for some $x' \in \mathcal{X}$, then $\hat{R}_{0,p}(x') \in (\hat{R}^{\mathbf{a}}_{0,p}(x'), G(x', \gamma)]$ for any $\gamma > 0$. This is contradictory because if $\sup_{y \in \mathcal{X}} \tilde{g}(x', y, \cdot)$ is upper semicontinuous, then $\lim_{\gamma \to 0^+} G(x', \gamma) = \hat{R}^{\mathbf{a}}_{0,p}(x')$. Thus, we have $\hat{\Delta}_{0,p} = 0$. Next, we assume that (ii) holds. We have for any $x \in \mathcal{X}$ that

$$\begin{aligned}
\hat{R}^{\mathbf{a}}_{\varepsilon,p}(x) &= \sup_{y \in \mathcal{X}} \inf_{\lambda \in \mathbb{R}_+} \sup_{\xi_{1:N} \in \Xi^N} \varepsilon^p \lambda + \hat{R}^{\mathbf{s}}_p(x, y, \xi_{1:N}, \lambda) \\
&= \sup_{\xi \in \Xi, y \in \mathcal{X}} g(x, y, \xi) = \sup_{\mathbb{P} \in \hat{\mathcal{P}}_{\varepsilon,p}} \mathbb{E}_{\xi \sim \mathbb{P}}[R^{\mathbf{p}}(x, \xi)] = \hat{R}^{\mathbf{p}}_{\varepsilon,p}(x),
\end{aligned}$$

where the second equality holds as $\frac{1}{N} \sum_{i=1}^{N} \|\xi - \hat{\xi}_i\|^p \le \varepsilon^p$ for any $\xi_{1:N} \in \Xi^N$. Since $\hat{R}^{\mathbf{a}}_{\varepsilon,p} \le \hat{R}_{\varepsilon,p} \le \hat{R}^{\mathbf{p}}_{\varepsilon,p}$, it is further deduced that $\hat{R}^{\mathbf{a}}_{\varepsilon,p} = \hat{R}_{\varepsilon,p}$, i.e., $\hat{\Delta}_{\varepsilon,p} = 0$. ■

Proposition 2 suggests that our model can cover both the *ex-ante* regret minimization model based on SAA,

$$\inf_{x \in \mathcal{X}} R^{\mathbf{a}}(x, \hat{\mathbb{P}}), \tag{8}$$

and the worst-case *ex-post* regret minimization model based on RO, $\inf_{x \in \mathcal{X}} \sup_{\xi \in \Xi} R^{\mathbf{p}}(x, \xi)$. The solution set of (8), if it exists, is identical to that of the expected cost minimization model based on SAA, $\inf_{x \in \mathcal{X}} \hat{f}_N(x) := \mathbb{E}_{\xi \sim \hat{\mathbb{P}}}[f(x, \xi)]$. However, for any $\varepsilon > 0$, solving either (1) or (6) typically does not provide a solution to M1. We numerically compare our model and M1 in Section V. In the following section, we analyze performance guarantees offered by our model.

III. Performance Guarantees

Minimizing the worst-case expected cost over a Wasserstein ball can provide both finite-sample and asymptotic performance guarantees [7], [8]. In this section, we show that our regret minimization model can similarly offer two performance guarantees, assuming that a solution to it is attainable. We denote any solution and the optimal value of (6) by $\hat{x}^\star_{\varepsilon,p,N}$ and $\hat{R}^\star_{\varepsilon,p,N}$, respectively. Further, we define $R : \mathcal{X} \to \mathbb{R}_+$, $R(x) := R^{\mathbf{a}}(x, \mathbb{P}^\star)$, to represent the true *ex-ante* regret. Note that we have $R^\star := \inf_{x \in \mathcal{X}} R(x) = 0$.

First, we prove that the *ex-ante* regret with respect to the underlying distribution of uncertainty can be bounded above with high probability using our model.

[2]In such cases, we can obtain a solution to (1) by solving the optimization problem $\inf_{x \in \mathcal{X}} f_1(x)$ without uncertainty as $g(x, y, \xi) = f_1(x) - f_1(y)$ is independent of ξ. Thus, depending on the tractability of $\inf_{x \in \mathcal{X}} f_1(x)$, we can obtain a solution to (1) in a computationally efficient way.

[3]While the choice between the *ex-ante* and *ex-post* regret in the context of distributionally robust regret minimization depends on the decision-maker's preference, (1) is computationally more challenging than M2 as only M2 is known to admit an inf-sup reformulation (see the proof of Proposition 1).

Proposition 3: Suppose that $\mathbb{E}_{\xi\sim\mathbb{P}^\star}[\exp(\|\xi\|^\alpha)]<\infty$ for some $\alpha>p$ and $p\neq m/2$. For any $\beta\in(0,1]$ and $N\geq 1$, there exists an $\varepsilon>0$ such that $\mathbb{P}^{\star N}[R(\hat{x}^\star_{\varepsilon,p,N})\leq\hat{R}^\star_{\varepsilon,p,N}]\geq 1-\beta$.

Proof: According to [8, Th. 18], there exists some $\varepsilon>0$ as a function of β and N such that $\hat{\mathcal{P}}_{\varepsilon,p}$ contains $\mathbb{P}^\star$ with probability at least $1-\beta$, i.e., $\mathbb{P}^{\star N}\left[\mathbb{P}^\star\in\hat{\mathcal{P}}_{\varepsilon,p}\right]\geq 1-\beta$.[4] This implies that $\mathbb{P}^{\star N}[R\leq\hat{R}^{\mathbf{a}}_{\varepsilon,p}]\geq 1-\beta$. As we have $\hat{R}^{\mathbf{a}}_{\varepsilon,p}\leq\hat{R}_{\varepsilon,p}$, the statement holds. ■

Subsequently, we show that our model is asymptotically optimal as the sample size tends to infinity. We define $h:\mathcal{X}\to\mathbb{R}$, $h(x):=\mathbb{E}_{\xi\sim\mathbb{P}^\star}[f(x,\xi)]$. Further, we let $\hat{f}^\star_N:=\inf_{x\in\mathcal{X}}\hat{f}_N(x)$ and $h^\star:=\inf_{x\in\mathcal{X}}h(x)$.

Proposition 4: Suppose that $g(x,y,\cdot)$ is Lipschitz continuous for any $x,y\in\mathcal{X}$ and $\hat{f}_N$ uniformly converges to h a.s. Then, we have

$$R\left(\hat{x}^\star_{\varepsilon_N,p,N}\right)\to R^\star \text{ a.s.} \quad \text{as } N\to\infty$$

for any sequence $\varepsilon_N\geq 0$ such that $\varepsilon_N\to 0$.

Proof: To explicitly express the dependence of $\hat{R}_{\varepsilon,p}$ and $\hat{R}^{\mathbf{a}}_{\varepsilon,p}$ on N, we use $\hat{R}_{\varepsilon,p,N}$ and $\hat{R}^{\mathbf{a}}_{\varepsilon,p,N}$ to represent them in this proof, respectively. Let $\overline{L}:=\sup_{x,y\in\mathcal{X}}L(x,y)$ where $L(x,y)$ denotes the Lipschitz constant of $g(x,y,\cdot)$. As $\hat{R}_{\varepsilon_N,p,N}(x)\leq\varepsilon_N^p\overline{L}+\sup_{y\in\mathcal{X},\xi_{1:N}\in\Xi^N}\hat{R}^{\mathbf{s}}_p\left(x,y,\xi_{1:N},\overline{L}\right)=\varepsilon_N^p\overline{L}+\sup_{y\in\mathcal{X}}\tilde{g}\left(x,y,\hat{\xi}_{1:N}\right)$ for any $x\in\mathcal{X}$, we have

$$\hat{R}_{\varepsilon_N,p,N}\leq\hat{R}^{\mathbf{a}}_{0,p,N}+\varepsilon_N^p\overline{L}, \tag{9}$$

which leads to $\sup_{x\in\mathcal{X}}|\hat{R}_{\varepsilon_N,p,N}(x)-R(x)|\leq\varepsilon_N\overline{L}+\sup_{x\in\mathcal{X}}|\hat{R}^{\mathbf{a}}_{0,p,N}(x)-R(x)|$. Meanwhile, $\hat{R}^{\mathbf{a}}_{0,p,N}$ uniformly converges to R a.s., because $\sup_{x\in\mathcal{X}}|\hat{R}^{\mathbf{a}}_{0,p,N}(x)-R(x)|\leq\sup_{x\in\mathcal{X}}|\hat{f}_N(x)-h(x)|+|\hat{f}^\star_N-h^\star|$, $\hat{f}_N$ uniformly converges to h a.s., and $\hat{f}^\star_N\to h^\star$ a.s. as indicated by the convergence of $\hat{f}_N$ [20, Proposition 5.2]. Thus, $\hat{R}_{\varepsilon_N,p,N}$ uniformly converges to R a.s., implying that $R(\hat{x}^\star_{\varepsilon_N,p,N})\to\hat{R}^\star_{\varepsilon_N,p,N}$ a.s. Taking the infimum of both sides in (9) for all $x\in\mathcal{X}$ yields $\hat{R}^\star_{\varepsilon_N,p,N}\leq\varepsilon_N^p\overline{L}$. This suggests that $\hat{R}^\star_{\varepsilon_N,p,N}\to 0$ a.s., as $0\leq\hat{R}^\star_{\varepsilon_N,p,N}$ by definition. Consequently, we have $R(x^\star_{\varepsilon_N,p,N})\to 0$ a.s. Since $R^\star{=}0$, the proof is complete. ■

For detailed descriptions of conditions for the convergence of $\hat{f}_N$ assumed in Proposition 4, the reader is referred to [2] and the references therein. In the following section, we discuss the tractability of our model.

IV. Solution Method

Our model (6) can be regarded as an RO problem to determine (x,λ) under uncertainty of $(y,\xi_{1:N})$, which is thus intractable in general [3].[5] For example, suppose that $\mathcal{X}\subset\mathbb{R}$ and Ξ are compact convex polytopes, $f(x,\xi)=|x-\xi^\top M\xi|$ for any positive-definite matrix $M\in\mathbb{R}^{m\times m}$, $\{\xi^\top M\xi:\xi\in\Xi\}\subseteq\mathcal{X}$ and $\varepsilon=D(\Xi)$. As Proposition 2 holds and $\inf_{x\in\mathcal{X},\xi\in\Xi}f(y,\xi)=0$, $\hat{x}^\star_{\varepsilon,p,N}=\frac{1}{2}\sup_{\xi\in\Xi}\xi^\top M\xi$ uniquely solves (6). Notably, maximizing the quadratic form $\xi^\top M\xi$ in ξ is NP-hard and intractable [21]. Nonetheless, even if it is intractable, a solution to an RO problem may be obtained by the cutting-plane method [22]. In this section, we show that despite the general intractability of (6), an optimal point within any given tolerance can be obtained using the cutting-plane method in finite iterations. This achievement is contingent on the conditions of $\mathcal{X}$, Ξ, f, p, and $\|\cdot\|$, but not ε. We first describe how the cutting-plane method can tackle (6) in general, assuming that a solution to it is attainable.

A. General Idea: Cutting-Plane Approach

By the cutting-plane method, (6) is decomposed into a master problem and a subproblem that are iteratively solved. The master problem is a relaxation of (6), which can be generally written as

$$\inf_{x\in\mathcal{X},\lambda\in\mathbb{R}_+,\eta\in\mathbb{R}}\left\{\varepsilon^p\lambda+\eta:(x,\lambda,\eta)\in\mathcal{M}_K\right\}, \tag{10}$$

where $\mathcal{M}_K$ is a convex polytope defined at each iteration step $K\geq 1$ such that $\mathcal{M}_K\supseteq\mathcal{M}:=\{(x,\lambda,\eta)\in\mathbb{R}^{n+2}:\eta\geq\hat{R}^{\mathbf{s}}_p(x,y,\xi_{1:N},\lambda)\ \forall(y,\xi_{1:N})\in\mathcal{X}\times\Xi^N\}$. Note that the solution set of (10) associated with (x,λ) for $\mathcal{M}_K=\mathcal{M}$ is identical to that of (6). For any $(x,\lambda)\in\mathcal{X}\times\mathbb{R}_+$, we consider the optimization problem

$$\sup_{y\in\mathcal{X},\xi_{1:N}\in\Xi^N}\hat{R}^{\mathbf{s}}_p(x,y,\xi_{1:N},\lambda). \tag{11}$$

Given any solution $\left(x^{(K)},\lambda^{(K)},\eta^{(K)}\right)$ to the master problem (10), if it exists, the subproblem is formulated to assess its optimality as (11) for $(x,\lambda)=(x^{(K)},\lambda^{(K)})$. Let U'_K denote the optimal value of the subproblem. Then, $L_K:=\varepsilon^p\lambda^{(K)}+\eta^{(K)}$ and $U_K:=\varepsilon^p\lambda^{(K)}+U'_K$ are respectively a lower bound and an upper bound of $\hat{R}^\star_{\varepsilon,p,N}$ associated with $x=x^{(K)}$. As such, if a convergence criterion based on L_K and U_K is met, e.g., $U_K-L_K\leq\rho$ for a given tolerance $\rho\geq 0$, then $x^{(K)}$ is returned as our decision and the iteration stops. Otherwise, a linear constraint, often referred to as a *cut*, is added to the master problem to refine $\mathcal{M}_K$ and obtain $\mathcal{M}_{K+1}$.

In principle, an optimal solution to (6) can be attained by the cutting-plane method in finitely many iterations if we can build $\mathcal{M}_K$, such that the solution sets of (10) and (6) associated with (x,λ) are guaranteed to be equal for some $K<\infty$, and obtain a solution to (10) for any $K<\infty$. As a special case, we show in the next subsection that our model for two-stage linear programming under right-hand side uncertainty or objective uncertainty is solvable in this sense, if an oracle based on the simplex method is available for exactly solving LPs and mixed-integer LPs (MILPs). Notably, two-stage linear programming highlights its practical relevance across various applications [23]–[25].

B. Special Case: Two-Stage Linear Programming

In this subsection, we make the following assumption in addition to Assumptions 1–3:

Assumption 4: (i) $\mathcal{X}$ is a compact convex polytope, and Ξ is a closed convex polytope (ii) For any $(x,\xi)\in\mathcal{X}\times\Xi$, $f(x,\xi)$ is equal to the optimal value of the LP to determine a decision vector z,

$$\inf_{z\in\mathbb{R}^{n_2}_+}\left\{(A\xi+a)^\top z:Bz\geq Cx+E\xi+b\right\}, \tag{12}$$

[4] Weaker conditions for this inequality to hold can be found in [19, Th. 2].

[5] We describe an optimization problem as *intractable* if no polynomial-time algorithm is available for it under the common assumption P≠NP [4].

where either $A \in \mathbb{R}^{n_2 \times m}$ or $E \in \mathbb{R}^{r \times m}$ is zero while $B \in \mathbb{R}^{r \times n_2}$, $C \in \mathbb{R}^{r \times n}$, $a \in \mathbb{R}^{n_2}$, and $b \in \mathbb{R}^r$. (iii) The Wasserstein distance is of order 1 and defined with the 1- or ∞-norm, i.e., $p = 1$ and $\|\cdot\| \in \{|\cdot|, \|\cdot\|_\infty\}$.

In what follows, we show that an optimal point of (6) within any given tolerance can be obtained using the cutting-plane method with the oracle in finite iterations under Assumption 4 when A is zero. We can similarly prove that the same holds when E is zero (see [13]).[6] Without loss of generality, we let $\mathcal{X} \subseteq \mathbb{R}^n_+$. As for Assumption 4-(iii), we first describe the method for the case with the 1-norm and briefly explain how it extends to the case with the ∞-norm.

The dual of (12) is written as the LP $\sup_{\nu \in \mathcal{N}} (Cx + E\xi + b)^\top \nu$ where ν and $\mathcal{N} := \{\nu \in \mathbb{R}^r_+ : B^\top \nu \leq a\}$ denote the dual decision vector and its feasible set, respectively. Thus, introducing auxiliary variables $q_i^+, q_i^- \in \mathbb{R}^m_+$ for each $i \leq N$ such that $\xi_i = \hat{\xi}_i + q_i^+ - q_i^-$, we can rewrite (11) as

$$\sup_{w \in \mathcal{W}, \nu_{1:N} \in \mathcal{N}^N} \phi\left(w, \nu_{1:N}, x, \lambda\right), \tag{13}$$

where we define $w := \left(y, z_{1:N}, q_{1:N}^+, q_{1:N}^-\right)$,

$$\begin{aligned} \mathcal{W} := \{ &\left(y, z_{1:N}, q_{1:N}^+, q_{1:N}^-\right) \in \mathcal{X} \times \mathbb{R}_+^{n_2 N + 2mN} : \\ &- Cy + Bz_i - E\left(q_i^+ - q_i^-\right) \geq E\hat{\xi}_i + b \quad \forall i \leq N, \\ &\hat{\xi}_i + q_i^+ - q_i^- \in \Xi \quad \forall i \leq N\}, \end{aligned}$$

and $\phi(w, \nu_{1:N}, x, \lambda) := \frac{1}{N}\sum_{i=1}^N \{-\lambda \mathbf{1}_m^\top (q_i^+ + q_i^-) - a^\top z_i + [Cx + E(\hat{\xi}_i + q_i^+ - q_i^-) + b]^\top \nu_i\}$. For any $(x, \lambda) \in \mathcal{X} \times \mathbb{R}_+$, (13) can be written as an LP in canonical form with respect to w for a fixed $\nu_{1:N}$ and vice versa. Thus, (13) is rewritten as

$$\sup_{w \in \mathcal{V}(\mathcal{W}), \nu_{1:N} \in \mathcal{V}(\mathcal{N}^N)} \phi\left(w, \nu_{1:N}, x, \lambda\right). \tag{14}$$

As a result, we can reformulate (6) equivalently as the LP

$$\inf_{x \in \mathcal{X}, \lambda \in \mathbb{R}_+, \eta \in \mathbb{R}} \left\{\varepsilon\lambda + \eta : (x, \lambda, \eta) \in \overline{\mathcal{M}}\right\}, \tag{15}$$

where $\overline{\mathcal{M}} := \{(x, \lambda, \eta) \in \mathbb{R}^{n+2} : \eta \geq \phi(w, \nu_{1:N}, x, \lambda)$ $\forall (w, \nu_{1:N}) \in \mathcal{V}(\mathcal{W}) \times \mathcal{V}(\mathcal{N}^N)\}$ is a convex polytope. Since it may be impractical to identify all the vertices of $\mathcal{W}$ and $\mathcal{N}^N$, we address (15) by the cutting-plane method and decompose it into the master problem (10) and the subproblem (14). With this aim, we let $\mathcal{M}_K = \{(x, \lambda, \eta) \in \mathbb{R}^{n+2} : \eta \geq \phi\left(w^{(k)}, \nu_{1:N}^{(k)}, x, \lambda\right) \forall k = 0, \ldots, K-1\}$ with any $\left(w^{(0)}, \nu_{1:N}^{(0)}\right) \in \mathcal{W} \times \mathcal{N}^N$ such that $q_{1:N}^+$ and $q_{1:N}^-$ are zero. For each $k \geq 1$, $\left(w^{(k)}, \nu_{1:N}^{(k)}\right)$ denotes any solution to (14) for $(x, \lambda) = \left(x^{(k)}, \lambda^{(k)}\right)$. We impose the initialization condition on $\left(w^{(0)}, \nu_{1:N}^{(0)}\right)$ to make (10) bounded. As (15) is a finite LP with $\overline{\mathcal{M}}$ defined by $|\mathcal{V}(\mathcal{W})||\mathcal{V}(\mathcal{N})|^N$ hyperplanes each of which corresponds to a cut, the iteration terminates in $O\left(|\mathcal{V}(\mathcal{W})||\mathcal{V}(\mathcal{N})|^N\right)$ steps [26, Proposition 1].

However, the subproblem (14) is not easy to solve due to the bilinear terms $\nu_i^\top E q_i^+$'s and $\nu_i^\top E q_i^-$'s in the objective function. To add a valid cut to the master problem, we find a vertex $w^{(K)}$ of $\mathcal{W}$ and that $\nu_{1:N}^{(K)}$ of $\mathcal{N}^N$ separately as follows: First, we obtain $\nu_{1:N}^{(K)}$ by solving an MILP equivalent of (13). This MILP can be acquired by explicitly imposing the Karush–Kuhn–Tucker condition for (12) and applying the Big-M method (see [13] and [26] for details). Subsequently, we obtain $w^{(K)}$ by solving (13) with ν fixed to $\nu_{1:N}^{(K)}$, which is an LP. Using the solution to the MILP reformulation of (13), we can locate a point $w' \in \mathcal{W}$ such that $\left(w', \nu_{1:N}^{(K)}\right)$ attains the optimal value of (13). This ensures the existence of $w^{(K)}$. Note that all the problems in the iteration process are an LP or MILP. Hence, we conclude that our model (6) is solvable.

When the ∞-norm is adopted instead of the 1-norm, the cutting-plane method can be applied similarly after replacing w, $\mathcal{W}$, and $\phi\left(w, \nu_{1:N}, x, \lambda\right)$ with $\tilde{w} := (w, \zeta_{1:N})$, $\tilde{\mathcal{W}} := \{(w, \zeta_{1:N}) \in \mathcal{W} \times \mathbb{R}_+^N : \zeta_i \mathbf{1}_m \geq q_i^+, \zeta_i \mathbf{1}_m \geq q_i^- \ \forall i \leq N\}$, and $\tilde{\phi}\left(\tilde{w}, \nu_{1:N}, x, \lambda\right) := \frac{1}{N}\sum_{i=1}^N \{-\lambda\zeta_i - a^\top z_i + [Cx + E(\hat{\xi}_i + q_i^+ - q_i^-) + b]^\top \nu_i\}$, respectively.

V. Numerical Experiment

We empirically examine the performance of our model on the newsvendor problem to determine the inventory level $x \in \mathcal{X} = [0, 100]$ before the demand $\xi \in \Xi = [0, 100]$ is revealed.[7] The cost function is defined as $f\left(x, \xi\right) = cx - q\min\left\{x, \xi\right\}$ where c=1 and q=5 denote the unit cost and price, respectively. Furthermore, we let $p = 1$ and choose the 1-norm, i.e., $\|\cdot\| = |\cdot|$. Note that $f\left(x, \xi\right)$ can be written as the optimal value of a two-stage LP in the form discussed in Section IV-B, for which our model is solvable with the cutting-plane method. Specifically, we have $f\left(x, \xi\right) = \inf_{z_1, z_2 \in \mathbb{R}_+}\{z_1 - z_2 : z_1 - z_2 \geq (c - q)x, z_1 - z_2 \geq cx - q\xi\}$. We model the underlying distribution of ξ as the normal distribution with mean μ=20, 80 and standard deviation 30 truncated over Ξ. Given a randomly generated sample set of size $N = 10$, we compare our model against M1 and M2[8] for $\varepsilon = eD\left(\Xi\right)$ with $e = 0, 0.0001, 0.001, \ldots, 1$ in terms of the maximum *ex-ante* regret over the Wasserstein ball and the *ex-ante* regret with respect to the underlying distribution. As the maximum *ex-ante* regret is difficult to calculate exactly, we evaluate an upper bound (UB) and a lower bound (LB) on it. Specifically, for any decision $x \in \mathcal{X}$ yielded by each model, the UB and LB are computed as $\hat{R}_{\varepsilon,p}\left(x\right)$ and $\sup_{\mathbb{P} \in \hat{\mathcal{P}}^{\mathrm{e}}_{\varepsilon,p}} R^{\mathbf{a}}_{\varepsilon,p}\left(x, \mathbb{P}\right)$, respectively, where $\hat{\mathcal{P}}^{\mathrm{e}}_{\varepsilon,p} \subseteq \hat{\mathcal{P}}_{\varepsilon,p}$ contains only *perturbed* empirical distributions (see [8, Th. 6]). Moreover, the true *ex-ante* regret is computed approximately using 100,000 randomly generated scenarios of ξ. For statistical robustness, we repeat the experiment with 100 independent sample sets.

Tables I and II show the results averaged over the 100 simulation runs. From Table I, we observe that the average UBs associated with our model are strictly lower than the average LBs associated with the others when $e \in (0, 1)$; the three models have the same solution for e=0, while our model and M1 coincide when $e = 1$ by construction. This implies that our model can effectively mitigate unexpected

[6] However, it is unclear whether the cutting-plane method for our model can converge under any other condition than Assumption 4.

[7] The source code of our implementation is available online: https://github.com/CORE-SNU/DR-Regret.

[8] For the newsvendor problem of our interest, M1 admits an LP reformulation [7, Corollary 5.1] while M2, which is of the inf-sup form (see the proof of Proposition 1), can be solved in the same manner as our model.

TABLE I

AVERAGE UPPER AND LOWER BOUNDS ON MAXIMUM *Ex-Ante* REGRET

(e,μ)		(0,20)	$(0.0001, 20)$	$(0.001, 20)$	$(0.01, 20)$	$(0.1, 20)$	(1,20)
M1	UB	**0.0000**	0.0499	0.4954	4.3375	48.0910	**80.0000**
	LB	0.0000	0.0495	0.4908	4.0923	47.0278	80.0000
M2	UB	**0.0000**	0.0500	0.4966	4.3611	33.5252	400.0000
	LB	0.0000	0.0500	0.4958	4.1357	30.5143	400.0000
Ours	UB	**0.0000**	**0.0307**	**0.3067**	**3.0670**	**29.6306**	**80.0000**
	LB	0.0000	0.0000	0.0362	2.2741	26.2649	80.0000
(e,μ)		(0,80)	$(0.0001, 80)$	$(0.001, 80)$	$(0.01, 80)$	$(0.1, 80)$	(1,80)
M1	UB	**0.0000**	0.0500	0.4815	4.1375	29.5060	**80.0000**
	LB	0.0000	0.0500	0.4789	3.8864	29.5060	80.0000
M2	UB	**0.0000**	0.0500	0.4819	3.9024	27.1683	400.0000
	LB	0.0000	0.0500	0.4789	3.6513	26.4267	400.0000
Ours	UB	**0.0000**	**0.0282**	**0.2824**	**2.8142**	**23.1424**	**80.0000**
	LB	0.0000	0.0004	0.0726	2.4572	22.5625	80.0000

TABLE II

AVERAGE TRUE *Ex-Ante* REGRET

(e,μ)	(0,20)	$(0.0001, 20)$	$(0.001, 20)$	$(0.01, 20)$	$(0.1, 20)$	(1,20)
M1	**4.2661**	3.5083	3.5083	3.5083	14.5975	**15.7801**
M2	**4.2661**	4.1674	4.1674	4.1674	4.0248	96.4340
Ours	**4.2661**	**2.4985**	**2.4985**	**2.4985**	**2.2567**	**15.7801**
(e,μ)	(0,80)	$(0.0001, 80)$	$(0.001, 80)$	$(0.01, 80)$	$(0.1, 80)$	(1,80)
M1	**2.9320**	2.5129	2.5530	2.6827	2.1488	**2.3963**
M2	**2.9320**	2.6402	2.6402	2.6479	2.8457	243.8004
Ours	**2.9320**	**2.4033**	**2.4033**	**2.4021**	**1.1958**	**2.3963**

surges in *ex-ante* regret compared to M1 and M2. Table II indicates that we can also reduce the *ex-ante* regret with respect to the underlying distribution of uncertainty by properly adjusting the Wasserstein ball's radius. However, finding an optimal radius for our model is a challenging task. Although cross-validation can reportedly yield a satisfactory radius for many Wasserstein DRO problems [7], [9], it involves dealing with multiple Wasserstein DRO problems and is thus computationally demanding. One possible alternative would be to extend the statistical inference method [27] designed with the focus on Wasserstein DRO problems for minimizing costs to our setting, which we leave as a future work.

VI. CONCLUDING REMARKS

We presented a novel distributionally robust *ex-ante* regret minimization model, whose objective function approximates the worst-case *ex-ante* regret over a Wasserstein ball. We proved that its approximation error is bounded, which even becomes zero under certain conditions. Furthermore, it offers performance guarantees on the true *ex-ante* regret. While a solution to our model for a class of two-stage LPs can be obtained by the cutting-plane method, it is unclear if this is the only class with a convergence guarantee. Moreover, it remains unanswered whether our model can be solved using alternative techniques. Hence, exploring the tractability of our model across a broader range of decision spaces, uncertainty sets, and cost functions is an important future research direction.

REFERENCES

[1] A. J. Kleywegt, A. Shapiro, and T. Homem-deMello, "The sample average approximation method for stochastic discrete optimization," *SIAM J. Optim.*, vol. 12, no. 2, pp. 479–502, 2002.

[2] S. Kim, R. Pasupathy, and S. G. Henderson, "A guide to sample average approximation," in *Handbook of Simulation Optimization*, M. C. Fu, Ed. New York, NY, USA: Springer, 2015, pp. 207–243.

[3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*, vol. 28. Princeton, NJ, USA: Princeton Univ., 2009.

[4] D. Bertsimas, D. B. Brown, and C. Caramanis, "Theory and applications of robust optimization," *SIAM Rev.*, vol. 53, no. 3, pp. 464–501, 2011.

[5] E. Delage and Y. Ye, "Distributionally robust optimization under moment uncertainty with application to data-driven problems," *Oper. Res.*, vol. 58, no. 3, pp. 595–612, 2010.

[6] F. Farokhi, "Distributionally-robust optimization with noisy data for discrete uncertainties using total variation distance," *IEEE Control Syst. Lett.*, vol. 7, pp. 1494–1499, 2023.

[7] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Math. Program.*, vol. 171, nos. 1–2, pp. 115–166, 2018.

[8] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, "Wasserstein distributionally robust optimization: Theory and applications in machine learning," *Operations Research & Management Science in the Age of Analytics* (TutORials in Operations Research). Catonsville, MD, USA: INFORMS, 2019, pp. 130–166.

[9] R. Gao and A. Kleywegt, "Distributionally robust stochastic optimization with Wasserstein distance," *Math. Oper. Res.*, vol. 48, no. 2, pp. 603–655, 2023.

[10] M. Poursoltani, E. Delage, and A. Georghiou, "Risk-averse regret minimization in multistage stochastic programs," *Oper. Res.*, pp. 1–26, Jan. 2023.

[11] L. J. Savage, "The theory of statistical decision," *J. Am. Stat. Assoc.*, vol. 46, no. 253, pp. 55–67, 1951.

[12] D. Bertsimas and I. Dunning, "Relative robust and adaptive optimization," *INFORMS J. Comput.*, vol. 32, no. 2, pp. 408–427, 2020.

[13] M. Poursoltani and E. Delage, "Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems," *Oper. Res.*, vol. 70, no. 5, pp. 2906–2930, 2022.

[14] A. Agarwal and T. Zhang, "Minimax regret optimization for robust machine learning under distribution shift," in *Proc. 35th Annu. Conf. Learn. Theory*, 2022, pp. 2704–2729.

[15] F. Al Taha, S. Yan, and E. Bitar, "A distributionally robust approach to regret optimal control using the Wasserstein distance," in *Proc. IEEE Conf. Decision Control*, 2023, pp. 2768–2775.

[16] J. Yue, B. Chen, and M.-C. Wang, "Expected value of distribution information for the newsvendor problem," *Oper. Res.*, vol. 54, no. 6, pp. 1128–1136, 2006.

[17] G. Perakis and G. Roels, "Regret in the newsvendor model with partial information," *Oper. Res.*, vol. 56, no. 1, pp. 188–203, 2008.

[18] D. Bertsimas, I. Dunning, and M. Lubin, "Reformulation versus cutting-planes for robust optimization: A computational study," *Comput. Manag. Sci.*, vol. 13, pp. 195–217, Apr. 2016.

[19] N. Fournier and A. Guillin, "On the rate of convergence in Wasserstein distance of the empirical measure," *Probab. Theory Relat. Fields*, vol. 162, pp. 707–738, Aug. 2015.

[20] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia, PA, USA: SIAM, 2009.

[21] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, vol. 174. New York, NY, USA: W. H. Freeman, 1979.

[22] J. E. Kelley Jr., "The cutting-plane method for solving convex programs," *J. Soc. Ind. Appl. Math.*, vol. 8, no. 4, pp. 703–712, 1960.

[23] O. Wilchfort and J. R. Lund, "Shortage management modeling for urban water supply systems," *J. Water Resour. Plan. Manag.*, vol. 123, no. 4, pp. 250–258, 1997.

[24] N. Noyan, "Risk-averse two-stage stochastic programming with an application to disaster management," *Comput. Oper. Res.*, vol. 39, no. 3, pp. 541–559, 2012.

[25] Y. Cho and I. Yang, "On affine policies for Wasserstein distributionally robust unit commitment," in *Proc. IEEE Conf. Decision Control*, 2022, pp. 7167–7173.

[26] B. Zeng and L. Zhao, "Solving two-stage robust optimization problems using a column-and-constraint generation method," *Oper. Res. Lett.*, vol. 41, no. 5, pp. 457–461, 2013.

[27] . J. Blanchet, K. Murthy, and V. A. Nguyen, "Statistical analysis of Wasserstein distributionally robust estimators," *Emerging Optimization Methods and Modeling Techniques with Applications* (TutORials in Operations Research). Catonsville, MD, USA: INFORMS, 2021, pp. 227–254.