

Nonlinear Distributionally Robust Optimization

Mohammed Rayyan Sheriff and Peyman Mohajerin Esfahani

ABSTRACT. This article focuses on a class of distributionally robust optimization (DRO) problems where, unlike the growing body of the literature, the objective function is potentially nonlinear in the distribution. Existing methods to optimize nonlinear functions in probability space use the Frechet derivatives, which present both theoretical and computational challenges. Motivated by this, we propose an alternative notion for the derivative and corresponding smoothness based on Gateaux (G)-derivative for generic risk measures. These concepts are explained via three running risk measure examples of variance, entropic risk, and risk on finite support sets. We then propose a G-derivative based Frank-Wolfe (FW) algorithm for generic nonlinear optimization problems in probability spaces and establish its convergence under the proposed notion of smoothness in a completely norm-independent manner. We use the set-up of the FW algorithm to devise a methodology to compute a saddle point of the nonlinear DRO problem. Finally, we validate our theoretical results on two cases of the *entropic* and *variance* risk measures in the context of portfolio selection problems. In particular, we analyze their regularity conditions and “sufficient statistic”, compute the respective FW-oracle in various settings, and confirm the theoretical outcomes through numerical validation.

Keywords. Gateaux derivative, norm-free-smoothness, Frank-Wolfe algorithm, saddle point

1. INTRODUCTION

Modern-day decision problems involve uncertainty in the form of a random variable ξ whose behavior is modeled via a probability distribution \mathbb{P}_o . A central object to formalize such decision-making problems under uncertainty is risk measures. The most popular risk measure is arguably the expected loss, yielding the classical decision-making problem of

$$\text{Stochastic Program: } \min_{x \in \mathcal{X}} \mathbb{E}_{\mathbb{P}_o} [\ell(x, \xi)], \quad (1)$$

where ℓ is the loss function of interest, and \mathcal{X} being the set of feasible decisions. The paradigm of stochastic programming (SP) relies on the assumption that the distribution \mathbb{P}_o is available (or at least up to its sufficient statistics), thereby the expectation can be computed for every decision $x \in \mathcal{X}$. A common practical challenge is, however, that the complete information of \mathbb{P}_o may not be available. Moreover, it might also be the case that the distribution is varying over a period of time which could be difficult to characterize. These limitations call for a more conservative risk measure to ameliorate the decision performance in such situations.

Date: October 29, 2024.

The authors are with the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands. Emails: M.R.Sheriff@tudelft.nl, P.MohajerinEsfahani@tudelft.nl. The authors acknowledge the fruitful discussions with Armin Eftekhari. This research is supported by the European Research Council (ERC) under the grant TRUST-949796.

An alternative framework is Robust Optimization (RO) where the decision-maker has only access to the support of uncertainty and takes the most conservative approach:

$$\text{Robust Optimization:} \quad \min_{x \in \mathcal{X}} \max_{\xi \in \Xi} \ell(x, \xi). \quad (2)$$

For many interesting examples, the RO min-max problem admits tractable reformulations that can be solved efficiently [3]. However, a generic RO problem is known to be computationally formidable as the inner maximization over ξ can be NP-hard. Moreover, if the support Ξ of the distribution is “large”, the results of RO tend to be extremely conservative.

Distributionally Robust Optimization (DRO). The SP and RO decision models represent two extreme cases of having full or bare minimum distributional information, respectively. In practice, however, we often have more information about the ground truth distribution than just its support. A typical example is when we have sample realizations $\{\hat{\xi}_i : i = 1, 2, \dots, N\}$. Looking at such settings through the lens of SP, one may construct a *nominal distribution* $\hat{\mathbb{P}}$ and use it as a proxy for \mathbb{P}_o in the SP (1). A standard data-driven nominal distribution is the discrete distribution $\hat{\mathbb{P}} = \sum_{i=1}^N \delta_{\hat{\xi}_i}$. The SP decision when $\mathbb{P}_o = \hat{\mathbb{P}}$ in (1) often admits a poor out-of-sample performance on a different dataset, a phenomenon that is also known as the “optimizer’s curse” or “overfitting” [45]. On the other hand, the RO viewpoint in (2) completely disregards the statistical information of \mathbb{P}_o available through the dataset $(\hat{\xi}_i)_i$, or any other form of prior information.

An attempt to bridge the SP and RO modeling frameworks gives rise to the paradigm of *Distributionally Robust Optimization* (DRO), which dates back to the Scarf’s seminal work on the ambiguity-averse newsvendor problem in 1958 [41]. The “ambiguity” set \mathcal{P} is a family of distributions that are close in some sense to the nominal distribution $\hat{\mathbb{P}}$, potentially including the true unknown distribution \mathbb{P}_o . With this in mind, the DRO problem is formulated as

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\ell(x, \xi)]. \quad (3)$$

On the one hand, if the ambiguity set is *very big*, possibly including all possible distributions, the DRO problem (3) reduces to the RO problem (2). On the other hand, if the ambiguity set \mathcal{P} is *very small*, potentially a singleton containing the nominal distribution, then the DRO problem reduces to the SP problem (1). In this light, the DRO framework (3) provides flexibility for the decision-maker to fill the gap between SP (1) and RO (2). The ambiguity set \mathcal{P} in (3) is typically constructed either based on the moments information [10, 17, 51], or a neighborhood of $\hat{\mathbb{P}}$ with respect to a notion of distance over probability distributions, e.g., Prohorov [12], Kullback-Leiber [20, 11], Wasserstein [20, 27, 5, 14, 16], Sinkhorn [50], to name but a few; see also the surveys [39, 23] and the references therein.

Linear DRO problems: An important feature of (3) is the linearity of the objective function in the distribution \mathbb{P} , which is also shared among all the literature mentioned above. The simplicity of this linearity in the inner maximization of (3) is the underlying driving force to develop tractable convex reformulations and computational solutions for various combinations of ambiguity sets and cost functions [38, 53, 40, 43, 7]. This line of research effectively translates the original infinite-dimensional DRO problem (3) to a tractable finite-dimensional one; see also the general optimal transport framework of [52], and the case of mean-covariance risk measure [34] in a financial context.

Nonlinear DRO (NDRO): Our goal here is to generalize the linear setting (3) to

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} F(x, \mathbb{P}), \quad (4)$$

where $F : \mathcal{X} \times \mathcal{P} \rightarrow \mathbb{R}$ is a generic, possibly nonlinear, function signifying a risk measure. We refer to the decision-making problem (4) as *Nonlinear Distributionally Robust Optimization (NDRO)*. There are several interesting NDRO examples including the variance [4], entropy [46], [8, Example 1, p. 14]. In this study, instead of focusing only on the outer decision x in (4), we aim to compute an approximate saddle point between the decision-maker and the nature presented by the distribution \mathbb{P} . A motivation supporting this effort is the fact that, unlike the worst-case distribution computed for a given decision, the saddle point distributions (also called Nash equilibria) naturally retain more realistic features [42]. From a computational perspective, the existing techniques deployed in linear DRO problems cannot be directly extended for the NDRO (4). The objective of this work is to precisely tackle this challenge, where we seek to devise a methodology along with mild regularity conditions under which a saddle point solution exists and can be computed.

Frank-Wolfe (FW) algorithm. It is a first-order method that only uses the information of the gradient to solve a constrained optimization problem [19, 13, 9]. In a nutshell, each iteration of the FW algorithm optimizes the linear function given by the gradient of the objective function and then takes a step towards the optimizer of the linear problem. On the contrary, other first-order methods like projected gradient descent require a quadratic function to be minimized at each iteration. Since solving a linear problem at each iteration is easier than a quadratic one, the iteration complexity of the FW algorithm is much simpler than other first-order methods. This is crucial for optimization problems over probability distributions since the complexity of projecting onto the ambiguity set can be as challenging as solving the original NDRO problem. In comparison, optimizing a cost function that is linear in distributions admits strong duality and tractable finite-dimensional reformulations for many interesting examples as seen in linear DRO. Therefore, we seek to use the principles of the Frank-Wolfe algorithm in the context of NDRO problems and devise an iterative procedure to compute a saddle point of (4).

The FW algorithm for optimization problems over probability distributions has already been introduced in [21], which uses Frechet derivatives and the associated notion of smoothness to establish convergence. We would also like to highlight the work of [25], where the canonical gradient ascent-descent algorithm [28, 18] for min-max problems is extended to infinite-dimensional spaces involving probability distributions using the Frechet derivatives and its smoothness. In a similar spirit, the recent work of [33] proposes a mirror-descent algorithm [29] for constrained nonlinear optimization problems over probability distributions using the Frechet derivative. However, Frechet derivatives are difficult to deal with in practice due to several prominent challenges including their (i) existence, (ii) finite representability, and (iii) norm consistency; see Section 3.3 for more details on this. Our focus in this study is to remedy this by proposing a FW algorithm based on an alternative G-derivative along with a completely norm-independent convergence analysis.

Contributions. The main contributions of this study are summarized as follows:

- (i) **Norm free smoothness in probability spaces.** We propose a novel G-derivative based notion of derivative for nonlinear risk measures (Definition 3.1), and the associated notion of

smoothness that is independent of the norm structure on the ambiguity set (Definition 3.6). Moreover, we also derive conditions on the function $F(x, \mathbb{P})$ such that the risk measure $\mathbb{P} \mapsto \min_{x \in \mathcal{X}} F(x, \mathbb{P})$ is smooth in the sense of the proposed notion (Lemma 5.5).

- (ii) **G-derivative based Frank-Wolfe algorithm.** We provide a Frank-Wolfe algorithm based on the proposed notion of derivative for optimizing nonlinear risk measures. Moreover, the classical proof techniques for the FW algorithm carry forward under the proposed notion of derivative and smoothness, resulting in apriori (Propositions 4.4) and aposteriori (Proposition 4.6) convergence guarantees that exist for finite-dimensional problems.
- (iii) **Saddle point seeking algorithm for NDRO problems.** For the potentially infinite-dimensional min-max problem of NDRO, we propose a FW-based algorithm to compute a saddle point (Algorithm 1), and also quantify its convergence properties (Theorem 5.6).
- (iv) **The entropic and variance risk measures.** We study the NDRO problem and our proposed algorithm for two cases of the entropic and variance risk measures in Sections 6 and 7, respectively. An interesting difference between these two examples is that in the context of portfolio selection, the variance preserves its “sufficient statistic” throughout the FW-algorithm while the entropic risk requires the knowledge of the entire distribution. We establish the required convex regularity conditions (Lemmas 6.3 and 7.2, respectively), and then provide a complete description of the corresponding FW-oracle (Lemmas 6.4 and 7.4, 7.6, respectively). Furthermore, for the minimum variance portfolio selection problem, in the special case of $\Xi = \mathbb{R}^n$ and the type-2 Wasserstein ambiguity set, we slightly extend the results of [4] by explicitly describing the saddle point of the minimum variance problem when the feasible portfolio set \mathcal{X} is *any* arbitrary compact set (Proposition 7.5). To facilitate the reproducibility of the numerical results, we also provide the open source repository [44] including the respective MATLAB code.

The rest of the paper is organized as follows: In Section 2, we discuss a generic nonlinear optimization problem over distributions with relevant examples. In Section 3, we introduce the notion of a directional derivative and the associated notion of smoothness. In Section 4, the FW-algorithm and its convergence guarantees are discussed. In Section 5, we introduce the NDRO problem, discuss a solution concept, and provide an FW-based algorithm to compute the solution. To illustrate the methodology (the required assumptions and algorithms), we discuss in detail the NDRO problems for two nonlinear risk measures of entropic and variance, respectively, in Sections 6 and 7. The numerical results are provided for both risk measures in the context of optimal portfolio selection. All of the technical proofs have been relegated towards the end of the article in Section 8.

Notations. The set of real-valued $n \times n$ matrices is denoted by $\mathbb{R}^{n \times n}$. The trace of a matrix M is denoted by $\text{tr}(M)$ and being positive semi-definite is denoted by $M \succeq 0$. For a function r defined over a finite-dimensional space, its gradient at x is denoted by $\nabla r(x)$. For a multivariate distribution \mathbb{P} , the first and second moments are denoted by the shorthand notation $\mu_{\mathbb{P}} := \mathbb{E}_{\mathbb{P}}[\xi]$ and $\Sigma_{\mathbb{P}} := \mathbb{E}_{\mathbb{P}}[\xi\xi^{\top}]$, respectively. Given two probability distributions \mathbb{P}, \mathbb{Q} , we also use $\mathbb{E}_{\mathbb{P}-\mathbb{Q}}[\cdot] := \mathbb{E}_{\mathbb{P}}[\cdot] - \mathbb{E}_{\mathbb{Q}}[\cdot]$. Other notations shall be defined whenever necessary.

2. WORST-CASE NONLINEAR RISK MEASURES

In this section, we focus our attention to the potentially infinite-dimensional optimization problem:

$$R^* := \sup_{\mathbb{P} \in \mathcal{P}} R(\mathbb{P}), \quad (5)$$

where $R : \mathcal{P} \rightarrow \mathbb{R}$ is a desired, possibly *nonlinear*, concave risk measure, and \mathcal{P} is the ambiguity set containing a family of probability distributions over $\Xi \subset \mathbb{R}^n$. Our main goal is to develop a framework with appropriate mathematical notions for a *Franke-Wolfe* (FW) like algorithm to solve problem (5) and investigate its convergence properties. Furthermore, since the FW algorithm for (5) operates in an infinite-dimensional setting, we also seek to derive its tractable finite-dimensional simplification to solve specific instances of (5).

Definition 2.1 (Regular risk (RR) measures & sufficient statistic). *A risk measure R in (5) is regular if it can be described as $R(\mathbb{P}) = r(\mathbb{E}_{\mathbb{P}}[L(\xi)])$, for some functions $L : \Xi \rightarrow \mathbb{R}^m$ that is integrable for all $\mathbb{P} \in \mathcal{P}$ and $r : \mathbb{R}^m \rightarrow \mathbb{R}$ that is concave and differentiable. We refer to $\mathbb{E}_{\mathbb{P}}[L(\xi)]$ the “sufficient statistic” as the risk value depends on the distribution \mathbb{P} only through this vector.*

A particular feature of the regular risks in Definition 2.1 is that its value is determined based on a finite-dimensional vector $\mathbb{E}_{\mathbb{P}}[L(\xi)]$. This is a concept close to “sufficient statistic”, which is particularly appealing when it comes to the computation of the worst-case risk in (5). We will return to this in Section 4 when discussing the FW algorithm in the space of probability distributions.

Example 2.2 (RR-examples). Throughout this study, we discuss three particular examples of the regular risk measures to showcase the concepts and our theoretical statements:

- (a) *Variance*: A popular risk measure is the variance associated with the distribution. Formally, considering ξ to be a \mathbb{P} -distributed random variable for $\mathbb{P} \in \mathcal{P}$, the associated variance is

$$V(\mathbb{P}) := \mathbb{E}_{\mathbb{P}}[\|\xi - \mathbb{E}_{\mathbb{P}}[\xi]\|_2^2]. \quad (6a)$$

Considering the functions $L : \Xi \rightarrow \mathbb{R}^{n \times n} \times \mathbb{R}^n$ and $r : \mathbb{R}^{n \times n} \times \mathbb{R}^n \rightarrow \mathbb{R}$

$$\begin{cases} L(\xi) = (\xi\xi^\top, \xi) \\ r(\Sigma, \mu) = \text{tr}(\Sigma) - \|\mu\|_2^2, \end{cases} \quad (6b)$$

one can observe that the variance is indeed an RR measure, since

$$V(\mathbb{P}) = \mathbb{E}_{\mathbb{P}}[\|\xi\|_2^2] - \|\mathbb{E}_{\mathbb{P}}[\xi]\|_2^2 = \text{tr}(\mathbb{E}_{\mathbb{P}}[\xi\xi^\top]) - \|\mathbb{E}_{\mathbb{P}}[\xi]\|_2^2 = r(\mathbb{E}_{\mathbb{P}}[L(\xi)]).$$

- (b) *Entropic risk*: Another interesting example of a nonlinear RR measure is the entropic risk of a multivariate distribution [46, Section 5]. If $\mathbb{P} \in \mathcal{P}$ is a distribution with marginals \mathbb{P}_j for $j = 1, 2, \dots, n$, (i.e., the j -th component ξ_j is \mathbb{P}_j distributed). The entropic risk $\mathcal{E}(\mathbb{P})$ associated with the distribution \mathbb{P} is defined as

$$\mathcal{E}(\mathbb{P}) := \sum_{j=1}^n \frac{1}{\theta_j} \log \left(\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j \xi_j}] \right), \quad (7a)$$

where $(\theta_j)_{j=1}^n$ is a collection of positive real numbers referred to as the *risk-aversion parameters*. This is indeed an RR measure, which can be seen by introducing the functions

$$\begin{cases} L(\xi) = (e^{-\theta_1 \xi_1}, e^{-\theta_2 \xi_2}, \dots, e^{-\theta_n \xi_n}) \\ r(z) = \sum_{j=1}^n \frac{1}{\theta_j} \log(z_j). \end{cases} \quad (7b)$$

(c) *Finite-support*: The final example is the case of Ξ being a finite set: $\Xi = \{\xi_i : i = 1, 2, \dots, N\}$. In this case, the simplex Δ^N is the set of all probability distributions on Ξ , and the ambiguity set of distributions is a subset, $\mathcal{P} \subset \Delta^N$. It turns out that *any* arbitrary risk measure R can be characterized as an RR-measure in the sense of Definition 2.1 by introducing appropriate functions $r : \mathbb{R}^N \rightarrow \mathbb{R}$, and $L : \Xi \rightarrow \mathbb{R}^N$. To see this, we first observe that Ξ being a finite set gives rise to the enumerating bijection $b : \Xi \rightarrow \{1, 2, \dots, N\}$ defined by $b(\xi) := \{i : \xi = \xi_i\}$; secondly, the matrix $M \in \mathbb{R}^{N \times N}$ given by $[M]_{ij} := (j/N)^i$ for $i, j = 1, 2, \dots, N$, is invertible. Then, considering the functions

$$\begin{cases} L(\xi) = ((b(\xi)/N), (b(\xi)/N)^2, \dots, (b(\xi)/N)^N)^\top \\ r(z) = R(M^{-1}z) \text{ for } z \in M(\mathcal{P}), \end{cases} \quad (8)$$

we have $\mathbb{E}_{\mathbb{P}}[L(\xi)] = M \cdot \mathbb{P}$ (viewing $\mathbb{P} \in \Delta^N$ as an element of \mathbb{R}^N), and as such, $R(\mathbb{P}) = r(\mathbb{E}_{\mathbb{P}}[L(\xi)])$ for all $\mathbb{P} \in \mathcal{P}$.

3. DERIVATIVES OF RISK MEASURES

A well-defined notion of the gradient is a fundamental quantity in developing iterative algorithms like that of FW to solve any optimization problem. For finite-dimensional convex problems, the FW algorithm optimizes the linear functional given by the gradient of the objective function over the feasible set at each iteration. Naturally, devising a similar algorithm for (5) requires at least a well-defined notion of *directional derivative* or G-derivative [2, (A.3), p. 152].

3.1. Gateaux directional derivatives

Definition 3.1 (G-derivative). *Given $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$, the Gateaux (G)-derivative $dR(\mathbb{P}; \mathbb{Q})$ of the risk measure R at \mathbb{P} in the direction \mathbb{Q} is defined as*

$$dR(\mathbb{P}; \mathbb{Q}) := \lim_{\gamma \downarrow 0} \frac{1}{\gamma} (R(\mathbb{P}_\gamma) - R(\mathbb{P})), \quad \text{where } \mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P}). \quad (9)$$

Whenever the above limit exists, we say that the function R is \mathbb{Q} -directionally differentiable at \mathbb{P} . Moreover, we say that the risk measure R is directionally differentiable on \mathcal{P} if it is \mathbb{Q} -directionally differentiable at every $\mathbb{P} \in \mathcal{P}$ and for all $\mathbb{Q} \in \mathcal{P}$.




FIGURE 1. Pictorial representation of the risk surface and its directional derivative.

We note that the G-directional derivative in Definition 3.1 does not rely on any metric underlying the space of probability distributions. This is in fact the main feature with respect to the alternative F-derivative (Frechet derivative) that will be discussed in Section 3.3. Figure 1 visualizes this directional derivative. Next lemma provides an explicit description of Definition 3.1 for RR measures.

Lemma 3.2 (Regular G-derivatives). *Suppose that the risk measure is regular, i.e., $R(\mathbb{P}) = r(\mathbb{E}_{\mathbb{P}}[L(\xi)])$ with ∇r denoting the gradient of function r . Then, for any $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$, we have*

$$dR(\mathbb{P}; \mathbb{Q}) = \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[\langle \nabla r(\mathbb{E}_{\mathbb{P}}[L(\xi)]), L(\xi) \rangle] = \langle \nabla r(\mathbb{E}_{\mathbb{P}}[L(\xi)]), \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[L(\xi)] \rangle. \quad (10)$$

Lemma 3.2 indicates that the G-derivative of an RR-measure is essentially characterized by the “sufficient statistic” vectors $\mathbb{E}_{\mathbb{P}}[L(\xi)]$ and $\mathbb{E}_{\mathbb{Q}}[L(\xi)]$ (cf. (10)). Thus, if an algorithm optimizes an RR measure using only their directional derivatives, it then requires tracking the evolution of this finite-dimensional sufficient statistic, allowing them to be implemented efficiently.

Example 3.3 (Regular G-derivatives). The directional derivative of an RR measure is completely characterized in terms of only a few finite-dimensional quantities that depend on the moments of the distribution, and the functions r and L .

- (a) *Variance*: With the underlying inner product $\langle (\Sigma, \mu), (\Sigma', \mu') \rangle := \text{tr}(\Sigma^\top \Sigma') + \mu^\top \mu'$, we recall from (6) that $\nabla r(\Sigma, \mu) = (\mathbb{I}_n, -2\mu)$. Then the \mathbb{Q} -directional derivative of the variance V , at \mathbb{P} can be calculated from (10) as

$$dV(\mathbb{P}; \mathbb{Q}) = \text{tr}(\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}}) - 2\mu_{\mathbb{P}}^\top(\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}). \quad (11)$$

- (b) *\mathcal{E} -measure*: With the canonical inner product $\langle z, \hat{z} \rangle = z^\top \hat{z}$, we recall from (7) that $\nabla r(z) = ((\theta_1 z_1)^{-1}, (\theta_2 z_2)^{-1}, \dots, (\theta_n z_n)^{-1})$ at every $z \in \mathbb{R}_+^n$. Thus, its \mathbb{Q} -directional derivative at \mathbb{P} is

$$d\mathcal{E}(\mathbb{P}; \mathbb{Q}) = \sum_{j=1}^n \frac{\mathbb{E}_{\mathbb{Q}_j-\mathbb{P}_j}[e^{-\theta_j \xi_j}]}{\theta_j \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j \xi_j}]}. \quad (12)$$

- (c) *Finite support*: Recall from (8) that $\mathbb{E}_{\mathbb{P}}[L(\xi)] = M \cdot \mathbb{P}$ for all $\mathbb{P} \in \mathcal{P}$. Then with the canonical inner product $\langle z, z' \rangle = z^\top z'$, we have $\nabla r(z) = (M^{-1})^\top \nabla R(M^{-1}z)$ for all $z \in M(\mathcal{P})$. Substituting these quantities in (10) and simplifying, we get

$$dR(\mathbb{P}; \mathbb{Q}) = (\nabla R(\mathbb{P}))^\top (\mathbb{Q} - \mathbb{P}). \quad (13)$$

It is to be observed that the matrix M has no relevance in the G-derivative as one would expect since the risk measure R is defined independent of the matrix M .

The G-derivative enjoys inherent properties that will be helpful to devise computational solutions.

Proposition 3.4 (G-derivative: properties). *For any $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$, and $\gamma \in [0, 1]$, let $\mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$, then the directional derivative $dR(\mathbb{P}; \cdot)$ in Definition 3.1 satisfies*

$$\text{Positively homogeneous:} \quad dR(\mathbb{P}; \mathbb{P}_\gamma) = \gamma dR(\mathbb{P}; \mathbb{Q}) \quad (14a)$$

$$\text{Upper bound for concave risk measures:} \quad R(\mathbb{Q}) - R(\mathbb{P}) \leq dR(\mathbb{P}; \mathbb{Q}). \quad (14b)$$

Remark 3.5 (Optimality conditions). *Similar to KKT conditions, F-derivative based first-order optimality conditions for both constrained and unconstrained versions of (5) are given in [24, Section 3]. With G-derivatives, the upper bound (14b) for the concave risk measure R immediately*

gives sufficient optimality conditions for (5). More precisely, if $\mathbb{P}^* \in \operatorname{argmax}_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}^*; \mathbb{Q})$, then applying (14b) for any $\mathbb{Q} \in \mathcal{P}$ yields

$$R(\mathbb{Q}) \leq R(\mathbb{P}^*) + dR(\mathbb{P}^*; \mathbb{Q}) \leq R(\mathbb{P}^*) + dR(\mathbb{P}^*; \mathbb{P}^*) = R(\mathbb{P}^*),$$

which establishes the optimality $\mathbb{P}^* \in \operatorname{argmax}_{\mathbb{Q} \in \mathcal{P}} R(\mathbb{Q})$.

3.2. Smoothness

The proposed approach to solve (5) builds on the Frank-Wolfe (FW) algorithm for finite-dimensional convex optimization problems [13]. It is a well-known fact that the primal sub-optimality in FW algorithm converges at a sub-linear rate $O(1/k)$ for optimization of “smooth” objective functions over compact feasible sets in finite-dimensional problems. A convex function is said to be smooth if it has Lipschitz continuous gradients with respect to some norm. The choice of norm in an infinite-dimensional setting can be problematic, as all the norms are not equivalent (unlike the finite-dimensional setting). To extend such convergence attributes for a FW like algorithm in the infinite-dimensional setting of (5), we first propose an appropriate notion of smoothness in terms of the directional derivatives (as in Definition 3.1) that is also norm-independent.

Definition 3.6 (G-smoothness). *The risk measure R is C -smooth if there exists a constant $C \geq 0$ such that for all $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$ and $\gamma \in [0, 1]$, we have the inequality*

$$dR(\mathbb{P}_\gamma; \mathbb{P}) + dR(\mathbb{P}; \mathbb{P}_\gamma) \leq \gamma^2 C \quad \text{where } \mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P}). \quad (15)$$

Connection to the existing notions of smoothness. The notion of smoothness in Definition 3.6 is a generalization of the canonical smoothness condition of Lipschitz continuous gradients [31, Section 2.1.5]. Formally, a function $f : \mathcal{D} \rightarrow \mathbb{R}$ with a finite-dimensional domain $(\mathcal{D}, \|\cdot\|)$ and gradients ∇f is said to be β -smooth if

$$\|\nabla f(x) - \nabla f(y)\|_* \leq \beta \|x - y\| \quad \text{holds for all } x, y \in \mathcal{D}. \quad (16)$$

It is further generalized (or relaxed) by the notion of Holder-smoothness (particularly, 1-Holder smooth) where it is required that

$$\langle \nabla f(x) - \nabla f(y), y - x \rangle \leq \beta \|x - y\|^2 \quad \text{holds for all } x, y \in \mathcal{D}. \quad (17)$$

Furthermore, if the set \mathcal{D} is $\|\cdot\|$ -bounded in addition, then the notion of Holder-smoothness (17) is sufficient to the requirement that there exists some constant $C \geq 0$ such that for every $x, y \in \mathcal{D}$, $\gamma \in [0, 1]$, the function f satisfies the inequality

$$\langle \nabla f(x) - \nabla f(x_\gamma), x_\gamma - x \rangle \leq \gamma^2 C \quad \text{where } x_\gamma := x + \gamma(y - x). \quad (18)$$

Expanding the left-hand side of (18) as $\langle \nabla f(x) - \nabla f(x_\gamma), x_\gamma - x \rangle = \langle \nabla f(x), x_\gamma - x \rangle + \langle \nabla f(x_\gamma), x - x_\gamma \rangle$, we see that the individual terms are simply the directional derivatives of f ; Definition 3.6 of the proposed notion of smoothness becomes apparent at once. We also highlight that the notion of smoothness in Definition 3.6 is closely related to the notion of “curvature coefficient” used in [19, (3)]. In fact, it can be easily shown that a function has a finite curvature coefficient if it is smooth in the sense of Definition 3.6.

Concave quadratic lower bound via smoothness. It is well established that in a finite-dimensional setting, smoothness conditions like (16) and (17) give rise to a global convex (resp. concave) quadratic upper bound (resp. lower bound). For the risk measure R specifically, a sample representation of such a quadratic lower bound is shown in Figure 2. The existence of such bounds guarantees that the curvature of the function is at most that of the quadratic bounds, which is crucial in concluding the convergence of the FW algorithm. In other words, if a function f is β -canonically smooth (as in (16)), then the following inequality holds:




FIGURE 2. Concave-quadratic lower bound via smoothness.

$$-\frac{\beta}{2} \|y - x\|^2 \leq f(y) - f(x) - \langle \nabla f(x), y - x \rangle \leq \frac{\beta}{2} \|y - x\|^2 \quad \text{for all } x, y \in \mathcal{D}.$$

The notion of smoothness in Definition 3.6 imposes similar quadratic bounds, but without using any norm. This is done by enforcing “quadratic-like” bounds to hold uniformly over all directions. The following Lemma establishes the smoothness of the class of RR measures.

Lemma 3.7 (Regular G -smoothness). *Suppose the RR measure $R(\mathbb{P}) = r(\mathbb{E}_{\mathbb{P}}[L(\xi)])$ satisfies*

- (i) *bounded diameter in expectations with $d := \sup_{\mathbb{P}, \mathbb{Q} \in \mathcal{P}} \|\mathbb{E}_{\mathbb{P}}[L(\xi)] - \mathbb{E}_{\mathbb{Q}}[L(\xi)]\| < +\infty$, and*
- (ii) *it admits a smooth r in the canonical sense (16), i.e.,*

$$\|\nabla r(u) - \nabla r(v)\|_* \leq \beta \|u - v\| \quad \text{for all } u, v \in \{\mathbb{E}_{\mathbb{P}}[L(\xi)] : \mathbb{P} \in \mathcal{P}\}.$$

Then, it is (βd^2) -smooth in the sense of Definition 3.6.

Example 3.8 (Regular G -smoothness). For the RR measures in Examples 2.2 and 3.3 with the same underlying norms, the smoothness constant β of the respective r -function is as follows:

- (a) *Variance:* Recalling the function r from (6), and that $\nabla r(\Sigma, \mu) = (\mathbb{I}_n, -2\mu)$, we get $\beta = 2$.
- (b) *Entropic risk:* Assume that there exists $b > 0$ such that $b \leq \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j \xi_j}]$ for all $j = 1, 2, \dots, n$, and $\mathbb{P} = (\mathbb{P}_j)_j \in \mathcal{P}$, and letting $\theta_{\min} := \min_{j \leq n} \theta_j$, we see that the function $r : [b, +\infty)^n \rightarrow \mathbb{R}$ as given in (7) is β -canonically smooth for $\beta = (b^2 \theta_{\min})^{-1}$.
- (c) *Finite support:* Assuming that the risk measure R has β' -Lipschitz continuous gradients on Δ^N with respect to some norm $\|\cdot\|$. It follows from (8), that $\beta = \beta' \|M^{-1}\|_o^2$, where, $\|M^{-1}\|_o = \sup_{\|u\| \leq 1} \|M^{-1}u\|$.

One of the advantages of the proposed notion of smoothness is that it allows us to employ inequalities and bound sets in a finite-dimensional space to guarantee the smoothness of the risk measures. Since all norms on finite-dimensional spaces are equivalent, establishing that $\beta, d < +\infty$ with respect to the same norm is not restrictive, even though this may give rise to dimension-dependent constants β and d .

3.3. Frechet derivatives

An important observation to be made is that the \mathbb{Q} -directional derivative $dR(\mathbb{P}; \mathbb{Q})$ of regular-risk measures is *affine* in \mathbb{Q} .¹ In principle, without any further assumptions on the risk measure R , its \mathbb{Q} -directional derivative need not be affine in \mathbb{Q} . Counter examples of functions that have well defined G-derivatives in all directions but that are nonlinear in the direction exist even among functions defined on \mathbb{R}^2 , let alone the infinite-dimensional setting of \mathcal{P} . A sufficient condition for the directional derivative to be affine in \mathbb{Q} is the existence of a stronger notion of derivative called the *Frechet-derivative* or *F-derivative*.

Definition 3.9 (Frechet-derivative). *The Frechet(F)-derivative of the risk measure R at $\mathbb{P} \in \mathcal{P}$ associated with a given norm $\|\cdot\|_{\mathcal{P}}$ on \mathcal{P} , is a function $\ell_{\mathbb{P}} : \Xi \rightarrow \mathbb{R}$ such that the mapping $\mathcal{P} \ni \mathbb{P}' \mapsto \mathbb{E}_{\mathbb{P}'}[\ell_{\mathbb{P}}(\xi)]$ is continuous w.r.t. $\|\cdot\|_{\mathcal{P}}$ and satisfies*

$$0 = \lim_{\|\mathbb{P}' - \mathbb{P}\|_{\mathcal{P}} \downarrow 0} \frac{R(\mathbb{P}') - R(\mathbb{P}) - \mathbb{E}_{\mathbb{P}' - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)]}{\|\mathbb{P}' - \mathbb{P}\|_{\mathcal{P}}}. \quad (19)$$

Smoothness with F-derivatives. If the F-derivative $\ell_{\mathbb{P}}$ were to exist at every $\mathbb{P} \in \mathcal{P}$, the canonical notion of smoothness (39) can be naturally extended into the infinite-dimensional setting. The risk measure R is *F-smooth* if there exists some $\beta \geq 0$ such that its F-derivative $\ell_{\mathbb{P}}$ satisfies

$$\text{F-smoothness: } \quad \|\ell_{\mathbb{P}} - \ell_{\mathbb{P}'}\|_{\mathcal{P}^*} \leq \beta \|\mathbb{P} - \mathbb{P}'\|_{\mathcal{P}} \quad \text{for all } \mathbb{P}, \mathbb{P}' \in \mathcal{P}, \quad (20)$$

where $\|\cdot\|_{\mathcal{P}^*}$ is the dual norm of $\|\cdot\|_{\mathcal{P}}$. If the risk measure in (5) is F-smooth, most of the convergence analysis due to smoothness in finite-dimensional convex problems simply carries forward to the infinite-dimensional setting right away. This provides a natural recipe to extend the FW algorithm into the infinite-dimensional setting of probability spaces and establish their convergence under F-smoothness. F-derivative based FW-algorithms in probability spaces have already been studied in the literature [21] with a slightly more general notion of smoothness than (20), and for a slightly more general class of risk measures than the usual concavity assumption. Our approach differs from [21] in the fact that we only make use of G-derivatives in both the development of the FW-algorithm and also in establishing its convergence based on only G-derivative based regularity conditions, which are simpler to deal with than F-derivatives.

Comparison with G-derivatives. It is not necessary for a function to have F-derivatives even if it has affine directional derivatives in all directions, such counterexamples exist even in a finite-dimensional setting. On the contrary, if the risk measure R has a well-defined F-derivative $\ell_{\mathbb{P}}$, then it can be shown that its \mathbb{Q} -directional derivatives also exist in all directions $\mathbb{Q} \in \mathcal{P}$. This is easily seen by considering $\mathbb{P}' = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$ for $\gamma \in [0, 1]$ in the definition (19) and observing that $\mathbb{E}_{\mathbb{P}' - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)] = \gamma \mathbb{E}_{\mathbb{Q} - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)]$. When $\mathbb{Q} \neq \mathbb{P}$, the limit $\|\mathbb{P}' - \mathbb{P}\|_{\mathcal{P}} \downarrow 0$ is achieved if and only if $\gamma \downarrow 0$. Then it follows from (19) that

$$0 = \lim_{\gamma \downarrow 0} \frac{R(\mathbb{P}_{\gamma}) - R(\mathbb{P}) - \gamma \mathbb{E}_{\mathbb{Q} - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)]}{\gamma \|\mathbb{Q} - \mathbb{P}\|_{\mathcal{P}}} = \frac{1}{\|\mathbb{Q} - \mathbb{P}\|_{\mathcal{P}}} \left(dR(\mathbb{P}; \mathbb{Q}) - \mathbb{E}_{\mathbb{Q} - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)] \right).$$

Since $\mathbb{Q} \neq \mathbb{P}$, we conclude that the \mathbb{Q} -directional derivative exists and $dR(\mathbb{P}; \mathbb{Q}) = \mathbb{E}_{\mathbb{Q} - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)]$. Moreover, this equality holds even if $\mathbb{Q} = \mathbb{P}$ since $dR(\mathbb{P}; \mathbb{P}) = 0 = \mathbb{E}_{\mathbb{P} - \mathbb{P}}[\ell_{\mathbb{P}}(\xi)]$.

¹A function $f : \mathcal{S} \rightarrow \mathbb{R}$ is said to affine if $f(x + \theta(y - x)) = f(x) + \theta(f(y) - f(x))$ for every $x, y \in \mathcal{S}$ and $\theta \in [0, 1]$.

The notion of F-derivative relies heavily on the underlying metric structure on \mathcal{P} , whereas, the notion (Definition 3.1) of G-derivatives is independent of it. To compare, for a G-derivative to exist along a given direction, it is only required for the limit in (9) to exist. However, the existence of an F-derivative requires that the limit in (9) is achieved uniformly over all possible directions.

Challenges with F-derivatives. Even though the notion of F-smoothness (20) is a natural extension of canonical smoothness (39) in an infinite-dimensional setting, working with F-derivatives is potentially challenging due to the following reasons:

- (i) **Existence:** It is a stronger requirement that the limit in (20) converges uniformly (w.r.t. $\|\cdot\|_{\mathcal{P}}$) in all directions, which often implies that an F-derivative might not even exist.
- (ii) **Finite representability:** An F-derivative is a function $\ell_{\mathbb{P}} : \Xi \rightarrow \mathbb{R}$ which is an infinite-dimensional object, so a priori, it is not clear as to whether it can be characterized in terms of a few finite-dimensional quantities.
- (iii) **Norm consistency:** Most importantly, it is often very difficult to establish the smoothness condition (20) of the F-derivatives in a specific norm. To elaborate further, we know that the FW algorithm in a finite-dimensional convex problem converges sub-linearly, if the feasible set is bounded and the objective function is smooth. Since all norms on finite-dimensional vector spaces are equivalent, the choices of norms for establishing the smoothness of the objective function, and the boundedness of the feasibility set are irrelevant (even though this could potentially give rise to dimensionally dependent constants). However, since no such equivalence exists between norms on an infinite-dimensional space, it becomes then necessary that the risk measure is F-smooth w.r.t. the same norm under which the ambiguity set is bounded, which is a much stronger condition to expect.

We close this discussion by providing an example wherein a risk measure has directional derivatives and is G-smooth, yet its F-derivatives do not exist. In particular, we argue that all RR measures that satisfy conditions of Lemma 3.7 are G-smooth, whereas their F-derivative does not exist if the corresponding L function (i.e., sufficient statistic) is discontinuous.

Example 3.10 (G-derivative vs F-derivative). Suppose the support set is $\Xi = [-1, +1]$, the ambiguity set \mathcal{P} contains all possible distributions supported on Ξ and it is equipped with a Wasserstein metric. Let the regular risk be defined by the sign function $L(\xi) = \text{sgn}(\xi)$ (with the convention that $\text{sgn}(0) = 0$) and $r(z) = z$, i.e., $R(\mathbb{P}) = \mathbb{E}_{\mathbb{P}}[\text{sgn}(\xi)]$.

- (a) *Existence of G-smoothness:* We conclude from Lemma 3.2 that the G-derivatives of the risk measure exist and are given by $dR(\mathbb{P}; \mathbb{Q}) = \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[\text{sgn}(\xi)]$. Moreover, since the support Ξ is bounded, condition (i) of Lemma 3.7 is satisfied with $d = 2$, and since $\nabla r(z) = 1$ for all $z \in [-1, +1]$, condition (ii) of Lemma 3.7 is also satisfied with $\beta = 0$. Consequently, the given regular risk is 0-smooth in the sense of Definition 3.6.
- (b) *Non-existence of F-derivatives:* Since $dR(\mathbb{P}; \mathbb{Q}) = \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[\text{sgn}(\xi)]$, if the F-derivative of the risk measure R were to exist, the mapping $\mathbb{Q} \mapsto \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[\text{sgn}(\xi)]$ must be continuous. However, for the sequence of distributions $\mathbb{Q}_n(\xi) = \delta(\xi - 1/n)$, we see that $\mathbb{E}_{\mathbb{Q}_n-\mathbb{P}}[\text{sgn}(\xi)] = 1 - \mathbb{E}_{\mathbb{P}}[\text{sgn}(\xi)]$ for all $n = 1, 2, \dots$, and they converge to the distribution $\mathbb{Q} = \delta(\xi)$ in the Wasserstein metric, for which we have $\mathbb{E}_{\mathbb{Q}-\mathbb{P}}[\text{sgn}(\xi)] = -\mathbb{E}_{\mathbb{P}}[\text{sgn}(\xi)] \neq 1 - \mathbb{E}_{\mathbb{P}}[\text{sgn}(\xi)] = \mathbb{E}_{\mathbb{Q}_n-\mathbb{P}}[\text{sgn}(\xi)]$. Therefore, the F-derivative of the risk measure does not exist.

More sophisticated examples of risk measures and ambiguity sets can be constructed following the same underlying idea of discontinuity of function L . Example 3.10 highlights the relevance of different notions of derivatives and the resulting regularity. The existence of F-derivatives and associated smoothness requires the derivative object to be continuous w.r.t. more variations of distributions whereas the notion of G-derivatives and smoothness considers variations only along a line joining any pair of distributions. Luckily, since the FW algorithm operates by taking convex combinations at each iteration, we only require the latter bounds which are considerably less restrictive.

3.4. G-derivatives in non-flat spaces

In the following, we consider a slightly general setting wherein the ambiguity set \mathcal{P} may not be convex, e.g., when $\mathcal{P} \subset \mathcal{N}(\mathbb{R}^n)$ is the set of n multivariate Gaussian distributions. In such a setting, given two distributions $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$, the line joining them $\mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$ is not contained in \mathcal{P} . Even though such examples cannot be directly handled in our setting, it turns out that a slight modification of the definition of the derivative allows us to take care of such scenarios. To this end, we assume that the ambiguity set \mathcal{P} is equipped with

- (a) *Metric*: $d : \mathcal{P} \times \mathcal{P} \rightarrow [0, +\infty[$
- (b) *Geodesics*: For any $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$, there exists a parametric curve $[0, 1] \ni \gamma \mapsto \mathbb{P}_\gamma^d \in \mathcal{P}$ such that $\mathbb{P}_0^d = \mathbb{P}$, $\mathbb{P}_1^d = \mathbb{Q}$, and $d(\mathbb{P}_\gamma^d, \mathbb{P}_{\gamma'}^d) = |\gamma - \gamma'| d(\mathbb{P}, \mathbb{Q})$.

For example, if $\mathcal{P} = \mathcal{N}(\mathbb{R}^n)$, then each distribution is uniquely identified by its first and second moments μ and Σ respectively. Then, for given two distributions $\mathbb{P} = \mathcal{N}(\mu, \Sigma)$ and $\mathbb{P}' = \mathcal{N}(\mu', \Sigma')$, we define the metric $d(\mathbb{P}, \mathbb{P}') := \sqrt{\|\mu - \mu'\|_2^2 + \|\Sigma - \Sigma'\|_F^2}$, and the associated geodesic is $\mathbb{P}_\gamma^d = \mathcal{N}(\mu_\gamma, \Sigma_\gamma)$, where $(\mu_\gamma, \Sigma_\gamma) = (1 - \gamma)(\mu, \Sigma) + \gamma(\mu', \Sigma')$.

G-derivatives and smoothness. Given a geodesic structure on \mathcal{P} , we define the associated G-derivatives by

$$dR(\mathbb{P}; \mathbb{Q}) := \lim_{\gamma \downarrow 0} \frac{1}{\gamma} \left(R(\mathbb{P}_\gamma^d) - R(\mathbb{P}) \right).$$

Observe that the above definition is slightly different from that of (9) where the convex combination \mathbb{P}_γ of distributions \mathbb{P} and \mathbb{Q} is replaced with its geodesic counterpart. Consequently, we can also define G-smoothness on non-flat spaces in a geodesic sense similar to Definition 3.6. We say that the risk measure R is C -smooth if there exists some $C \geq 0$ such that the inequality holds

$$dR(\mathbb{P}; \mathbb{P}_\gamma^d) + dR(\mathbb{P}_\gamma^d; \mathbb{P}) \leq \gamma^2 C, \quad \text{for any } \mathbb{P}, \mathbb{Q} \in \mathcal{P} \text{ and } \gamma \in [0, 1].$$

Geodesic concavity. We can define a notion of concavity using the geodesic structure that is slightly more general than the usual notion. We say that a risk measure R is geodesically concave on \mathcal{P} if the mapping $[0, 1] \ni \gamma \mapsto R(\mathbb{P}_\gamma^d)$ is concave for every $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$. It turns out that both of the results from Proposition 3.4 hold true for derivatives defined in a geodesic sense. This allows us to consider DRO problems like

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{N}} \mathbb{E}_{\mathbb{P}}[f(x, \xi)] - \lambda G(\mathbb{P}, \widehat{\mathbb{P}}),$$

where $G(\mathbb{P}, \widehat{\mathbb{P}})$ is the Gelbrich distance (or any other moment-based distance). We wish to emphasize that the FW algorithm must also be adapted to work with the geodesic derivatives instead.

4. THE FRANK-WOLFE ALGORITHM

Given the notion of G-derivative as in Definition 3.1, the Frank-Wolfe (FW) algorithm for (5) is an iterative procedure that involves solving the optimization problem

$$\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q}), \tag{21}$$

for a given $\mathbb{P} \in \mathcal{P}$, at each iteration. The FW-problem (21) is linear if and only if the G-derivative $dR(\mathbb{P}; \mathbb{Q})$ is affine in \mathbb{Q} , for every $\mathbb{P} \in \mathcal{P}$, translating the problem into the linear DRO class in (3). This is indeed the case for many interesting risk measures as seen in Example 3.3. Moreover, in all finite-dimensional optimization problems, the corresponding FW-problem (21) is always linear, which need not be the case for a generic infinite-dimensional optimization problem like (5).

Implementation of the FW algorithm only requires a well-defined notion of the G-derivative, and as seen in Lemma 3.2, such objects can be computed by means of a few finite-dimensional quantities in several problems of interest. Moreover, even though the FW-problem (21) is infinite-dimensional in nature, it admits tractable finite-dimensional convex reformulations for many relevant applications similar to linear DROs [23]. This is a compelling reason to investigate FW methods for optimization problems over probability distributions, particularly in the case of nonlinear risk measures as in (5).




FIGURE 3. FW-oracle.

4.1. FW oracle

The FW-oracle is a set-valued mapping $\mathcal{F} : \mathcal{P} \times [0, 1] \rightarrow 2^{\mathcal{P}}$ defined as

$$\mathcal{F}(\mathbb{P}, \gamma) := \left\{ \mathbb{Q}' \in \mathcal{P} \text{ such that } \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q}) \leq \gamma \delta C + dR(\mathbb{P}; \mathbb{Q}') \right\}. \tag{22}$$

Given the current iterate $\mathbb{P} \in \mathcal{P}$, the FW algorithm involves solving the FW-problem (21) to obtain its approximate solution $\mathbb{Q}' \in \mathcal{P}$, then the current iterate \mathbb{P} is updated by moving it towards \mathbb{Q}' as shown in Figure 3.

Additive accuracy of the oracle. The parameter $\delta > 0$ is an arbitrary positive number signifying the accuracy of the FW oracle, and C is the smoothness constant of the risk measure R as in Definition 3.6. It must be observed that in an iterative scheme to solve (5), it is typical that the stepsize sequence (γ_k) is monotonically decreasing and converges to 0. Therefore, it is also required that the FW-oracle solves the sub-problems (21) up to a greater precision as the iterations progress.

The Frank-Wolfe-gap. We refer to the quantity $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q})$ as the *Frank-Wolfe (FW)-gap* at \mathbb{P} , and it is crucial in defining a posteriori stopping criteria for the FW-algorithm. Along with

the distribution \mathbb{Q}' , we assume that the FW-oracle also provides access to the quantity $dR(\mathbb{P}; \mathbb{Q}')$, which is an approximate of the FW-gap at \mathbb{P} .

Assumption 4.1 (Accessibility of FW oracle). *We assume that the FW oracle (22) is computationally available, i.e., there exists a tractable approach to find a feasible solution from the set (22).*

To develop our algorithm for NDRO problems, we assume the accessibility of the FW oracle in Assumption 4.1. However, given a DRO problem at hand, one needs to ensure that the corresponding FW oracle is indeed computationally available. When the directional derivatives are linear in \mathbb{Q} , the corresponding FW problem is a linear worst-case distribution problem, akin to the linear DRO problems (3). The properties of the loss function $\ell(\xi)$ (suppressing the decision variable x in (3)) under which the respective linear DRO enjoys a tractable reformulation have been extensively studied in the literature. Next, we provide an example of this kind.

Remark 4.2 (Tractable FW oracle). *For the RR measures in Definition 2.1, the FW oracle is a linear DRO (3) with the loss function $\ell(\xi) = \langle \nabla r(\mathbb{E}_{\mathbb{P}}[L(\xi)]), L(\xi) \rangle$ (cf. (10)). If this loss can, for instance, be described as a sum of pointwise maximum of concave functions (i.e., $\ell(\xi) = \sum_{t \leq T} \max_{k \leq K} \ell_{tk}(\xi)$), we know that the linear DRO (3) under the Wasserstein ambiguity set \mathcal{P} and the norm transportation cost $\|\cdot\|$ has a tractable reformulation [27, Theorem 6.1]. In this light, a sufficient condition for the RR measure is when each element of the sufficient statistic vector $L(\xi)$ constitutes a pointwise maximum of finitely many concave functions and $\nabla r \geq 0$.*

We note that the sufficient statistics of the variance and entropic risks in Example 2.2 do not meet the piecewise concavity condition of Remark 4.2; they are convex quadratic and exponential function, respectively, over the desired support sets. It is interesting to see that in the case of entropic risk, the positivity $\nabla r \geq 0$ is, however, fulfilled (see Example 3.3(b)). Besides the tractability setting of Remark 4.2, the literature includes various other combinations of ambiguity sets and functionals of probability distributions supported on finite support [38], the moment-based ambiguity sets [10, 17, 51], and the metric-based ambiguity set [15, 23], to name but a few.

Lemma 4.3 (FW-one-step-bounds). *Consider (5) with a risk measure that is C -smooth in the sense of Definition 3.6 for some $C \geq 0$, and let R^* be its optimal value. Let \mathcal{F} be the corresponding FW-oracle as in (22) with an arbitrary accuracy parameter $\delta \geq 0$. For any $\mathbb{P} \in \mathcal{P}$, $\gamma \in [0, 1]$, and $\mathbb{Q}' \in \mathcal{F}(\mathbb{P}, \gamma)$, let $\mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q}' - \mathbb{P})$ be the one-step-ahead FW update from \mathbb{P} with a stepsize of γ . Then we have*

$$R^* - R(\mathbb{P}) \leq \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q}) \leq \frac{1}{\gamma} \left(R(\mathbb{P}_\gamma) - R(\mathbb{P}) \right) + \gamma C(1 + \delta). \quad (23)$$

Rearranging the inequality (23), the one-step improvement in sub-optimality is seen to be

$$R^* - R(\mathbb{P}_\gamma) \leq \gamma^2 C(1 + \delta) + (1 - \gamma)(R^* - R(\mathbb{P})). \quad (24)$$

4.2. FW convergence guarantees

The FW-algorithm seeks to solve the optimization problem (5) by iteratively solving the FW-problem (21) using a FW-oracle as in (22). More precisely, given a step size sequence $(\gamma_k)_k \subset [0, 1]$

and a distribution $\mathbb{P}_0 \in \mathcal{P}$, the FW-algorithm generates a sequence of distributions $(\mathbb{P}_k)_k$ such that

$$\text{FW-algorithm: } \mathbb{P}_{k+1} = \mathbb{P}_k + \gamma_k(\mathbb{Q}_k - \mathbb{P}_k), \quad \text{where } \mathbb{Q}_k \in \mathcal{F}(\mathbb{P}_k, \gamma_k). \quad (25)$$

It must be observed that the implementation of the FW algorithm (25) does not depend on the choice of a norm on the ambiguity set \mathcal{P} . Therefore, it is desirable to have a *norm-independent* analysis of the FW algorithm. To this end, we take inspiration from [19], which has a similar analysis for the convergence of FW-algorithm for finite-dimensional problems by working with the notion of *curvature co-efficient* instead of the canonical smoothness (16) used in [13]. It turns out that the notion of smoothness defined in Definition 3.6 is also amenable to a similar analysis of the FW-algorithm; with the advantage that (15) is more in-line as a generalized notion of smoothness from (16) and (17).

Proposition 4.4 (Apriori bounds). *Consider (5) with a risk measure that is C -smooth in the sense of Definition 3.6 for some $C \in [0, +\infty)$, and let R^* be its optimal value. For $k = 1, 2, \dots$, let $(\mathbb{P}_k)_k$ be the sequence of distributions obtained from the FW-algorithm (25) with a step size sequence $\gamma_k = \frac{2}{k+2}$ and some $\mathbb{P}_0 \in \mathcal{P}$. Then*

$$R^* - R(\mathbb{P}_k) \leq \frac{4C}{k+2}(1+\delta) \quad \text{for all } k \geq 1. \quad (26)$$

Regarding the FW algorithm (25), an important observation to be made is that the “complexity” of the distribution \mathbb{P}_k may, in general, increase with each update of FW-iteration. To make this more clear, suppose that the distributions of \mathbb{P}_k and \mathbb{Q}_k are both discrete with potentially different support sets. Then, it is straightforward to see that the support of the next iteration \mathbb{P}_{k+1} is the union of the two supports, and as such, its cardinality increases in each iteration. However, this issue can be avoided in the case of regular risk measures.

Remark 4.5 (Sufficient statistics & reduced FW update). *When the risk measure is regular in the sense of Definition 2.1 (i.e., $R(\mathbb{P}) = r(\mathbb{E}_{\mathbb{P}}[L(\xi)])$), we recall from (10) that the G -derivative $dR(\mathbb{P}; \mathbb{Q})$ is characterized entirely in terms of the finite-dimensional sufficient statistic vector $\mathbb{E}_{\mathbb{P}}[L(\xi)]$. The FW algorithm (25) reduces to the finite-dimensional update*

$$\mu_{k+1} = \mu_k + \gamma_k(\nu_k - \mu_k), \quad \text{where } \nu_k \in \{\mathbb{E}_{\mathbb{Q}}[L(\xi)] \in \mathbb{R}^m : \mathbb{Q} \in \mathcal{F}(\mathbb{P}, \gamma)\} \quad (27)$$

in which $\mathcal{F}(\mathbb{P}, \gamma)$ is the FW-oracle (22), and the initial condition is $\mu_0 = \mathbb{E}_{\mathbb{P}}[L(\xi)]$. When the FW-oracle (22) is computationally available (e.g., the tractable cases in Remark 4.2), it suffices to follow the reduced finite-dimensional update rule (27), instead of the infinite-dimensional update (25). In other words, the worst-case risk (5) can be solved via $\sup_{\mathbb{P} \in \mathcal{P}} R(\mathbb{P}) = r(\mu^*)$, where μ^* is the convergent point of the reduced FW-algorithm in (27). Thus, the finite vector $\mathbb{E}_{\mathbb{P}}[L(\xi)]$ is indeed “sufficient” for computing the worst-case risk via the FW algorithm.

We note that the sufficient statistics in Remark (4.5) becomes more involved when an additional decision such as x in (4) must be determined. This shall be addressed in the next section.

4.3. FW-gap based termination and aposteriori bounds

Suppose that we know some $C' \geq 0$ that satisfies the smoothness condition (15). Then for any given $\varepsilon > 0$, if the FW algorithm is run for $K \geq \left\lceil \frac{4C'(1+\delta)}{\varepsilon} \right\rceil - 2$ iterations with the stepsize sequence $\gamma_k = \frac{2}{k+2}$, then the last iterate \mathbb{P}_K is guaranteed to be ε sub-optimal in objective value. Thus, in principle, it suffices to only know some upper bound $C' \geq C$ for the smoothness constant. However, if finding C exactly is challenging, and the known upper bound C' is not tight; the theoretically guaranteed number of iterations required for ε -sub-optimality may not be practical. In such a setting, it turns out that the optimal value of the FW-problem (21) called the *FW-gap* provides a good measure to define aposteriori stopping criteria. Moreover, we will also see later for NDRO problems that terminating the FW algorithm when the FW-gap is small provides worst-case performance bounds in the context of DRO problems. We follow the analysis of [19, 9] by considering the FW algorithm under two regimes of stepsize sequence to obtain provable upper bounds on the FW-gap towards later iterations.

Since the oracle employed to solve the linear minimization sub-problems at each iteration is only accurate to some specified precision, the actual value of the FW-gap $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q})$ is never known exactly. However, at each iteration k , the FW-oracle does provide its approximate estimate $dR(\mathbb{P}; \mathbb{Q}_k)$, which satisfies the inequality $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q}) \leq \delta \gamma_k C + dR(\mathbb{P}; \mathbb{Q}_k)$. Thus, for a given value of $\varepsilon > 0$, we terminate the FW procedure by examining the quantity $dR(\mathbb{P}; \mathbb{Q}_k)$ such that the desired FW-gap inequality: $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}; \mathbb{Q}) \leq \varepsilon$, is satisfied.

Proposition 4.6 (Aposteriori bounds). *Consider (5) with a risk measure that is C -smooth in the sense of Definition 3.6 for some $C \geq 0$, and let R^* be its optimal value. Let $K \geq 1$ and let $(\mathbb{P}_k)_k$ be the sequence obtained from the FW-algorithm (25) using a diminishing stepsize $\gamma_k := \frac{2}{k+2}$ for $k = 0, 1, \dots, K-1$, and then a constant stepsize $\gamma_k = \frac{2}{K+2}$ for $k = K, K+1, \dots, 2K+1$. Finally, let $g_k = dR(\mathbb{P}; \mathbb{Q}_k)$ for $k = 1, 2, \dots, 2K+1$ be the sequence of approximate FW-gaps. There exists $\hat{k} \in \{K, K+1, \dots, 2K+1\}$ such that*

$$R^* - R(\mathbb{P}_{\hat{k}}) \leq \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbb{Q}) \leq \frac{2C(2+3\delta)}{K+2}, \quad (28)$$

and every such \hat{k} is recognised by verifying the inequality $g_{\hat{k}} \leq \frac{4C}{K+2}(1+\delta)$.

Remark 4.7 (Explicit error bounds). *Consider the setting of Proposition 4.6 with $K = K(\varepsilon) := \left\lceil \frac{2C}{\varepsilon}(2+3\delta) \right\rceil - 2$ for an $\varepsilon > 0$. Then, there exists $\hat{k} \in \{K(\varepsilon), K(\varepsilon)+1, \dots, 2K(\varepsilon)+1\}$ such that*

$$\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbb{Q}) \leq \varepsilon,$$

and every such \hat{k} is recognised by verifying the inequality $g_{\hat{k}} \leq \varepsilon \frac{2+2\delta}{2+3\delta}$.

4.4. FW stepsize selection

The last part of this section discusses two particular features of the stepsize rule (two-regimes and diminishing behavior) in Proposition 4.6.

Two-regimes stepsize. The two-regimes for γ_k in Proposition 4.6 turns out to be crucial to obtain provable guarantees that the FW-gap is bounded above in the later iterations. Even though such certificates are of independent interest in their own right, having such upper bounds is also essential in the context of DRO problems. An upper bound on the FW-gap at iteration k ensures that the performance of the decision x_k for the worst-case distribution is not “too-bad”.

Different stepsize selection for the FW algorithm. For generic feasible sets, the Frank-Wolfe algorithm requires that the stepsize sequence $(\gamma_k)_k$ be diminishing. Even though the risk measure is smooth, the direction \mathbf{Q}'_k may change discontinuously around the optimal solution if the ambiguity set \mathcal{P} has flat faces (like Wasserstein-1 balls). Therefore, the Frank-Wolfe algorithm in general requires diminishing stepsize to converge, and might not converge with constant stepsize unlike gradient descent algorithms. Even with the standard diminishing stepsize of $\gamma_k = 2/k+2$, the canonical FW algorithm is plagued with the *zig-zag* phenomena where the iterates keep oscillating around the optimal solution. To remedy these challenges of FW in finite-dimensional convex problems, various adaptive stepsize sequences have been proposed in the literature with provably better guarantees under some additional assumptions. In the following, we describe the main ideas of these stepsize selection rules in the context of (5).

- (a) *Demyanov-Rubinov (DR) stepsize:* A very interesting stepsize selection is due to [49, 48]

$$\gamma_k = \min \left\{ \frac{g_k}{2C}, 1 \right\}, \quad (29)$$

which adaptively tunes the stepsize γ_k using the current value of the FW-gap.

- (b) *Backtracking:* Suppose, it is relatively easy to evaluate the risk measure R , then one can employ the backtracking based DR step size due to [36, 35]. This stepsize selection rule adaptive tunes the smoothness constant C locally along the line segment $(\mathbf{P}_k, \mathbf{Q}'_k)$. This allows the backtracking step size to take larger steps than other methods. Specific way to adaptively tune the smoothness constant may vary, the specific rule in [36] can be summarised as follows. For fixed constants $\eta \in (0, 1), \tau \geq 1$, at each iteration the stepsize rule is

$$\begin{cases} \gamma_k = \min \left\{ \frac{g_t}{2\tau^t C_k}, 1 \right\} \text{ for the largest } t = 1, 2, \dots \text{ such that} \\ R(\mathbf{P}_{k+1}) \geq R(\mathbf{P}_k) + \gamma_k g_k - \gamma^2 \tau^t C_k. \end{cases} \quad (30)$$

The smoothness constant is also updated as $C_{k+1} = \eta \tau^t C_k$. Another variant [35, Assumption 6.1, Algorithm 1] of backtracking-based step size selection in the FW algorithm gives rise to linear/geometric convergence with further assumptions on the feasible set.

- (c) *Exact line search:* If it is easier to optimize the risk measure R on the line segment $(\mathbf{P}_k, \mathbf{Q}'_k)$, then one can also select the stepsize by exactly maximizing

$$\gamma_k = \operatorname{argmax}_{\gamma \in [0,1]} R(\mathbf{P}_k + \gamma(\mathbf{Q}'_k - \mathbf{P}_k)). \quad (31)$$

5. NONLINEAR DISTRIBUTIONALLY ROBUST OPTIMIZATION

Let us recall that a generic DRO problem is formulated as the min-max problem

$$F^* := \inf_{x \in \mathcal{X}} \sup_{\mathbf{P} \in \mathcal{P}} F(x, \mathbf{P}), \quad (32)$$

where $\mathcal{X} \subset \mathbb{R}^n$ is a closed convex set, denoting the set of feasible decisions, and \mathcal{P} denotes a given ambiguity set of distributions. A DRO problem is said to be *feasible* if $F^* < +\infty$, which happens if and only if there exists some $x \in \mathcal{X}$ such that $\sup_{\mathbb{P} \in \mathcal{P}} F(x, \mathbb{P}) < +\infty$. Our objective is to develop a framework to solve a generic DRO problem (32). Particularly, with emphasis on the case when $F(x, \mathbb{P})$ is nonlinear in \mathbb{P} for every $x \in \mathcal{X}$, to which we refer to (32) as a *nonlinear distributionally robust optimization* (NDRO) problem.

The way the variable x enters the risk F in (32) may have an impact on the scalability of the proposed FW algorithm. In particular, in the case of the regular risks in Definition 2.1, the important feature is whether the decision x influences the sufficient statistics of the risk (cf. Remark 4.5). This consideration leads to two classes of regular risk measures:

$$(i) F(x, \mathbb{P}) = r(x, \mathbb{E}_{\mathbb{P}}[L(\xi)]) \quad \text{and} \quad (ii) F(x, \mathbb{P}) = r(\mathbb{E}_{\mathbb{P}}[L(x, \xi)]). \quad (33)$$

It should be noted that class (i) in (33) is a special form of class (ii). For regular risk measures in class (i), the sufficient statistic $\mathbb{E}_{\mathbb{P}}[L(\xi)]$ is *not* influenced by the decision x , whereas this is not the case for general regular risk measures in class (ii). This subtle difference makes a significant impact on whether the FW algorithm can benefit from the notion of sufficient statistic described in Remark 4.5. The variance and entropic risks investigated in this article are indeed different in view of this feature.

Example 5.1 (Regular NDRO examples). The analogous examples of Example 2.2 in the NDRO context are the following:

(a) *Variance*: The variance $V(x, \mathbb{P})$ (cf. (6)) associated with the distribution \mathbb{P} and a decision x is

$$V(x, \mathbb{P}) := x^\top (\Sigma_{\mathbb{P}} - \mu_{\mathbb{P}} \mu_{\mathbb{P}}^\top) x, \quad \text{where} \quad \Sigma_{\mathbb{P}} := \mathbb{E}_{\mathbb{P}}[\xi \xi^\top] \text{ and } \mu_{\mathbb{P}} := \mathbb{E}_{\mathbb{P}}[\xi]. \quad (34)$$

(b) *Entropic risk measure*: The entropic risk measure $\mathcal{E}(x, \mathbb{P})$ (cf. (7)) of a multivariate distribution \mathbb{P} and decision x . If \mathbb{P} is a distribution with marginals \mathbb{P}_j for $j = 1, 2, \dots, n$, (i.e., the j -th component ξ_j is \mathbb{P}_j distributed). For a given set of positive *risk-aversion parameters* $(\theta_j)_{j=1}^n$ in $(0, +\infty)$, the associated risk is defined as

$$\mathcal{E}(x, \mathbb{P}) := \sum_{j=1}^n \frac{1}{\theta_j} \log \left(\mathbb{E}_{\mathbb{P}_j} [e^{-\theta_j x_j \xi_j}] \right). \quad (35)$$

It is straightforward to see that the NDRO variance risk (34) belongs to the class (i) in (33), and thus keeping its sufficient statistic (i.e., the first two moments) intact, while the NDRO entropic risk (35) belongs to the class (ii) in (33) where the vector $L(x, \xi)$ depends inseparably on x . Looking ahead at (64), we see that the proposed FW Algorithm 1 applied to variance risk measure simplifies to iterations over only sufficient statistic (μ, Σ) and x .

NDRO Dual problem. Associated with the DRO problem (32) is its dual problem:

$$\text{Dual problem: } F_* := \sup_{\mathbb{P} \in \mathcal{P}} \inf_{x \in \mathcal{X}} F(x, \mathbb{P}). \quad (36)$$

In general, we have *weak-duality* $F_* \leq F^*$ relating the optimal values of the primal problem (32) and its dual (36). If $F_* = F^*$ specifically, we say that *strong duality* holds between (32) and (36).

Moreover, suppose the DRO problem (32) and its dual (36) admit the solutions x^* and \mathbb{P}^* , i.e.,

$$x^* \in \operatorname{argmin}_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} F(x, \mathbb{P}) \quad \text{and} \quad \mathbb{P}^* \in \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}} \inf_{x \in \mathcal{X}} F(x, \mathbb{P}),$$

Then, the pair (x^*, \mathbb{P}^*) is said to be a *saddle point* solution to the problems (32) and (36), which is also characterized by the condition

$$\max_{\mathbb{P} \in \mathcal{P}} F(x^*, \mathbb{P}) = F(x^*, \mathbb{P}^*) = \min_{x \in \mathcal{X}} F(x, \mathbb{P}^*).$$

The existence of a saddle point is sufficient for strong duality to hold, however, it is not necessary. Therefore, whenever strong duality holds, we consider the slightly relaxed notion of an ε -sub-optimal saddle points as a solution concept for (32).

Definition 5.2 (ε -saddle point). Given $\varepsilon \geq 0$, a pair $(x_\varepsilon, \mathbb{P}_\varepsilon) \in \mathcal{X} \times \mathcal{P}$ is an ε -saddle point of the DRO problem (32), if it satisfies

$$-\varepsilon + \sup_{\mathbb{P} \in \mathcal{P}} F(x_\varepsilon, \mathbb{P}) \leq F(x_\varepsilon, \mathbb{P}_\varepsilon) \leq \varepsilon + \inf_{x \in \mathcal{X}} F(x, \mathbb{P}_\varepsilon). \quad (37)$$

It is worth noting that if $(x_\varepsilon, \mathbb{P}_\varepsilon)$ is an ε -saddle point, then we have the inequalities:

$$\begin{aligned} \sup_{\mathbb{Q} \in \mathcal{P}} F(x_\varepsilon, \mathbb{Q}) &\leq 2\varepsilon + \inf_{y \in \mathcal{X}} F(y, \mathbb{P}_\varepsilon) \leq 2\varepsilon + \inf_{y \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{P}} F(y, \mathbb{Q}) &&= 2\varepsilon + F^*, \quad \text{and} \\ \inf_{y \in \mathcal{X}} F(y, \mathbb{P}_\varepsilon) &\geq -2\varepsilon + \sup_{\mathbb{Q} \in \mathcal{P}} F(x_\varepsilon, \mathbb{Q}) \geq -2\varepsilon + \sup_{\mathbb{Q} \in \mathcal{P}} \inf_{y \in \mathcal{X}} F(y, \mathbb{Q}) &&= -2\varepsilon + F_*. \end{aligned}$$

In other words, if $(x_\varepsilon, \mathbb{P}_\varepsilon)$ is an ε -saddle point, then both x_ε and \mathbb{P}_ε are at most 2ε -sub-optimal to the DRO problem (32) and its dual (36) respectively. Consequently, the decision x_ε is guaranteed to be at most 2ε worse from the best decision that could have been made in a DRO framework.

Solving the minimization over x in the dual-problem (36) results in a maximization problem (potentially nonlinear) over the distributions

$$\sup_{\mathbb{P} \in \mathcal{P}} R(\mathbb{P}), \quad \text{where } R(\mathbb{P}) := \inf_{x \in \mathcal{X}} F(x, \mathbb{P}). \quad (38)$$

Denoting $x(\mathbb{P}) := \operatorname{argmin}_{x \in \mathcal{X}} F(x, \mathbb{P})$ (whenever a minimizer exists), for every $\mathbb{P} \in \mathcal{P}$, the proposed method to compute an ε -saddle point of the DRO problem generates a sequence $(x_k, \mathbb{P}_k)_k$, where $x_k \in x(\mathbb{P}_k)$ for each k , and the sequence of distributions $(\mathbb{P}_k)_k$ is obtained by applying the FW algorithm to the maximization problem (38). If a pair (x', \mathbb{P}') satisfies $\mathbb{P}' \in \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}} R(\mathbb{P})$ and $x' \in x(\mathbb{P}')$, then it is not guaranteed that $x' \in \operatorname{argmin}_{y \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} F(y, \mathbb{P})$ unless $x(\mathbb{P}')$ is unique. It so turns out that the regularity assumptions on F and \mathcal{P} , required for the algorithm convergence, also ensure uniqueness.

5.1. NDRO: continuity, derivatives and smoothness

Our proposed method to solve the NDRO problem (32) by applying the FW algorithm to (38) requires that the risk measure R therein has well defined G-derivatives that are also smooth in the sense of Definition 3.6. This is not guaranteed apriori. In the following, we impose some regularity assumptions on the function F that guarantee the required smoothness of the risk measure R .

Assumption 5.3 (NDRO smoothness). Let $\mathbb{P}_\gamma := \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$ as in (9) and denote $F_x(\cdot) := F(x, \cdot)$, for every $x \in \mathcal{X}$. We assume that there exists positive constants α, C_1, C_2 such that

- (i) *Continuous function: The mapping $\mathcal{X} \times [0, 1] \ni (x, \gamma) \mapsto F(x, \mathbb{P}_\gamma)$ is proper, convex-concave, and continuous for all $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$.*
- (ii) *Continuous derivatives: The function $F_x(\mathbb{P})$ is directionally differentiable on \mathcal{P} , and the G-derivative $dF_x(\mathbb{P}; \mathbb{Q})$ is C_1 -Lipschitz continuous in x uniformly over $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$, i.e.,*

$$dF_x(\mathbb{P}; \mathbb{Q}) - dF_y(\mathbb{P}; \mathbb{Q}) \leq C_1 \|x - y\|, \quad \forall x, y \in \mathcal{X} \quad \forall \mathbb{P}, \mathbb{Q} \in \mathcal{P}. \quad (39)$$

- (iii) *Smoothness: The function $F_x(\mathbb{P})$ is C_2 -smooth in the sense of Definition 3.6, uniformly over $x \in \mathcal{X}$.*
- (iv) *Strong convexity: The function $F(x, \mathbb{P})$ is α -strongly convex in x w.r.t. the norm $\|\cdot\|$, uniformly over all $\mathbb{P} \in \mathcal{P}$, i.e.,*

$$\frac{\alpha}{2} \|x - y\|^2 \leq F(y, \mathbb{P}) - F(x, \mathbb{P}) - \langle \nabla_1 F(x, \mathbb{P}), y - x \rangle, \quad \forall x, y \in \mathcal{X}, \quad \forall \mathbb{P}, \mathbb{Q} \in \mathcal{P}.^2 \quad (40)$$

Remark 5.4 (Choice of norm on \mathcal{X}). *It must be noted that the norm $\|\cdot\|$ considered in the strong convexity assumption (40) and the continuity assumption (39) is identical. Considering an identical norm is not restrictive since all norms on \mathcal{X} are equivalent. However, using such equivalence often makes the resulting constants α, C_1 to be dimension dependent (of \mathcal{X}). We emphasize here that the smoothness constant given in Lemma 5.5, requires that the constants α , and C_1 that satisfy conditions (39) and (40), to satisfy with an identical norm.*

For now, we assume that these conditions for the abstract problem (32) are satisfied. However, for specific problems like the entropic or variance risk minimization (Sections 6 and 7, respectively), we will determine verifiable conditions whenever possible so that the conditions in Assumption 5.3 are indeed satisfied, (see Lemmas 6.3 and 7.2). Informally, condition (39) together with the smoothness condition is akin to saying that the directional derivatives of $F_x(\mathbb{P})$ are similar to being “Lipschitz continuous” with respect to both x and \mathbb{P} . Moreover, the strong convexity assumption implies that the mapping $\mathbb{P} \mapsto x(\mathbb{P})$ is also similar to being “Lipschitz continuous”. These two consequences together, yield the smoothness of R . The following lemma formally establishes this deduction in a norm-independent (in \mathbb{P}) analysis.

Lemma 5.5 (NDRO-derivative properties). *Let the function F satisfy Assumptions 5.3 with constants $\alpha > 0$, and $C_1, C_2 \geq 0$, then the following holds for the risk measure R as defined in (38)*

- (i) *Danskin’s theorem: The risk measure R is directionally differentiable on \mathcal{P} , and for any $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$ its \mathbb{Q} -directional derivative $dR(\mathbb{P}; \mathbb{Q})$ at \mathbb{P} , is given by*

$$dR(\mathbb{P}; \mathbb{Q}) = dF_{x(\mathbb{P})}(\mathbb{P}; \mathbb{Q}). \quad (41)$$

- (ii) *Smoothness: The risk measure R is C -smooth in the sense of Definition 3.6 for*

$$C = C_2 + \frac{C_1}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right).$$

² $\nabla_1 F(x, \mathbb{P}) := (\partial F / \partial x)(x, \mathbb{P})$ denotes the partial derivative of F w.r.t. x evaluated at (x, \mathbb{P}) .

5.2. Frank-Wolfe based algorithm for the NDRO problem.

Let \mathbb{P}_k for $k = 0, 1, 2, \dots$, be the sequence of iterates generated by the FW-algorithm (25) for R as defined in (38). Assume that the FW oracle is δ -accurate, for some arbitrary $\delta > 0$. To produce a solution x_ε to the NDRO problem (32) for a given $\varepsilon > 0$, one must decide the number of iterations $K(\varepsilon)$, the stepsize sequence γ_k for $k = 0, 1, \dots, K(\varepsilon)$, and the stopping criteria for the FW-algorithm. Algorithm 1 and the related next theorem present the main result of the article that provides a solution to the NDRO problem (32).

Algorithm 1: FW algorithm for NDRO problem (32)

Input: A distribution $\mathbb{P} \in \mathcal{P}$, positive real numbers ε and C , and access to a FW-oracle corresponding to the function F .

Output: The final decision variable x_ε and worst case distribution \mathbb{P}_ε .

1 **Initialization:** $\mathbb{P}_0 := \mathbb{P}$

diminishing stepsize regime

2 **for:** $k = 0, 1, 2, \dots, K(\varepsilon) := \lceil \frac{2C}{\varepsilon}(2 + 3\delta) \rceil - 2$

stepsize: $\gamma_k := \frac{2}{k+2}$

find $(x_k, \mathbb{Q}_k) \in \mathcal{X} \times \mathcal{P}$ such that

$$F(x_k, \mathbb{P}_k) = \min_{x \in \mathcal{X}} F(x, \mathbb{P}_k) \quad \text{and} \quad \sup_{\mathbb{Q} \in \mathcal{P}} dF_{x_k}(\mathbb{P}_k; \mathbb{Q}) \leq \delta \gamma_k C + dF_{x_k}(\mathbb{P}_k; \mathbb{Q}_k) \quad (42)$$

$$g_k = dF_{x_k}(\mathbb{P}_k; \mathbb{Q}_k) \quad \text{and} \quad \mathbb{P}_{k+1} := \mathbb{P}_k + \gamma_k(\mathbb{Q}_k - \mathbb{P}_k)$$

3 **end for**

constant stepsize regime

4 **for** $k = K(\varepsilon) + 1, \dots, 2K(\varepsilon) + 1$

stepsize: $\gamma_k := \frac{2}{K(\varepsilon)+2}$

if $g_k > \varepsilon \frac{2+2\delta}{2+3\delta}$, **do**

find $(x_k, \mathbb{Q}_k) \in \mathcal{X} \times \mathcal{P}$ based on (42)

$$g_k = dF_{x_k}(\mathbb{P}_k; \mathbb{Q}_k) \quad \text{and} \quad \mathbb{P}_{k+1} := \mathbb{P}_k + \gamma_k(\mathbb{Q}_k - \mathbb{P}_k)$$

else

Output: $x_\varepsilon := x_k$ and $\mathbb{P}_\varepsilon := \mathbb{P}_k$ and **end for.**

5 **end for**

Theorem 5.6 (NDRO solution). *Consider the Nonlinear DRO problem (32) under the setting of Assumptions 5.3. Then the following holds*

(i) Strong duality:

$$F^* = \min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} F(x, \mathbb{P}) = \sup_{\mathbb{P} \in \mathcal{P}} \min_{x \in \mathcal{X}} F(x, \mathbb{P}) = F_* \quad (43)$$

(ii) Saddle point computation: *Given any $\varepsilon > 0$, the pair $(x_\varepsilon, \mathbb{P}_\varepsilon)$ computed from Algorithm 1 is an ε -saddle point in the sense of Definition 5.2.*

Approaching an NDRO problem via the FW-based approach in Algorithm 1 requires the following three key ingredients:

- (i) *Regularity conditions of Assumption (5.3)*: To ensure the convergence of the FW algorithm, we must ensure that the corresponding G-derivatives exist, and satisfy regularity conditions of Assumption 5.3, particularly (ii) *continuous derivatives* and (iii) *smoothness*.
- (ii) *Solver for $\min_{x \in \mathcal{X}} F(x, \mathbb{P})$ in (42)*: Since it is a finite-dimensional convex problem, any well known first-order methods like ISTA, FISTA [1], Nesterov's Accelerated Gradient Descent [30], Extra Gradient [22] can be applied. Moreover, since the risk function F is assumed to be strongly convex in x (or at least with a regulariser), these first-order methods solve the corresponding minimization problems with geometric convergence.
- (iii) *Feasibility of the FW oracle (42)*: The FW worst-case distribution problem must admit tractable reformulations. In fact, when the directional derivatives are linear in the distribution, the corresponding FW problem is indeed tractable for several interesting choices of risk measures and ambiguity sets, as discussed in [38, 23, 15].

5.3. Slower convergence without strong-convexity

In this case, we assume that the function F satisfies conditions (ii) and (iii) of Assumption 5.3. However, it may not be necessarily strongly convex in x . For example, the variance risk measure is strongly convex if and only if the smallest eigenvalue of the matrix $(\Sigma_{\mathbb{P}} - \mu_{\mathbb{P}}\mu_{\mathbb{P}}^{\top})$ is bounded away from 0 uniformly over $\mathbb{P} \in \mathcal{P}$, which might not be the case. Even in such a setting, we desire to develop methods that compute an ε -saddle point of F using the setup of Algorithm 1 for any $\varepsilon > 0$. We take inspiration from the smoothing techniques in the optimization literature [32] for smoothing a non-smooth convex function and devise similar techniques that work with a suitable strongly-convex approximation F_{ε} , of F , and still use Algorithm 1 to compute an ε -saddle point of F . To this end, we assume that the set \mathcal{X} is also bounded in addition to being closed and, thus, compact. Many common examples of \mathcal{X} like the simplex Δ^n , satisfy the compactness assumption.

For any given an $\varepsilon > 0$, $x \in \mathcal{X}$, and $\mathbb{P} \in \mathcal{P}$, let $F_{\varepsilon}(x, \mathbb{P}) := F(x, \mathbb{P}) + (\varepsilon/B_x^2) \|x\|^2$. We propose to solve the following min-max problem in place of (32)

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} F_{\varepsilon}(x, \mathbb{P}), \quad (44)$$

where $B_x := \max_{x \in \mathcal{X}} \|x\|$. It is apparent at once that $F_{\varepsilon}(x, \mathbb{P})$ is $(2\varepsilon/B_x^2)$ -strongly convex in x , uniformly over $\mathbb{P} \in \mathcal{P}$, and consequently satisfies all the conditions in Assumption 5.3. Thus, employing Algorithm 1 with F_{ε} computes a pair (x', \mathbb{P}') that satisfies the inequalities

$$-\varepsilon + \sup_{\mathbb{Q} \in \mathcal{P}} F_{\varepsilon}(x', \mathbb{Q}) \leq F_{\varepsilon}(x', \mathbb{P}') \leq \varepsilon + \inf_{y \in \mathcal{X}} F_{\varepsilon}(y, \mathbb{P}'). \quad (45)$$

It turns out that such a pair (x', \mathbb{P}') is an ε -saddle point of F as well. This is easily seen by observing that on the one hand, we have

$$\begin{aligned} (\varepsilon/B_x^2) \|x'\|^2 + \sup_{\mathbb{Q} \in \mathcal{P}} F(x', \mathbb{Q}) &= \sup_{\mathbb{Q} \in \mathcal{P}} F_{\varepsilon}(x', \mathbb{Q}) \leq \varepsilon + F_{\varepsilon}(x', \mathbb{P}') \quad (\text{from (45)}) \\ &= \varepsilon + (\varepsilon/B_x^2) \|x'\|^2 + F(x', \mathbb{P}'). \end{aligned}$$

Thus, we immediately get $\sup_{\mathbb{Q} \in \mathcal{P}} F(x', \mathbb{Q}) \leq \varepsilon + F(x', \mathbb{P}')$. On the other hand, since the minimization over x is solved exactly in Algorithm 1, we have

$$\begin{aligned} F_\varepsilon(x', \mathbb{P}') &= \min_{y \in \mathcal{X}} F_\varepsilon(y, \mathbb{P}') = \min_{y \in \mathcal{X}} F(y, \mathbb{P}') + (\varepsilon/B_x^2) \|y\|^2 \\ &\leq \varepsilon + \min_{y \in \mathcal{X}} F(y, \mathbb{P}') \quad \text{since } \|y\| \leq B_x \text{ for all } y \in \mathcal{X}. \end{aligned}$$

Thus, it is apparent that the pair (x', \mathbb{P}') is an ε -saddle point of the function F as well.

Corollary 5.7 (Slower convergence). *Consider a function $F(x, \mathbb{P})$ that is not necessarily strongly convex in x . Then for any desired precision $\varepsilon > 0$, an ε -saddle point of F can be computed by applying Algorithm 1 with $K(\varepsilon) = \lceil \frac{2C_\varepsilon}{\varepsilon}(2 + 3\delta) \rceil - 2$, to the strongly convex approximate function F^ε , where $C_\varepsilon = C_2 + \frac{C_1 B_x^2}{4\varepsilon} \left(C_1 + \sqrt{C_1^2 + \frac{8\varepsilon}{B_x^2} C_2} \right)$.*

Since the strong convexity parameter of F_ε itself depends on ε , we conclude from Lemma 5.5 that the risk measure $R_\varepsilon(\mathbb{P}) := \min_{y \in \mathcal{X}} F_\varepsilon(y, \mathbb{P})$, also has an ε dependent smoothness constant $C_\varepsilon = C_2 + \frac{C_1 B_x^2}{4\varepsilon} \left(C_1 + \sqrt{C_1^2 + \frac{8\varepsilon}{B_x^2} C_2} \right)$. Now, Algorithm 1 terminates in $O(K(\varepsilon))$ iterations, where $K(\varepsilon) = O(C_\varepsilon/\varepsilon)$. Since $C_\varepsilon = O(1/\varepsilon)$, it is easily seen that for non-strongly convex functions F , applying Algorithm 1 to its regularized strongly-convex approximation F_ε takes $O(1/\varepsilon^2)$ iterations to compute an ε -saddle point of F . To compare, recall that for a strongly-convex function F , Algorithm 1 takes $O(1/\varepsilon)$ iterations to compute an ε -saddle point. This trade-off between speed of convergence and precision in the approximation is a typical occurrence in standard smoothing techniques as well.

6. ENTROPIC RISK PORTFOLIO SELECTION

This section is dedicated to study the proposed methodology and its required conditions for the entropic risk (35). To this end, let $\xi(t)$, $t = 1, 2, \dots, T$ be i.i.d. samples drawn from some unknown underlying distribution \mathbb{P}_o supported on $\Xi = \mathbb{R}^n$. We assume that the individual components $\xi_i(t)$, $i = 1, 2, \dots, n$, are independently distributed from each other. Let $\widehat{\mathbb{P}}_i := \frac{1}{T} \sum_{t=1}^T \delta(\xi_i(t))$ for $i = 1, 2, \dots, n$, and let $\widehat{\mathbb{P}} = \prod_{i=1}^n \widehat{\mathbb{P}}_i$. For $c > 0$ and any two distributions \mathbb{P}, \mathbb{Q} supported on \mathbb{R} , let $\text{WD}_c(\mathbb{P}, \mathbb{Q})$ be a Wasserstein distance between them defined as

$$\text{WD}_c(\mathbb{P}, \mathbb{Q}) := \begin{cases} \sup_{\pi} & \mathbb{E}_{\pi} \left[e^{c|u-v|} \right] \quad \text{where } (u, v) \text{ are } \pi\text{-jointly distributed} \\ \text{subject to} & \mathbb{P}(u) = \int_v \pi(u, v) dv \quad \text{and} \quad \mathbb{Q}(v) = \int_u \pi(u, v) du. \end{cases} \quad (46)$$

For each $j = 1, 2, \dots, n$, consider $\mathcal{W}_c((\widehat{\mathbb{P}}_j, \rho)) := \{\mathbb{P} : \text{WD}_c(\mathbb{P}, \widehat{\mathbb{P}}_j) \leq \rho\}$ and let $\mathcal{P}_c = \prod_{j=1}^n \mathcal{W}_c((\widehat{\mathbb{P}}_j, \rho))$. Let $\theta_j \in (0, 1)$ for $j = 1, 2, \dots, n$ be the risk aversion parameters, and for $c > \max\{\theta_j : j = 1, 2, \dots, n\}$, we seek to solve

$$\min_{x \in \Delta^n} \sup_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x, \mathbb{P}) := \sum_{j=1}^n \frac{1}{\theta_j} \log \left(\mathbb{E}_{\mathbb{P}_j} [e^{-\theta_j x_j \xi_j}] \right). \quad (47)$$

We will study (47) in detail for the various notions of the directional derivative, smoothness, and the resulting FW-oracle with its tractable formulations.

6.1. Regularity conditions

Recalling the entropic risk (7), we note that in the optimization (47) the mapping $\mathbb{P} \mapsto \mathcal{E}(x, \mathbb{P})$ is an RR measure in the sense of Definition 2.1 with functions L and r given by

$$L(x, \xi) = \left(e^{-\theta_1 x_1 \xi_1}, e^{-\theta_2 x_2 \xi_2}, \dots, e^{-\theta_n x_n \xi_n} \right) \quad \text{and} \quad r(z) = \sum_{j=1}^n \frac{1}{\theta_j} \log(z_j).$$

Thus, in view of Lemma 3.2 and using the short-hand notation $\mathcal{E}_x(\cdot) := \mathcal{E}(x, \cdot)$, we can write the directional derivatives of the risk measure \mathcal{E}_x as

$$d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}) = \sum_{j=1}^n \frac{\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j} [e^{-\theta_j x_j \xi_j}]}{\theta_j \mathbb{E}_{\mathbb{P}_j} [e^{-\theta_j x_j \xi_j}]} \quad \text{for every } \mathbb{P}, \mathbb{Q} \in \mathcal{P}. \quad (48)$$

It turns out that for any saddle point (x^*, \mathbb{P}^*) of the problem (47), \mathbb{P}^* belongs to a strictly smaller set \mathcal{P}'_c contained in \mathcal{P}_c and is of bounded support. The following Definition characterizes the smaller ambiguity set and the subsequent Lemma 6.2 formalizes this assertion.

Definition 6.1 (Restricted ambiguity set). *For each $j = 1, 2, \dots, n$, let*

$$\underline{\xi}_j := -\rho - \frac{\log(T)}{c} + \min_{t=1, \dots, T} \xi_j(t) \quad \text{and} \quad \bar{\xi}_j = \max_{t=1, \dots, T} \xi_j(t),$$

let $\mathcal{W}'_j \subset \mathcal{W}_c((\widehat{\mathbb{P}}_j, \rho))$ be the collection of distributions supported on $[\underline{\xi}_j, \bar{\xi}_j]$, and let $\mathcal{P}'_c := \prod_{j=1}^n \mathcal{W}'_j$.

Lemma 6.2 (Restricted ambiguity set). *A pair (x^*, \mathbb{P}^*) is a saddle point of (47) if and only if it is also a saddle point of*

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}'_c} \mathcal{E}(x, \mathbb{P}), \quad (49)$$

The crucial consequence of Lemma 6.2 is that it allows us to conclude regularity conditions for the risk measure $\mathcal{E}(x, \mathbb{P})$ by restricting our analysis to the smaller ambiguity set \mathcal{P}'_c . It turns out that the *continuous derivatives* and *smoothness* conditions (ii) and (iii) of Assumption 5.3 respectively, are not satisfied for the entropic risk portfolio optimization problem (47) on the entire ambiguity set \mathcal{P} but only on the smaller set \mathcal{P}'_c . Lemma 6.2 ensures that we can indeed restrict the ambiguity set from \mathcal{P}_c to \mathcal{P}'_c without losing any optimal solution. The following lemma establishes the regularity conditions.

Lemma 6.3 (Entropic risk regularity conditions). *Consider the entropic risk portfolio optimization problem (47). Let $\mathcal{E}_x(\cdot) := \mathcal{E}(x, \cdot)$ for every $x \in \mathcal{X}$, then the following assertions hold*

(i) Continuous derivatives: *The directional derivatives $d\mathcal{E}_x(\mathbb{P}; \mathbb{Q})$, satisfy*

$$d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}) - d\mathcal{E}_y(\mathbb{Q}; \mathbb{P}) \leq \|x - y\|_2 \sqrt{\sum_{j=1}^n (\bar{\xi}_j - \underline{\xi}_j)^2 e^{4\theta_j(\bar{\xi}_j - \underline{\xi}_j)}}, \quad \forall x, y \in \mathcal{X} \quad \forall \mathbb{P}, \mathbb{Q} \in \mathcal{P}'_c.$$

(ii) Smoothness: *The risk measure \mathcal{E}_x is C -smooth in the sense of Definition 3.6 on \mathcal{P}'_c and uniformly over $x \in \mathcal{X}$ for $C = \sum_{j=1}^n \frac{1}{\theta_j} (e^{\theta_j(\bar{\xi}_j - \underline{\xi}_j)} - 1)^2$.*

Besides the *smoothness* and *continuity* conditions, the uniform strong-convexity assumption is extremely difficult to verify for the entropic risk measure. Therefore, as discussed in Section 5.3, we add a strongly-convex regularizer of x to apply Algorithm 1.

6.2. The FW Oracle

The FW problem for minimum entropic risk portfolio selection (47), at a given distribution \mathbb{P} is

$$\sup_{\mathbb{Q} \in \mathcal{P}_c} \sum_{j=1}^n \frac{1}{\theta_j} \frac{\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}.$$
 (50)

Since $\mathbb{Q} = \Pi_{j=1}^n \mathbb{Q}_j$ and $\mathbb{Q}_j \in \mathcal{W}_c((\hat{\mathbb{P}}_j, \rho))$, the FW problem (50) for entropic risk is separable into n different problems, and for each $j = 1, 2, \dots, n$, we have

$$\sup_{\mathbb{Q}_j \in \mathcal{W}_c((\hat{\mathbb{P}}_j, \rho))} \mathbb{E}_{\mathbb{Q}_j}[e^{-\theta_j x_j \xi_j}].$$
 (51)

Lemma 6.4 (Entropic risk FW oracle). *Let $\theta > 0$, $x \in [0, 1]$, and $z(t)$, $t = 1, 2, \dots, T$, be any arbitrary collection of real numbers and let $\hat{\mathbb{P}} = \frac{1}{T} \sum_{t=1}^T \delta(z(t))$ be a given discrete distribution. Consider the following problem*

$$\sup_{\mathbb{Q} \in \mathcal{W}_c((\hat{\mathbb{P}}, \rho))} \mathbb{E}_{\mathbb{Q}}[e^{-\theta x z}].$$
 (52)

Let $(Z(t))_t$ be a non-increasing permutation of $(z(t))_t$, i.e., we have $Z(1) \geq Z(2) \geq \dots \geq Z(T)$. Let $T' \in \{1, 2, \dots, T\}$ be the smallest integer such that $(Te^{c\rho} - T')e^{\frac{-c\theta x Z(T')}{c-\theta x}} \geq \sum_{t=T'+1}^T e^{\frac{-c\theta x Z(t)}{c-\theta x}}$, define

$$\left\{ \begin{array}{l} \eta^* := \frac{\theta x}{c} \left(\frac{1}{Te^{c\rho} - T'} \sum_{t=T'+1}^T e^{\frac{-c\theta x Z(t)}{c-\theta x}} \right)^{\frac{c-\theta x}{\theta x}} \quad \text{and} \\ z^*(t) := \min \left\{ z(t), \frac{cz(t) + \log(c\eta^*/\theta x)}{c - \theta x} \right\} \quad \text{for } t = 1, 2, \dots, T. \end{array} \right.$$
 (53)

(i) Optimal solution: *The discrete distribution $\mathbb{Q}^* = \frac{1}{T} \sum_{t=1}^T \delta(z^*(t))$ is the optimal solution \mathbb{Q}^* to the linear worst-case distribution problem (52).*

(ii) Lower and upper bounds: *For any $t = 1, 2, \dots, T$ we have*

$$\underline{z} := -\rho - \frac{\log(T)}{c} + \min_{t=1, \dots, T} z(t) \leq z^*(t) \leq \bar{z} := \max_{t=1, \dots, T} z(t).$$

Corollary 6.5 (FW restricted ambiguity set). *For any $x \in \Delta^n$ and $\mathbb{P} \in \mathcal{P}$, we have*

$$\arg \max_{\mathbb{Q} \in \mathcal{P}_c} d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}) = \arg \max_{\mathbb{Q}' \in \mathcal{P}'_c} d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}').$$
 (54)

6.3. Simulation results

We validate the convergence properties of Algorithm 1 for the NDRO problem of the entropic risk portfolio selection (47) with unrestricted support (i.e., $\Xi = \mathbb{R}^n$). For each $j = 1, 2, \dots, n$, the samples $\xi_j(t)$, for $t = 1, 2, \dots, T = 2n$, are drawn randomly from bi-exponential distribution with density $f_j(\cdot) = \lambda_j e^{-\lambda_j |\cdot|}$ with each parameter λ_j drawn uniformly from $[0, 1]$, and the distribution $\hat{\mathbb{P}}$ is uniformly distributed over the samples drawn. The risk aversion parameter θ_j are drawn uniformly from $[0, 1]$ for each $j = 1, 2, \dots, n$, and we select $c = 1$ to define the Wasserstein distance in (46) and the resulting ambiguity set \mathcal{P} . For $n = 250$ and various values of ρ (the radius of the ambiguity set), we solve (47). The proposed FW-based method (Algorithm 1 with the FW-oracle of Proposition 6.4) is implemented in MATLAB on a Macbook Air (M1 with 8GB RAM), wherein the minimization problem: $\min_{x \in \Delta^n} \mathcal{E}(x, \mathbb{P}_k)$ is solved by running the FISTA algorithm [1] until convergence.

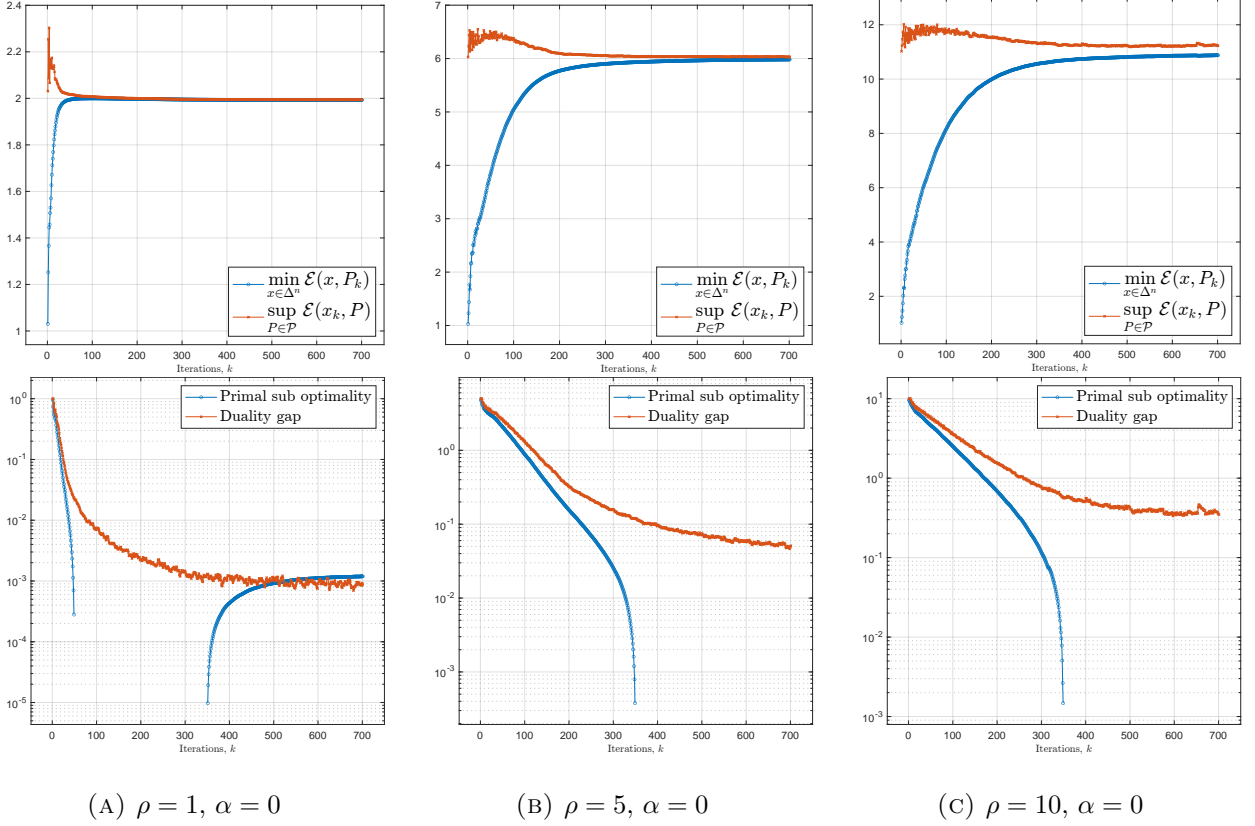


FIGURE 4. Convergence plots for Algorithm 1 with $K(\varepsilon) = 350$, applied to (47).

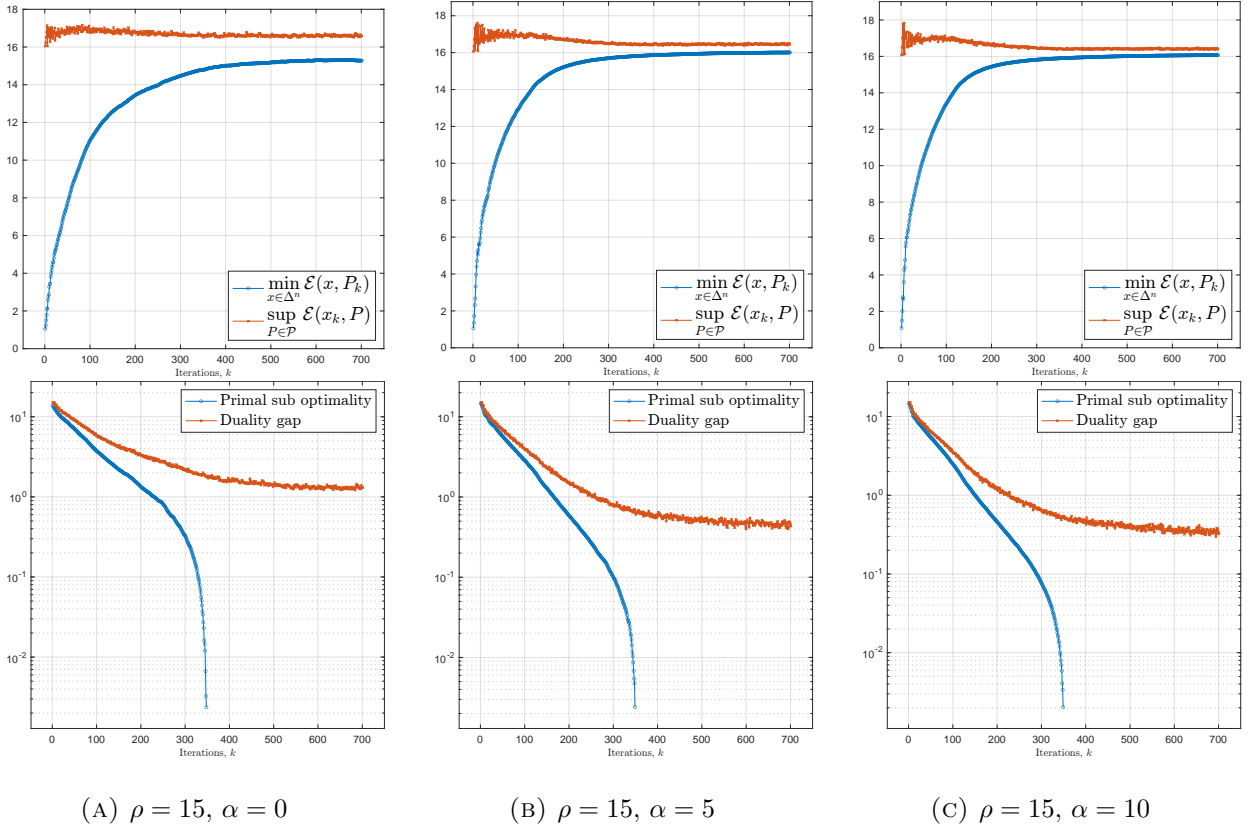
For three distinct values of $\rho = 1, 5$, and 10 , Figure 4 shows the convergence of Algorithm 1 with $K(\varepsilon) = 350$, on the entropic risk portfolio selection problem (47). The plots on the top show the evolution of the *Primal* and *Dual* functions: $\min_{x \in \Delta^n} \mathcal{E}(x, \mathbb{P}_k)$ and $\sup_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x_k, \mathbb{P})$ respectively, w.r.t. the iteration k of the algorithm. Whereas, the plots at the bottom of the figure show the evolution of the two sub-optimality metrics:

$$\left\{ \begin{array}{l} \text{primal sub-optimality} : \quad \mathcal{E}^* - \min_{x \in \Delta^n} \mathcal{E}(x, \mathbb{P}_k) \quad \text{and} \\ \text{duality gap} : \quad \sup_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x_k, \mathbb{P}) - \min_{x \in \Delta^n} \mathcal{E}(x, \mathbb{P}_k). \end{array} \right.$$

Since Algorithm 1 explicitly solves the minimization over x in each iteration (see (42)), the primal function is readily available. However, this is not the case with the dual function which needs to be computed independently at each iteration by solving $\sup_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}_{x_k}(\mathbb{P})$ while keeping the current iterate x_k fixed.

With explicit regularization of x . We observe in Figure 4 that, as the value of ρ increases, the convergence of the algorithm becomes extremely slow. This can be attributed to the fact that the strong-convexity parameter becomes extremely small, and thus the resulting smoothness constant is prohibitively large, making the convergence slow. To remedy this, as suggested in Section 5.3, we compute an approximate saddle point of (47) by applying the FW-algorithm to the explicitly regularized min-sup problem

$$\min_{x \in \Delta^n} \sup_{\mathbb{P} \in \mathcal{P}_c} \frac{\alpha}{2} \|x\|_2^2 + \mathcal{E}(x, \mathbb{P}). \quad (55)$$

FIGURE 5. Convergence plots for Algorithm 1 with $K(\varepsilon) = 350$, applied to (55).

Since $\|x\|_2 \leq \|x\|_1 = 1$ for all $x \in \Delta^n$, we know that an ε -saddle point of (70) can be computed for any $\varepsilon > 0$ by applying Algorithm 1 to (71) with $\alpha = 2\varepsilon$.

In Figure 5, we show the convergence plots of Algorithm 1 when applied to the explicitly regularized problem (55). We consider the same data set from Figure 4 but with a slightly larger value of $\rho = 15$, making the conditioning of the problem even worse than that for $\rho = 10$. We show the convergence plots of the algorithm for three different values of $\alpha = 0, 5, 10$. Clearly, for $\alpha = 0$, the problem is not regularized, and the convergence is bad (even worse than that for $\rho = 10$ from Figure 4c). Then the effect of explicit regularization and how it improves the regularity and convergence can be clearly seen in Figure 5b and 5c.

7. MINIMUM VARIANCE PORTFOLIO SELECTION

The minimum variance portfolio selection (34) is one of the textbook examples of NDRO [26, 37, 4, 34]. This section is dedicated to studying the NDRO regularity of this example and the different aspects of the proposed algorithm in this context. To this end, let $\hat{\xi}_i: i = 1, 2, \dots, N$ be i.i.d. samples drawn from some unknown underlying distribution \mathbb{P}_o , and let $\hat{\mathbb{P}}$ be the nominal distribution that is uniformly distributed over the samples $(\hat{\xi}_i)_i$. Let $\mathcal{P} = \mathcal{W}_m(\hat{\mathbb{P}}, \rho) := \{\mathbb{Q} : \text{WD}_m(\hat{\mathbb{P}}, \mathbb{Q}) \leq \rho\}$, where WD_m is the m -th order Wasserstein distance between the distributions,

induced by the transportation cost coming from a norm $\|\cdot\|$ on Ξ . More specifically, we have

$$\text{WD}_m(\mathbb{P}, \mathbb{Q}) := \begin{cases} \sup_{\pi} & \left(\mathbb{E}_{\pi}[\|u - v\|^m] \right)^{1/m} \quad \text{where } (u, v) \text{ are } \pi\text{-jointly distributed} \\ \text{subject to} & \mathbb{P}(u) = \int_v \pi(u, v) dv \quad \text{and} \quad \mathbb{Q}(v) = \int_u \pi(u, v) du. \end{cases}$$

Associated with the distribution \mathbb{P} , let $\Sigma_{\mathbb{P}} := \mathbb{E}_{\mathbb{P}}[\xi\xi^{\top}]$ and $\mu_{\mathbb{P}} := \mathbb{E}_{\mathbb{P}}[\xi]$ denote its first and second moments respectively. Finally, let $\mathcal{X} \subset \Delta^n$ denote the set of feasible actions, then we seek to investigate the min-sup problem:

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho)} V(x, \mathbb{P}) := x^{\top} (\Sigma_{\mathbb{P}} - \mu_{\mathbb{P}}\mu_{\mathbb{P}}^{\top})x. \quad (56)$$

We shall study in detail the notion of the directional derivative, smoothness, and the resulting FW-oracle with its tractable formulations for (56) under different conditions.

7.1. Regularity conditions

We recall that the variance risk (6) and note that the mapping $\mathbb{P} \mapsto V_x(\mathbb{P})$ is an RR measure in the sense of Definition 2.1 with the respective sufficient statistic L and the function r

$$L(\xi) = (\xi\xi^{\top}, \xi) \quad \text{and} \quad r(x, (\Sigma, \mu)) = x^{\top}\Sigma x - (x^{\top}\mu)^2.$$

Unlike the risk entropic in Section 6, the variance sufficient statistic L above is not influenced by the decision x (cf. the two regular risk classes in (33)). This is a particularly useful feature to control the complexity of the FW iteration (25) by only tracing the finite-dimensional sufficient statistic; see also Remark 4.5. In view of Lemma 3.2 and using the risk functions the short-hand notation $V_x(\cdot) := V(x, \cdot)$, we can describe the directional derivatives of V_x as

$$dV_x(\mathbb{P}; \mathbb{Q}) = \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[x^{\top}(\xi - 2\mu_{\mathbb{P}})\xi^{\top}x] \quad \text{for every } \mathbb{P}, \mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho). \quad (57)$$

Assumption 7.1 (Variance moments bound). *There exists constants $B_{\Sigma}, B_{\mu} \geq 0$ such that*

$$\|\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}}\|_o \leq B_{\Sigma} \quad \text{and} \quad \|\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}\|_2 \leq B_{\mu} \quad \text{for all } \mathbb{P}, \mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho), \quad (58)$$

where $\Sigma_{\mathbb{P}} = \mathbb{E}_{\mathbb{P}}[\xi\xi^{\top}]$ and $\mu_{\mathbb{P}} = \mathbb{E}_{\mathbb{P}}[\xi]$.

Lemma 7.2 (Variance regularity conditions). *Consider the minimum variance portfolio optimization problem (56), and suppose that Assumption 7.1 holds with constants B_{Σ}, B_{μ} . Let $V_x(\cdot) := V(x, \cdot)$ for every $x \in \mathcal{X}$, then the following assertions hold*

- (i) Continuous derivatives: For $C_1 = 2\left(B_{\Sigma} + 2(B_{\mu} + \|\hat{\mu}\|_2)^2 + B_{\mu}^2\right)$, the directional derivatives $dV_x(\mathbb{P}; \mathbb{Q})$, satisfy

$$dV_x(\mathbb{P}; \mathbb{Q}) - dV_x(\mathbb{Q}; \mathbb{P}) \leq C_1 \|x - y\|_2, \quad \forall x, y \in \mathcal{X} \quad \forall \mathbb{P}, \mathbb{Q} \in \mathcal{P}.$$

- (ii) Smoothness: The risk measure V_x is $(2B_{\mu}^2)$ -smooth in the sense of Definition 3.6, uniformly over $x \in \mathcal{X}$.

If Assumption 7.1 holds, then Lemma 7.2 ensures that the *continuous derivatives* and *smoothness* conditions (ii) and (iii) of Assumption 5.3 respectively, are satisfied for the minimum variance portfolio optimization problem (56). However, the uniform strong-convexity assumption need not

hold in general. For instance, if the ambiguity set is large (i.e., if ρ is large), then all the data points ξ_i can be perturbed within their respective ρ -neighborhoods such that the variance corresponding to the perturbed points is rank deficient. Thus, the strong-convexity condition in Assumption 5.3 is not satisfied at the distribution supported over the perturbed points. In such cases, adding an explicit strongly-convex regularizer to apply Algorithm 1 (as discussed in Section 5.3) not only provides theoretical guarantees for convergence but also improves the speed of the algorithm in our observation; see the numerical simulations concerning Figure 7.

7.2. The Frank-Wolfe algorithm

The FW-Oracle. Let us consider the FW-problem arising in the NDRO problem of minimum variance portfolio selection (56),

$$\sup_{\mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho)} dV_x(\mathbb{P}; \mathbb{Q}). \quad (59)$$

Since the directional derivatives $dV_x(\mathbb{P}; \mathbb{Q})$ are affine in \mathbb{Q} for every pair (x, \mathbb{P}) , the corresponding FW-problem is linear. However, the existence and characterization results of the solution to the FW-problem (59) change depending on the interplay of (i) the Wasserstein distance type m , (ii) the transportation cost $\|\cdot\|$, and (iii) the support set Ξ (unbounded, or compact). Therefore, the specific settings for which the corresponding FW-oracle is easy to describe are discussed later in the section. To this end, we simplify (59), and study its dual problem which is used later in the characterization of a solution to (59).

It is a simple algebraic exercise to verify that

$$\begin{aligned} dV_x(\mathbb{P}; \mathbb{Q}) &= \mathbb{E}_{\mathbb{Q}-\mathbb{P}} \left[x^\top ((\xi - 2\mu_{\mathbb{P}})\xi^\top) x \right] \\ &= \mathbb{E}_{\mathbb{Q}} \left[(\xi - \mu_{\mathbb{P}})^\top (xx^\top) (\xi - \mu_{\mathbb{P}}) \right] - x^\top (\Sigma_{\mathbb{P}} - \mu_{\mathbb{P}}\mu_{\mathbb{P}}^\top) x. \end{aligned}$$

Since the distribution \mathbb{P} is constant in (59), eliminating terms that only depend on \mathbb{P} does not affect the set of maximizers. Thus, we have

$$\operatorname{argmax}_{\mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho)} dV_x(\mathbb{P}; \mathbb{Q}) = \operatorname{argmax}_{\mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho)} \mathbb{E}_{\mathbb{Q}} \left[(\xi - \mu_{\mathbb{P}})^\top (xx^\top) (\xi - \mu_{\mathbb{P}}) \right]. \quad (60)$$

We now focus on a generic version of the linear worst-case distribution problem:

$$\sup_{\mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho)} \mathbb{E}_{\mathbb{Q}} [(\xi - v)^\top (xx^\top) (\xi - v)], \quad (61)$$

for any given $x, v \in \mathbb{R}^n$. We know that [23, Theorem 7] (61) admits an equivalent dual formulation given by

$$\inf_{\eta \geq 0} \eta \rho^m + \frac{1}{N} \sum_{i=1}^N \begin{cases} \sup_{q_i} & (q_i + \xi_i - v)^\top (xx^\top) (q_i + \xi_i - v) - \eta \|q_i\|^m \\ \text{subject to} & q_i + \xi_i \in \Xi \quad \text{for all } i = 1, 2, \dots, N. \end{cases} \quad (62)$$

For $\eta \geq 0$ and $\xi \in \Xi$, let

$$q(\eta, \xi, v) := \begin{cases} \operatorname{argmax}_q & ((x^\top q) + x^\top (\xi - v)^\top)^2 - \eta \|q\|^m \\ \text{subject to} & q + \xi \in \Xi, \end{cases} \quad (63)$$

whenever an optimal solution exists. If the dual problem (62) admits an optimal solution η_x , and $q(\eta_x, \xi_i, v)$ exists for each $i = 1, 2, \dots, N$; then for any collection $q'_i \in q(\eta_x, \xi_i, v)$, $i = 1, 2, \dots, N$, the discrete distribution $\mathbb{Q}_x(\xi) := \frac{1}{N} \sum_{i=1}^N \delta(\xi - \xi_i - q'_i)$ is a maximizer for the linear worst-case distribution problem (61).

Tractable one step FW-update. We observe that in the minimum variance problem (56), the only information needed pertaining to a distribution $\mathbb{Q} \in \mathcal{W}_m(\widehat{\mathbb{P}}, \rho)$ is its first and second order moments $\mu_{\mathbb{Q}} = \mathbb{E}_{\mathbb{Q}}[\xi]$, $\Sigma_{\mathbb{Q}} = \mathbb{E}_{\mathbb{Q}}[\xi\xi^\top]$ respectively. Thus, in order to solve the DRO problem (56) via the FW-algorithm (Algorithm 1), it is apparent that it suffices to track the evolution of these moments rather than the entire distribution (which gets increasingly difficult with the iterations). To this end, a single FW-update (42), in Algorithm 1, for the minimum-variance problem (56) can be summarised in terms of the finite-dimensional quantities:

$$\left\{ \begin{array}{l} \text{Solve for } x \quad : \quad x_k \in \underset{x \in \mathcal{X}}{\operatorname{argmin}} x^\top (\Sigma_k - \mu_k \mu_k^\top) x \\ \text{FW-oracle} \quad : \quad \begin{cases} \eta_k = \eta_{x_k} \quad \text{and} \quad q_i \in q(\eta_k, \xi_i, \mu_k) \quad \text{for all } i = 1, 2, \dots, N, \\ \mu'_k = \frac{1}{N} \sum_{i=1}^N (\xi_i + q_i) \quad \text{and} \quad \Sigma'_k = \frac{1}{N} \sum_{i=1}^N (\xi_i + q_i)(\xi_i + q_i)^\top, \end{cases} \\ \text{FW-update} \quad : \quad \mu_{k+1} = \mu_k + \gamma_k (\mu'_k - \mu_k), \quad \Sigma_{k+1} = \Sigma_k + \gamma_k (\Sigma'_k - \Sigma_k). \end{array} \right. \quad (64)$$

The tractable formulation (64) of the FW algorithm for variance particularly highlights the discussion in Remark 4.5 and (33) on the sufficient statistic for variance. In particular, it must be noted that the entire information of the distribution is characterized by (μ, Σ) , and therefore, the FW algorithm (Algorithm 1) can be simplified as (64) which is an iteration over only finite-dimensional quantities (μ, Σ) and x .

We now focus on specific settings of (56) for which the FW-oracle is easily characterized.

Case 1: unconstrained support. We consider the setting: $\Xi = \mathbb{R}^n$, $\mathcal{X} \in \Delta^n$ being any compact subset, and the transportation cost $\|\cdot\|$ to define the Wasserstein ambiguity set (and let $\|\cdot\|_*$ be the associated dual norm).

Lemma 7.3. *Consider the maximization problem (63) for $\Xi = \mathbb{R}^n$, and $\eta \geq 0$, $x, v, \xi \in \mathbb{R}^n$, and let J denote its optimal value. Then the following assertions hold*

- (i) (Unbounded): $J = +\infty$, and no optimal solution exists for (63) in the following cases: ($m < 2$), ($m > 2, \eta = 0$), ($m = 2, \eta < \|x\|_*^2$), and ($m = 2, \eta = \|x\|_*^2, x^\top(\xi - v) \neq 0$).
- (ii) (Bounded): $J < +\infty$, and the optimal solution $q(\eta, \xi, v)$ exists for (63) in all other cases, and in particular,
 - (a) ($m > 2, \eta > 0$), both $J < +\infty$, and $q(\eta, \xi, v)$ exists.
 - (b) ($m = 2, \eta > \|x\|_*^2$), then we have

$$J = \frac{\eta |x^\top(\xi - v)|^2}{\eta - \|x\|_*^2}, \quad q(\eta, \xi, v) = \frac{\|x\|_* (x^\top(\xi - v))}{\eta - \|x\|_*^2} \bar{q}_x \text{ for } \bar{q}_x \in \underset{\|\bar{q}\| \leq 1}{\operatorname{argmin}} x^\top \bar{q}.$$

- (c) ($m = 2, \eta = \|x\|_*^2, x^\top(\xi - v) = 0$), then $J = 0$ and $q(\eta, \xi, v) = \{s\bar{q}_x : s \geq 0\}$.

Lemma 7.4 (Variance FW-oracle: unconstrained support). *Consider the linear worst-case-distribution problem (61) and its dual (62), under the setting, $\Xi = \mathbb{R}^n$, $m = 2$, and any $x \in \mathcal{X}$. The solutions*

(\mathbb{Q}_x, η_x) to (61) and its dual (62) are given by

$$\left\{ \begin{array}{l} \eta_x = \|x\|_*^2 + \frac{\|x\|_*}{\rho} \sqrt{\frac{1}{N} \sum_{i=1}^N |x^\top (\xi_i - v)|^2}, \\ \mathbb{Q}_x(\xi) = \frac{1}{N} \sum_{i=1}^N \delta(\xi - (\xi_i + q'_i)) \quad \text{for any } q'_i \in q(\eta_x, \xi_i, v), \quad i = 1, 2, \dots, N. \end{array} \right. \quad (65)$$

It turns out that for the special case of $m = 2$, the problem (56) with unconstrained support reduces to a simple empirical risk minimization problem. This was already discovered in [4] specifically for the case when $\mathcal{X} = \mathcal{X}(\bar{\alpha}) := \{x \in \Delta^n : \mathbb{E}_{\mathbb{Q}}[x^\top \xi] \geq \bar{\alpha}, \text{ for all } \mathbb{Q} \in \mathcal{P}\}$, for any $\bar{\alpha}$. We generalize it slightly by showing that a similar conclusion holds for any compact set $\mathcal{X} \subset \Delta^n$. Moreover, we provide an alternate proof completely based on first-order optimality conditions and the FW-oracle, for the min-max problem (56). We emphasize here that since it is shown explicitly that the min-max problem (56) reduces to an empirical minimization problem, a solution can be computed without the need to run the FW-algorithm (64) iteratively. More importantly, this also alleviates the need to ensure that the convex regularity assumptions hold for this special case of (56).

Proposition 7.5 (Variance saddle point). *Consider the minimum variance portfolio optimization problem (56) in the setting of $m = 2$, and $\Xi = \mathbb{R}^n$. Let*

$$x^* \in \operatorname{argmin}_{x \in \mathcal{X}} \sqrt{\langle x, (\hat{\Sigma} - \hat{\mu}\hat{\mu}^\top)x \rangle} + \rho \|x\|_*, \quad (66)$$

then there exists some $\bar{q}_{x^*} \in \operatorname{argmax}_{\|\bar{q}\| \leq 1} \langle x^*, \bar{q} \rangle$ such that for

$$\mathbb{P}^* := \frac{1}{N} \sum_{i=1}^N \delta(\xi - (\xi_i + q_i^*)), \quad \text{where } q_i^* = \frac{\rho \langle x^*, \xi_i - \hat{\mu} \rangle}{\sqrt{\langle x^*, (\hat{\Sigma} - \hat{\mu}\hat{\mu}^\top)x^* \rangle}} \bar{q}_{x^*} \quad \forall i \leq N, \quad (67)$$

the pair (x^*, \mathbb{P}^*) is a saddle point to (56).

Case 2: ellipsoidal support. We consider the setting: $\Xi = \mathcal{E}_M := \{\xi : \langle \xi, M\xi \rangle \leq 1\}$ for some $M \succeq 0$, $m = 2$, and the underlying transportation cost defining the Wasserstein distance to be $\|\cdot\| = \|\cdot\|_2$. Under this setting, the FW problem in (42) involves maximizing a quadratic function subject to convex ellipsoidal constraints. Even if this problem is non-convex in general, it admits a tractable reformulation as an SDP via the celebrated *S-procedure* [6, Appendix B], [47].

Lemma 7.6 (Variance FW-oracle: ellipsoidal support). *Consider the minimum variance portfolio optimization problem (56) under the setting of the ellipsoidal support. The corresponding dual problem (62) of the linear worst-case-distribution problem (61) is equivalent to the SDP*

$$\left\{ \begin{array}{l} \min_{\eta \in \mathbb{R}, \lambda, \theta \in \mathbb{R}^N} \quad \eta \rho^2 - \frac{1}{N} \sum_{i=1}^N \theta_i \\ \text{subject to} \quad \left\{ \begin{array}{l} \eta \geq 0, \text{ and } \lambda_i \geq 0 \quad \forall i \leq N, \\ \left[\begin{array}{cc} \eta \mathbb{I}_n - xx^\top + \lambda_i M & (x^\top)v - \eta \xi_i \\ (x^\top)v^\top - \eta \xi_i^\top & \eta \|\xi_i\|_2^2 - (x^\top v)^2 - \lambda_i - \theta_i \end{array} \right] \succeq 0, \quad \forall i \leq N. \end{array} \right. \end{array} \right. \quad (68)$$

Moreover, for any solution $(\eta^*, \lambda^*, \theta^*)$, to the SDP (68), the pair (\mathbf{Q}_x, η_x) given by

$$\eta_x = \eta^*, \quad \text{and} \quad \mathbf{Q}_x(\xi) = \frac{1}{N} \sum_{i=1}^N \delta \left(\xi - (\eta^* \mathbb{I}_n - xx^\top + \lambda_i^* M)^{-1} (\eta^* \xi_i - xx^\top v) \right), \quad (69)$$

is a solution to (61) and its dual (62) respectively.

7.3. Simulation results

We validate the convergence attributes of Algorithm 1 for the NDRO problem of the minimum variance portfolio selection (56) with ellipsoidal support (i.e., $\Xi = \mathcal{E}_M \subset \mathbb{R}^n$). The positive definite matrix M that characterizes the support \mathcal{E}_M , is generated randomly and it is ensured to be reasonably well conditioned. The samples ξ_i , for $i = 1, 2, \dots, N = 2n$, are drawn randomly from \mathcal{E}_M , and the distribution $\widehat{\mathbb{P}}$ is uniformly distributed over the samples. Then for $n = 25$ and various values of ρ (the radius of the ambiguity set), we solve

$$V^* := \min_{x \in \Delta^n} \sup_{\mathbb{P} \in \mathcal{W}_2(\widehat{\mathbb{P}}, \rho)} V(x, \mathbb{P}). \quad (70)$$

The proposed FW-based method (Algorithm 1 with the simplified FW-update (64)) is implemented in MATLAB, wherein the SDP (68) corresponding to the FW-oracle and the minimization problem: $\min_{x \in \Delta^n} V(x, \mathbb{P}_k)$ are solved using the `cvx` solver.

First, Figure 6 shows the convergence of Algorithm 1 for $K(\varepsilon) = 75$ iterations, and for three distinct values of $\rho = 0.1, 0.5$, and 1. The plots on the top show the evolution of the primal and dual functions: $\min_{x \in \Delta^n} V(x, \mathbb{P}_k)$ and $\sup_{\mathbb{P} \in \mathcal{W}_2(\widehat{\mathbb{P}}, \rho)} V(x_k, \mathbb{P})$ respectively, w.r.t. the iteration k of the algorithm. Whereas, the plots at the bottom of the figure show the evolution of the two sub-optimality metrics:

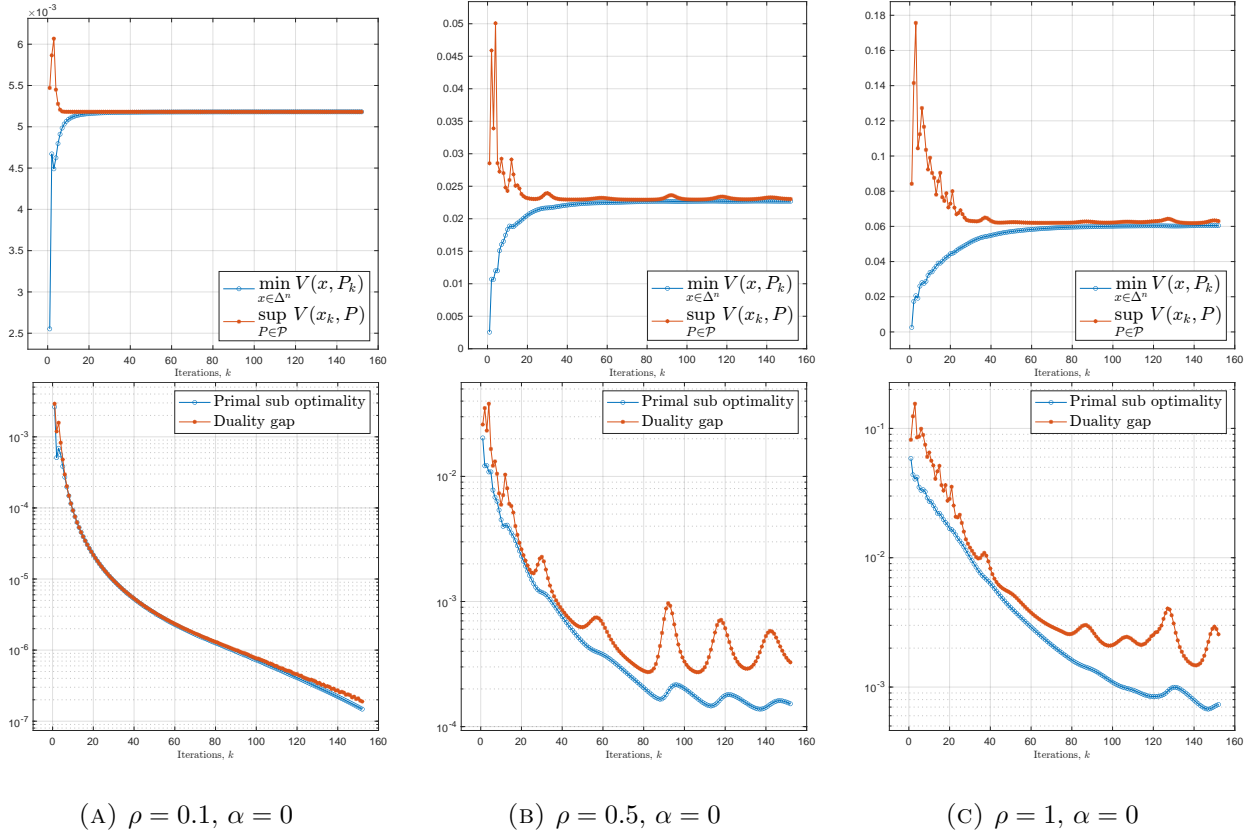
$$\left\{ \begin{array}{ll} \text{primal sub-optimality} & : \quad V^* - \min_{x \in \Delta^n} V(x, \mathbb{P}_k) \quad \text{and} \\ \text{duality gap} & : \quad \sup_{\mathbb{P} \in \mathcal{W}_2(\widehat{\mathbb{P}}, \rho)} V(x_k, \mathbb{P}) - \min_{x \in \Delta^n} V(x, \mathbb{P}_k). \end{array} \right.$$

Since Algorithm 1 explicitly solves the minimization over x in each iteration (see (42)), the primal function is readily available. However, this is not the case with the dual function which needs to be computed independently at each iteration for the current iterate x_k . We compute it by running the FW-update (64) for several iterations ($\sim K(\varepsilon)$) independently with x_k held fixed, this amounts to applying the FW-algorithm (25) for the risk measure $V_{x_k}(\cdot)$.

With explicit regularization of x . We observe in Figure 6 that, as the value of ρ increases, the convergence of the algorithm becomes less smooth, and also slower. Perhaps, this can be attributed to the fact that the strong-convexity assumption in Assumption 5.3 fails. This is so because, a larger value of ρ allows the samples $(\xi_i)_i$ to be perturbed in such a way that the variance matrix of the perturbed points is rank deficient. Thus, the function $V(x, \mathbb{P})$ is not strongly convex in x for \mathbb{P} corresponding to the perturbed points.

To remedy this, as suggested in Section 5.3, we compute an approximate saddle point of (70) by applying the FW-algorithm to the explicitly regularized min-sup problem

$$\min_{x \in \Delta^n} \sup_{\mathbb{P} \in \mathcal{W}_2(\widehat{\mathbb{P}}, \rho)} \frac{\alpha}{2} \|x\|_2^2 + V(x, \mathbb{P}). \quad (71)$$

FIGURE 6. Convergence plots for Algorithm 1 with $K(\varepsilon) = 75$, applied to (70).

Since $\|x\|_2 \leq \|x\|_1 = 1$ for all $x \in \Delta^n$, we know that an ε -saddle point of (70) can be computed for any $\varepsilon > 0$ by applying Algorithm 1 to (71) with $\alpha = 2\varepsilon$. We emphasize that both the primal function: $\min_{x \in \Delta^n} V(x, P_k)$ and the dual function: $\sup_{P \in \mathcal{W}_2(\hat{P}, \rho)} V(x_k, P)$ are not accessible in the implementation of Algorithm 1 for the explicitly regularized objective function. Thus, these quantities are computed separately at each iteration of the algorithm to record the primal sub-optimality and the duality gap.

In Figure 7, we show the convergence plots of Algorithm 1 when applied to the explicitly regularized problem (71). We consider the same data set from Figure 6 but with a slightly larger value of $\rho = 1.5$, making the conditioning of the problem even worse than that for $\rho = 1$.

We show the convergence plots of the algorithm for three different values of $\alpha = 0, 0.1, 1$. Clearly, for $\alpha = 0$, the problem is not regularized, and the convergence is bad (even worse than that for $\rho = 1$ from Figure 6c). Then the effect of explicit regularization and how it improves the regularity and convergence can be clearly seen in Figure 7. For $\alpha = 0.1$ and 1, we see that the solution computed is an ε -saddle point for the respective value of ε (i.e., $\varepsilon = 0.05$ for $\alpha = 0.1$ and $\varepsilon = 0.5$ for $\alpha = 1$). In addition, Figures 7b and 7c highlight the trade-off between the speed of convergence and worst-case performance metrics. For the larger value of $\alpha = 1$, due to better regularity, the algorithm is faster which can be seen in Figure 7c that both the curves of primal (blue) and dual (red) functions achieve within one percent of their final value in less than 20 iterations. Whereas, with $\alpha = 0.1$, even after 60 iterations, the primal function (blue curve) is only within 3.5 percent of its final value. However,

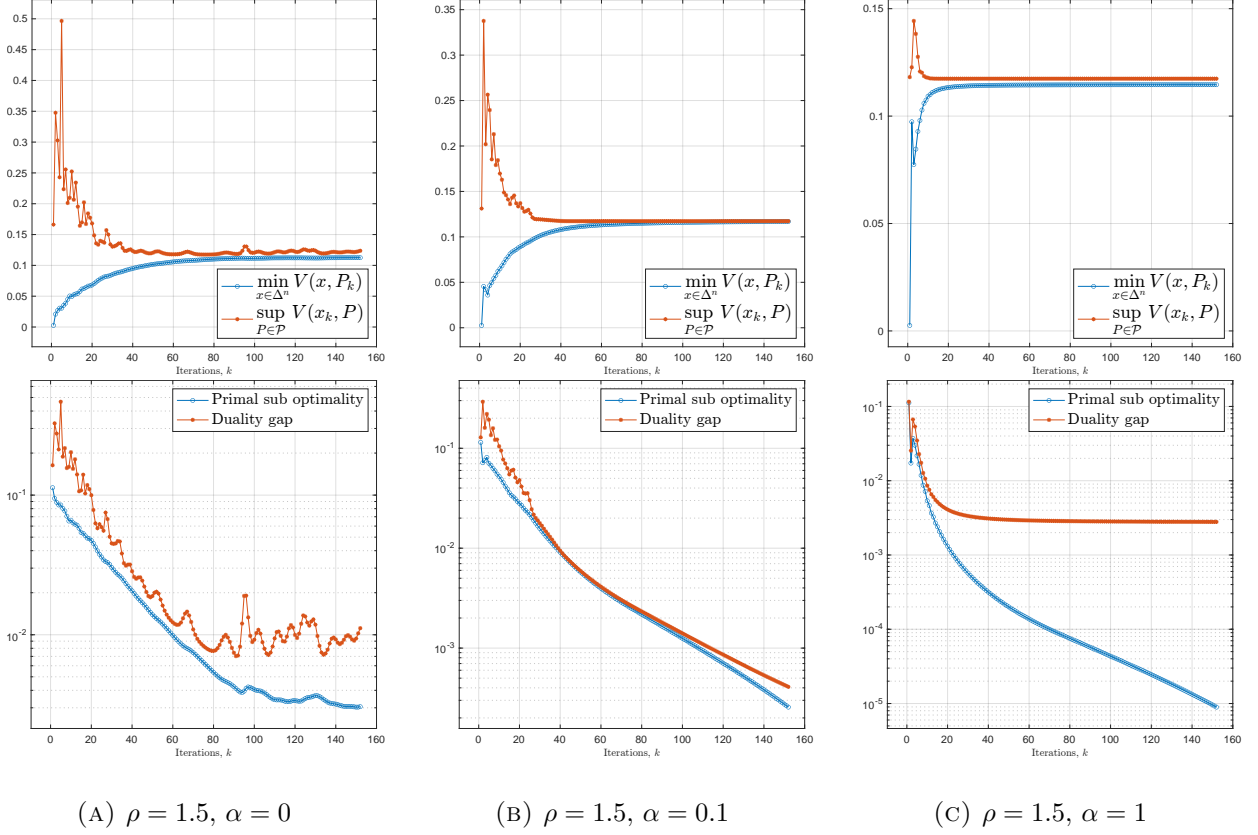


FIGURE 7. Convergence plots for Algorithm 1 with $K(\varepsilon) = 75$, applied to the explicitly regularized approximation (71).

this improved speed comes at a cost as clearly evident from Figures 7b and 7c that the worst-case cost for $\alpha = 0.1$ (0.11731) is smaller than that for $\alpha = 1$ (0.117403). Furthermore, even though the solution computed with both $\alpha = 0, 1$ and 1, is within the respective sub-optimality levels, the duality gap however, actually goes to zero for $\alpha = 0.1$, which is not the case for $\alpha = 1$.

8. TECHNICAL PROOFS

8.1. Proofs of Section 3 (derivatives)

In this part, we cover the technical proofs of the theoretical statements in Section 3.

Proof of Lemma 3.2. From the definition 3.1, we see that

$$\begin{aligned}
 dR(\mathbb{P}; \mathbb{Q}) &= \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} (R(\mathbb{P} + \varepsilon(\mathbb{Q} - \mathbb{P})) - R(\mathbb{P})) \\
 &= \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} (r(\mathbb{E}_{\mathbb{P}}[L(\xi)] + \varepsilon(\mathbb{E}_{\mathbb{Q}-\mathbb{P}}[L(\xi)])) - r(\mathbb{E}_{\mathbb{P}}[\xi])) \\
 &= \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} (r(\mathbb{E}_{\mathbb{P}}[\xi]) + \varepsilon \langle \nabla r(\mathbb{E}_{\mathbb{P}}[L(\xi)]), \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[L(\xi)] \rangle + o(\varepsilon^2) - r(\mathbb{E}_{\mathbb{P}}[\xi])) \\
 &= \langle \nabla r(\mathbb{E}_{\mathbb{P}}[L(\xi)]), \mathbb{E}_{\mathbb{Q}-\mathbb{P}}[L(\xi)] \rangle
 \end{aligned}$$

$$= \mathbb{E}_{\mathbf{Q}-\mathbf{P}} [\langle \nabla r(\mathbb{E}_{\mathbf{P}}[L(\xi)]), L(\xi) \rangle].$$

□

Proof of Proposition 3.4. From Definition 3.1, we have

$$\begin{aligned} dR(\mathbf{P}; \mathbf{P}_\gamma) &= \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} \left(R(\mathbf{P} + \varepsilon(\mathbf{P}_\gamma - \mathbf{P})) - R(\mathbf{P}) \right) = \lim_{\varepsilon \downarrow 0} \frac{\gamma}{(\varepsilon\gamma)} \left(R(\mathbf{P} + \varepsilon\gamma(\mathbf{Q} - \mathbf{P})) - R(\mathbf{P}) \right) \\ &= \gamma \lim_{\varepsilon' \downarrow 0} \frac{1}{\varepsilon'} \left(R(\mathbf{P} + \varepsilon'(\mathbf{Q} - \mathbf{P})) - R(\mathbf{P}) \right) \\ &= \gamma dR(\mathbf{P}; \mathbf{Q}), \end{aligned}$$

which establishes (14a). To establish (14b), we see that

$$\begin{aligned} dR(\mathbf{P}; \mathbf{Q}) &= \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} \left(R(\mathbf{P} + \varepsilon(\mathbf{Q} - \mathbf{P})) - R(\mathbf{P}) \right) \geq \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} \left(R(\mathbf{P}) + \varepsilon(R(\mathbf{Q}) - R(\mathbf{P})) - R(\mathbf{P}) \right) \\ &\geq R(\mathbf{Q}) - R(\mathbf{P}), \end{aligned}$$

where the first inequality is due to concavity of $\mathbf{P} \mapsto R(\mathbf{P})$. This completes the proof. □

Proof of Lemma 3.7. For any $\mathbf{P}, \mathbf{Q} \in \mathcal{P}$ with $\mathbf{P}_\gamma = \mathbf{P} + \gamma(\mathbf{Q} - \mathbf{P})$ and $\gamma \in [0, 1]$ we see that

$$\begin{aligned} dR(\mathbf{P}_\gamma; \mathbf{P}) + dR(\mathbf{P}; \mathbf{P}_\gamma) &= \langle \nabla r(\mathbb{E}_{\mathbf{P}}[L(\xi)]), \mathbb{E}_{\mathbf{P}_\gamma}[L(\xi)] - \mathbb{E}_{\mathbf{P}}[L(\xi)] \rangle + \langle \nabla r(\mathbb{E}_{\mathbf{P}_\gamma}[L(\xi)]), \mathbb{E}_{\mathbf{P}}[L(\xi)] - \mathbb{E}_{\mathbf{P}_\gamma}[L(\xi)] \rangle \\ &= \langle \nabla r(\mathbb{E}_{\mathbf{P}}[L(\xi)]) - \nabla r(\mathbb{E}_{\mathbf{P}_\gamma}[L(\xi)]), \mathbb{E}_{\mathbf{P}_\gamma}[L(\xi)] - \mathbb{E}_{\mathbf{P}}[L(\xi)] \rangle \\ &\leq \beta \|\mathbb{E}_{\mathbf{P}_\gamma}[L(\xi)] - \mathbb{E}_{\mathbf{P}}[L(\xi)]\|^2 \leq \gamma^2 \beta \|\mathbb{E}_{\mathbf{Q}}[L(\xi)] - \mathbb{E}_{\mathbf{P}}[L(\xi)]\|^2 \leq \gamma^2 (\beta d^2). \end{aligned}$$

□

8.2. Proofs of Section 4 (FW-algorithm)

In this part, we cover the technical proofs of the theoretical statements in Section 4.

Proof for Lemma 4.3. For \mathbf{P} and \mathbf{P}_γ , we see that the inequalities hold.

$$\begin{aligned} R(\mathbf{P}) - R(\mathbf{P}_\gamma) &\leq dR(\mathbf{P}_\gamma; \mathbf{P}) && \text{from (14b)} \\ &= dR(\mathbf{P}_\gamma; \mathbf{P}) + dR(\mathbf{P}; \mathbf{P}_\gamma) - dR(\mathbf{P}; \mathbf{P}_\gamma) \\ &\leq \gamma^2 C - \gamma dR(\mathbf{P}; \mathbf{Q}) && \text{from smoothness (15), and (14a)} \\ &\leq \gamma^2 C(1 + \delta) - \gamma dR(\mathbf{P}; \mathbf{Q}) && \text{from (22)}. \end{aligned}$$

Rearranging terms and using (14b), we get

$$R(\mathbf{Q}) - R(\mathbf{P}) \leq dR(\mathbf{P}; \mathbf{Q}) \leq \frac{1}{\gamma} \left(R(\mathbf{P}_\gamma) - R(\mathbf{P}) \right) + \gamma C(1 + \delta),$$

wherein, by considering the supremum over $\mathbf{Q} \in \mathcal{P}$, the inequalities in (23) follow at once. □

Proof for Proposition 4.4. Let $(\mathbf{P}_k)_k$ be the sequence of Frank-Wolfe iterates updated as in (25). From the one-step inequality (24), we have

$$R^* - R(\mathbf{P}_{k+1}) \leq \gamma_k^2 C(1 + \delta) + (1 - \gamma_k)(R^* - R(\mathbf{P}_k)) \quad \text{for every } k \geq 0. \quad (72)$$

Using (72), we prove the primal convergence of FW-algorithm following the classical approach [19, 9] by the method of induction.

Base case: For $k = 1$, since $\mathbb{P}_1 = \mathbb{P}_0 + \gamma_0(\mathbb{Q}_0 - \mathbb{P}_0)$ we deduce from (72) using $\gamma = \gamma_0 = 1$ that

$$R^* - R(\mathbb{P}_1) \leq C(1 + \delta) < \frac{4C}{3}(1 + \delta).$$

Thus, (26) holds for $k = 1$.

Induction step: Now, assuming that (26) holds for some $k \geq 1$, we deduce from (72) that

$$\begin{aligned} R^* - R(\mathbb{P}_{k+1}) &\leq \gamma_k^2 C(1 + \delta) + (1 - \gamma_k)(R^* - R(\mathbb{P}_k)) \\ &\leq \frac{4C(1 + \delta)}{(k + 2)^2} + \frac{k}{k + 2} \frac{4C(1 + \delta)}{k + 2} = \frac{4C(1 + \delta)}{k + 2} \frac{1 + k}{k + 2} \\ &\leq \frac{4C(1 + \delta)}{k + 3}, \quad \text{since } (k + 2)^2 \leq (k + 1)(k + 3). \end{aligned}$$

Thus, we conclude that (26) indeed holds for every $k \geq 1$. The proof is now complete. \square

Proof for Proposition 4.6. Following the ideas of [9, 19], we first prove by the method of contradiction that there exists \hat{k} such that

$$\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbb{Q}) \leq \frac{4C(1 + \delta)}{K + 2} \quad \text{for some } \hat{k} \in \{K, \dots, 2K + 1\}. \quad (73)$$

Suppose that $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_k; \mathbb{Q}) > \frac{4C(1 + \delta)}{K + 2}$ for every $k \in \{K, \dots, 2K + 1\}$, then by considering $\gamma = \frac{2}{K + 2}$ in (23) and rearranging terms, we see that for every $k \in \{K, \dots, 2K + 1\}$, we have the inequality

$$\begin{aligned} R(\mathbb{P}_k) - R(\mathbb{P}_{k+1}) &\leq \frac{4}{(K + 2)^2} C(1 + \delta) - \frac{2}{K + 2} \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_k; \mathbb{Q}) \\ &< \frac{4}{(K + 2)^2} C(1 + \delta) - \frac{8}{(K + 2)^2} C(1 + \delta) \\ &< -\frac{4}{K + 2} C(1 + \delta). \end{aligned}$$

Summing these inequalities for $k = K, \dots, 2K + 1$, we get $R(\mathbb{P}_K) - R(\mathbb{P}_{2K+2}) < -(4/K+2)C(1+\delta)$. Combining this with the sub-optimality of $R(\mathbb{P}_K)$ from (26) finally gives

$$R^* - R(\mathbb{P}_{2K+2}) < -\frac{4}{K + 2} C(1 + \delta) + \frac{4}{K + 2} C(1 + \delta) < 0,$$

which is a contradiction. Therefore, there must exist some $\hat{k} \in \{K, \dots, 2K + 1\}$, such that (73) holds.

For every \hat{k} such that (73) holds, we also necessarily have

$$g_{\hat{k}} = dR(\mathbb{P}_{\hat{k}}; \mathbb{Q}_{\hat{k}}) \leq \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbb{Q}) \leq \frac{4C(1 + \delta)}{K + 2}. \quad (74)$$

We emphasize that at each iteration k , the FW-oracle solves the FW-problems upto an additive accuracy of $\delta\gamma_{\hat{k}}C$. Therefore, we do not have access to the exact value of $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_k; \mathbb{Q})$, but only its approximation g_k . Consequently, it is only possible to verify whether the upper bound (74) for $g_{\hat{k}}$ is satisfied for some \hat{k} , and impossible to verify whether (73) holds. Moreover, even if (74) holds for some \hat{k} , it is not necessary that (73) also holds since $g_{\hat{k}} \leq \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbb{Q})$. However,

since $g_{\hat{k}}$ is approximately equal to the FW-gap, an upper bound on $g_{\hat{k}}$ gives the following slightly worse upper bound on the FW-gap

$$\sup_{\mathbf{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbf{Q}) \leq \delta \gamma_{\hat{k}} C + dR(\mathbb{P}_{\hat{k}}; \mathbf{Q}_{\hat{k}}) \leq \frac{2\delta C}{K+2} + \frac{4C(1+\delta)}{K+2} \leq \frac{2C(2C+3\delta)}{K+2}.$$

Finally, since $R(\mathbf{Q}) - R(\mathbb{P}) \leq dR(\mathbb{P}; \mathbf{Q})$ for every $\mathbf{Q} \in \mathcal{P}$ (follows from (14b)), taking supremum over $\mathbf{Q} \in \mathcal{P}$ immediately gives the inequality

$$R^* - R(\mathbb{P}) \leq \sup_{\mathbf{Q} \in \mathcal{P}} dR(\mathbb{P}_{\hat{k}}; \mathbf{Q}).$$

This completes the proof of the proposition. \square

8.3. Proofs of Section 5 (NDRO)

In this part we cover the technical proofs of the theoretical statements in Section 5.

Proof of Lemma 5.5. Since the set \mathcal{X} is closed and the function $F(x, \mathbb{P})$ is continuous and α -strongly convex in x for every $\mathbb{P} \in \mathcal{P}$, we conclude that the minimizer $\operatorname{argmin}_{x \in \mathcal{X}} F(x, \mathbb{P})$ exists and unique. Consequently, $x(\mathbb{P}) := \operatorname{argmin}_{x \in \mathcal{X}} F(x, \mathbb{P})$ is a singleton for every $\mathbb{P} \in \mathcal{P}$.

Proof of Lemma 5.5(i): Danskin's theorem. Letting $\mathbb{P}_\gamma := \mathbb{P} + \gamma(\mathbf{Q} - \mathbb{P})$ for $\gamma \in [0, 1]$, consider the mappings

$$\begin{cases} \mathcal{X} \times [0, 1] \ni (x, \gamma) \mapsto f(x, \gamma) := F(x, \mathbb{P}_\gamma) & \text{and} \\ [0, 1] \ni \gamma \mapsto g(\gamma) := \min_{x \in \mathcal{X}} f(x, \gamma). \end{cases}$$

It is easily seen that $g(\gamma) = R(\mathbb{P}_\gamma)$. Moreover, for the one sided derivatives of g at $\gamma \in [0, 1]$ defined: $dg(\gamma; 1) := \lim_{\delta \downarrow 0} \frac{1}{\delta}(g(\gamma + \delta) - g(\gamma))$, we easily verify that $dR(\mathbb{P}; \mathbf{Q}) = dg(0; 1)$. Using the short-hand notation $f_x(\cdot) := f(x, \cdot)$, we also verify similarly that $df_x(0; 1) = dF_x(\mathbb{P}; \mathbf{Q})$, for every $x \in \mathcal{X}$. For every $\gamma \in [0, 1]$, let $x_\gamma := \operatorname{argmin}_{x \in \mathcal{X}} f(x, \gamma)$, we know from the Danskin's theorem [2, (A.22), p. 154] that the one sided derivatives $dg(\gamma; 1)$ at γ , are given by

$$dg(\gamma; 1) = df_{x_\gamma}(\gamma; 1) = \lim_{\delta \downarrow 0} \frac{1}{\delta}(f(x_\gamma, \gamma + \delta) - f(x_\gamma, \gamma)).$$

Collecting everything, we have $dR(\mathbb{P}; \mathbf{Q}) = dg(0; 1) = df_{x_0}(0; 1) = dF_{x(\mathbb{P})}(\mathbb{P}; \mathbf{Q})$.

Proof of Lemma 5.5(ii): Smoothness. Recall that $\nabla_1 F(x, \mathbb{P}) := (\partial F / \partial x)(x, \mathbb{P})$ denotes the partial derivative of F w.r.t. x evaluated at (x, \mathbb{P}) . Since $x(\mathbb{P}) = \operatorname{argmin}_{x \in \mathcal{X}} F(x, \mathbb{P})$, the first-order optimality conditions give

$$\langle \nabla_1 F(x(\mathbb{P}), \mathbb{P}), y - x(\mathbb{P}) \rangle \geq 0 \quad \text{for all } y \in \mathcal{X}. \quad (75)$$

Due to α -strong-convexity of $F(y, \mathbb{P})$ in y , we have

$$\begin{aligned} \frac{\alpha}{2} \|x(\mathbb{P}) - x(\mathbf{Q})\|^2 &\leq F(x(\mathbf{Q}), \mathbb{P}) - F(x(\mathbb{P}), \mathbb{P}) - \langle \nabla_1 F(x(\mathbb{P}), \mathbb{P}), x(\mathbf{Q}) - x(\mathbb{P}) \rangle \\ &\leq F(x(\mathbf{Q}), \mathbb{P}) - F(x(\mathbb{P}), \mathbb{P}) \quad \text{from (75) with } y = x(\mathbf{Q}). \end{aligned}$$

Similarly, α -strong-convexity of $F(\cdot, \mathbf{Q})$ gives us

$$\frac{\alpha}{2} \|x(\mathbf{Q}) - x(\mathbb{P})\|^2 \leq F(x(\mathbb{P}), \mathbf{Q}) - F(x(\mathbf{Q}), \mathbf{Q}).$$

Combining the two inequalities, we infer that the inequality

$$\begin{aligned}
\alpha \|x(\mathbb{P}) - x(\mathbb{Q})\|^2 &\leq F(x(\mathbb{Q}), \mathbb{P}) - F(x(\mathbb{P}), \mathbb{P}) + F(x(\mathbb{P}), \mathbb{Q}) - F(x(\mathbb{Q}), \mathbb{Q}) \\
&= F(x(\mathbb{Q}), \mathbb{P}) - F(x(\mathbb{Q}), \mathbb{Q}) + F(x(\mathbb{P}), \mathbb{Q}) - F(x(\mathbb{P}), \mathbb{P}) \\
&\leq dF_{x(\mathbb{Q})}(\mathbb{Q}; \mathbb{P}) + dF_{x(\mathbb{P})}(\mathbb{P}; \mathbb{Q}) && \text{from (14b)} \\
&= dR(\mathbb{Q}; \mathbb{P}) + dR(\mathbb{P}; \mathbb{Q}) && \text{from (41),}
\end{aligned} \tag{76}$$

holds for every $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$. On the one hand, for $\mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$, $\gamma \in [0, 1]$, we have

$$\begin{aligned}
dR(\mathbb{P}; \mathbb{P}_\gamma) + dR(\mathbb{P}_\gamma; \mathbb{P}) &= dF_{x(\mathbb{P})}(\mathbb{P}; \mathbb{P}_\gamma) + dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}_\gamma; \mathbb{P}) && \text{from (41)} \\
&= dF_{x(\mathbb{P})}(\mathbb{P}; \mathbb{P}_\gamma) - dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}; \mathbb{P}_\gamma) + dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}; \mathbb{P}_\gamma) + dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}_\gamma; \mathbb{P}) \\
&= \gamma \left(dF_{x(\mathbb{P})}(\mathbb{P}; \mathbb{Q}) - dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}; \mathbb{Q}) \right) + \left(dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}; \mathbb{P}_\gamma) + dF_{x(\mathbb{P}_\gamma)}(\mathbb{P}_\gamma; \mathbb{P}) \right) && \text{from (14a)} \\
&\leq \gamma C_1 \|x(\mathbb{P}_\gamma) - x(\mathbb{P})\| + \gamma^2 C_2,
\end{aligned} \tag{77}$$

where the last inequality is due to (39) and the smoothness condition (iii) of Assumption 5.3. On the other hand, considering $\mathbb{Q} = \mathbb{P}_\gamma$ in (76), we have

$$\alpha \|x(\mathbb{P}_\gamma) - x(\mathbb{P})\|^2 \leq dR(\mathbb{P}_\gamma; \mathbb{P}) + dR(\mathbb{P}; \mathbb{P}_\gamma). \tag{78}$$

Collecting (77) and (78) together, we see that

$$\alpha \|x(\mathbb{P}) - x(\mathbb{P}_\gamma)\|^2 \leq \gamma C_1 \|x(\mathbb{P}) - x(\mathbb{P}_\gamma)\| + \gamma^2 C_2.$$

On rearranging and simplifying terms, it is now easily verified that

$$\left(\|x(\mathbb{P}) - x(\mathbb{P}_\gamma)\| - \frac{\gamma}{2\alpha} \left(C_1 - \sqrt{C_1^2 + 4\alpha C_2} \right) \right) \left(\|x(\mathbb{P}) - x(\mathbb{P}_\gamma)\| - \frac{\gamma}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right) \right) \leq 0,$$

which is only true if

$$\frac{\gamma}{2\alpha} \left(C_1 - \sqrt{C_1^2 + 4\alpha C_2} \right) \leq \|x(\mathbb{P}) - x(\mathbb{P}_\gamma)\| \leq \frac{\gamma}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right).$$

The lower bound is irrelevant since it is negative. However, the upper bound is non-trivial, and employing it in (77) finally gives

$$\begin{aligned}
dR(\mathbb{P}; \mathbb{P}_\gamma) + dR(\mathbb{P}_\gamma; \mathbb{P}) &\leq \gamma^2 \frac{C_1}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right) + \gamma^2 C_2 \\
&\leq \gamma^2 \left(C_2 + \frac{C_1}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right) \right).
\end{aligned}$$

Since $C_2 + \frac{C_1}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right) < +\infty$, we conclude that the risk measure R is C -smooth in the sense of Definition 3.6 for every $C \geq C_2 + \frac{C_1}{2\alpha} \left(C_1 + \sqrt{C_1^2 + 4\alpha C_2} \right)$. The proof is now complete. \square

Proof for Theorem 5.6. We establish assertion (i) of the theorem assuming assertion (ii) holds which is independently proved later.

Proof for Theorem 5.6 (i): Strong duality. Assuming assertion (ii) holds, we see that

$$\min_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{P}} F(x, \mathbb{Q}) \leq \sup_{\mathbb{Q} \in \mathcal{P}} F(x_\varepsilon, \mathbb{Q}) \leq \varepsilon + \min_{x \in \mathcal{X}} F(x, \mathbb{P}_\varepsilon) \leq \varepsilon + \sup_{\mathbb{Q} \in \mathcal{P}} \min_{x \in \mathcal{X}} F(x, \mathbb{Q}),$$

holds for every $\varepsilon > 0$, and thus, we have $\min_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{P}} F(x, \mathbb{Q}) \leq \sup_{\mathbb{Q} \in \mathcal{P}} \min_{x \in \mathcal{X}} F(x, \mathbb{Q})$. This, together with weak duality: $\sup_{\mathbb{Q} \in \mathcal{P}} \min_{x \in \mathcal{X}} F(x, \mathbb{Q}) \leq \min_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathcal{P}} F(x, \mathbb{Q})$ proves assertion (i).

Proof for Theorem 5.6 (ii): Saddle point computation. Let us recall that

$$R(\mathbb{P}) = \min_{x \in \mathcal{X}} F(x, \mathbb{P}) \quad \text{and} \quad x(\mathbb{P}) = \operatorname{argmin}_{x \in \mathcal{X}} F(x, \mathbb{P}).$$

From (41), since $dR(\mathbb{P}; \mathbb{Q}) = dF_{x(\mathbb{P})}(\mathbb{P}; \mathbb{Q})$, the iterates $(\mathbb{P}_k)_k$ obtained from Algorithm 1 can be equivalently regarded as the ones obtained from the FW-algorithm (25) for the maximization problem: $\sup_{\mathbb{P} \in \mathcal{P}} R(\mathbb{P})$ under the setting of Proposition 4.6. Therefore, we conclude from Proposition 4.6, and more specifically from Remark 4.7, we know that there exists a $K(\varepsilon) \leq \widehat{k} \leq 2K(\varepsilon) + 1$ such that $dR(\mathbb{P}_{\widehat{k}}; \mathbb{Q}_{\widehat{k}}) \leq \varepsilon \frac{2+2\delta}{2+3\delta}$.

We also conclude from Remark 4.7 that $\sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_{\widehat{k}}; \mathbb{Q}) \leq \varepsilon$ for any $K(\varepsilon) \leq \widehat{k} \leq 2K(\varepsilon) + 1$ satisfying $dR(\mathbb{P}_{\widehat{k}}; \mathbb{Q}_{\widehat{k}}) \leq \varepsilon \frac{2+2\delta}{2+3\delta}$. In particular, for $(x_\varepsilon, \mathbb{P}_\varepsilon)$ to be the output of Algorithm 1, we know that

$$\sup_{\mathbb{Q} \in \mathcal{P}} dF_{x_\varepsilon}(\mathbb{P}_\varepsilon; \mathbb{Q}) = \sup_{\mathbb{Q} \in \mathcal{P}} dR(\mathbb{P}_\varepsilon; \mathbb{Q}) \leq \varepsilon.$$

Consequently, for any $\mathbb{Q} \in \mathcal{P}$, we have

$$F(x_\varepsilon, \mathbb{Q}) - F(x_\varepsilon, \mathbb{P}_\varepsilon) \leq dF_{x_\varepsilon}(\mathbb{P}_\varepsilon; \mathbb{Q}) \leq \varepsilon,$$

where the first inequality follows from (14b) for $F(x_\varepsilon, \cdot)$. Finally, taking the supremum over $\mathbb{Q} \in \mathcal{P}$ we conclude $\sup_{\mathbb{Q} \in \mathcal{P}} F(x_\varepsilon, \mathbb{Q}) \leq \varepsilon + F(x_\varepsilon, \mathbb{P}_\varepsilon)$; which together with the fact that $F(x_\varepsilon, \mathbb{P}_\varepsilon) = \min_{x \in \mathcal{X}} F(x, \mathbb{P}_\varepsilon)$ implies $(x_\varepsilon, \mathbb{P}_\varepsilon)$ being indeed an ε -saddle point in the sense of Definition 5.2. The proof is now complete. \square

8.4. Proof of Section 6 (entropic risk)

We first prove the results on the FW oracle (Lemma 6.4) for the entropic risk portfolio selection problem (47) and then prove the results on regularity conditions (Lemma 6.2 and 6.3).

Proofs for the FW-oracle. We shall first establish two key lemmas that will be later used to prove Lemma 6.4.

Lemma 8.1. *Let $c > \theta > 0$, $\eta \geq 0$, $x \in [0, 1]$, and $\xi \in \mathbb{R}$. Consider*

$$\sup_{q \in \mathbb{R}} L(\eta, q) := e^{-\theta x(\xi+q)} - \eta e^{c|q|}. \quad (79)$$

The following assertions hold:

- (i) *If $\eta = 0$, the maximization problem (79) is unbounded.*
- (ii) *If $\eta > 0$, the maximization problem admits a unique optimal solution q_η , given by*

$$q_\eta = \min \left\{ 0, \frac{\theta x \xi + \log(c\eta/\theta x)}{c - \theta x} \right\}. \quad (80)$$

Proof of Lemma 8.1. We emphasize that the maximization problem (79) is *non-convex*. Even then, we shall establish the conditions for (79) to admit an optimal solution and also explicitly characterize

it. To this end, if $\eta = 0$, we see that

$$\lim_{q \rightarrow -\infty} L(0, q) = e^{-\theta x \xi} \lim_{q \rightarrow -\infty} e^{-\theta x q} = +\infty.$$

Therefore, assertion (i) of the lemma follows at once.

Now, if $\eta > 0$, then we see that

$$\begin{cases} \lim_{q \rightarrow +\infty} L(\eta, q) = \lim_{q \rightarrow +\infty} e^{-\theta x(z+q)} - \eta e^{c q} = 0 - \eta(+\infty) & = -\infty, \\ \lim_{q \rightarrow -\infty} L(\eta, q) = \lim_{q \rightarrow -\infty} e^{-c q} \left(e^{-\theta x z} e^{(c-\theta x)q} - \eta \right) = (+\infty)(0 - \eta) & = -\infty. \end{cases}$$

Therefore, we conclude that whenever $\eta > 0$, (79) admits an optimal solution $q_\eta \in \mathbb{R}$. Moreover, at the optimal solution q_η , we know that the necessary optimality conditions must be satisfied

$$0 \in \frac{\partial}{\partial q} L(\eta, q_\eta) = -\theta x e^{-\theta x(z+q_\eta)} - \eta c e^{c|q_\eta|} \operatorname{sgn}(q_\eta). \quad (81)$$

From (81), it is immediately evident that $q_\eta \leq 0$. Moreover, we also see that

- If $(c\eta/\theta x) \geq e^{-\theta x z}$ - we see that

$$\frac{\partial}{\partial q} L(\eta, q) = \theta x e^{-\theta x q} \left((c\eta/\theta x) e^{-(c-\theta x)q} - e^{-\theta x z} \right) \geq 0 \text{ for all } q < 0, \text{ and}$$

$0 \in \frac{\partial}{\partial q} L(\eta, 0) = \eta c[-1, +1] - \theta x e^{-\theta x z}$. Therefore, $q_\eta = 0$ is the only point satisfying the necessary optimality condition (81), and consequently the unique optimal solution to (52).

- Similarly if $(c\eta/\theta x) < e^{-\theta x z}$ - it is easily verified that for $q_\eta = \frac{\theta x z + \log(\eta c/\theta x)}{c-\theta x}$, is the unique point which satisfies the necessary first-order optimality condition (81), and consequently, is the unique optimal solution to (52).

We finally note that the cases above can be compressed as $q_\eta = \min \left\{ 0, \frac{\theta x z + \log(\eta c/\theta x)}{c-\theta x} \right\}$. The proof is now complete. \square

Lemma 8.2. *Let c, θ , and x be as in Lemma 8.1, and let $Z(t)$, $t = 1, 2, \dots, T$ be a non-decreasing sequence of real numbers. Consider the following minimization problem*

$$\inf_{\eta \geq 0} J(\eta) := \eta e^{c\rho} + \frac{1}{T} \sum_{t=1}^T \sup_{q(t) \in \mathbb{R}} e^{-\theta x(Z(t)+q(t))} - \eta e^{c|q(t)|}. \quad (82)$$

The minimization problem (82) admits a unique optimal solution η^* given by

$$\eta^* := \frac{\theta x}{c} \left(\frac{1}{T e^{c\rho} - T'} \sum_{t=T'+1}^T e^{\frac{-c\theta x Z(t)}{c-\theta x}} \right)^{\frac{c-\theta x}{\theta x}}, \quad (83)$$

where $T' \in \{1, 2, \dots, T\}$ is the smallest integer such that $(T e^{c\rho} - T') e^{\frac{-c\theta x Z(T')}{c-\theta x}} \geq \sum_{t=T'+1}^T e^{\frac{-c\theta x Z(t)}{c-\theta x}}$.

Proof of Lemma 8.2. Since $\eta \mapsto \log(c\eta/\theta x)$ is monotonically increasing and eventually positive, we observe from (80) that $\lim_{\eta \rightarrow +\infty} q_\eta = 0$. Consequently, we see that

$$\lim_{\eta \rightarrow +\infty} J(\eta) = \lim_{\eta \rightarrow +\infty} \eta (e^{c\rho} - 1) + \frac{1}{T} \sum_{t=1}^T e^{-\theta x z(t)} = +\infty.$$

Similarly, as $\eta \downarrow 0$, we conclude from assertion (i) of Lemma 8.1 that $q_\eta < 0$, consequently, we see

$$\lim_{\eta \downarrow 0} J(\eta) = \lim_{\eta \downarrow 0} \eta e^{c\rho} + \eta^{\frac{-\theta x}{c-\theta x}} \left(\frac{c}{\theta x} - 1 \right) (\theta x/c)^{\frac{-c}{c-\theta x}} \frac{1}{T} \sum_{t=1}^T e^{\frac{-c\theta x Z(t)}{c-\theta x}} = +\infty.$$

Since the mapping $\eta \mapsto J(\eta)$ is convex, we know that there exists some $\eta^* > 0$ such that $\eta^* = \operatorname{argmin}_{\eta \geq 0} J(\eta)$, i.e., an optimal solution to (82) exists. To characterise an optimal solution η^* , we observe that q_η is unique for any $\eta > 0$, and from Danskin's theorem, we conclude that $\frac{d}{d\eta} J(\eta) = e^{c\rho} - \frac{1}{T} \sum_{t=1}^T e^{c|q_\eta(t)|}$. Therefore, the optimal solution η^* is such that $T e^{c\rho} = \sum_{t=1}^T e^{c|q_{\eta^*}(t)|}$. Finding such an η^* where the equality holds is not straightforward.

For $\eta(t) = (\theta x/c) e^{-\theta x Z(t)}$, $t = 1, 2, \dots, T$, it is easily verified that $q_{\eta_t}(s) = 0$ for all $s \leq t$. Thus,

$$\frac{d}{d\eta} J(\eta_t) = e^{c\rho} - \frac{1}{T} \left(t + \sum_{s=t+1}^T e^{\frac{c\theta x(Z(t)-Z(s))}{c-\theta x}} \right).$$

Let $T' \in \{1, 2, \dots, T\}$ be the smallest integer such that $\frac{d}{d\eta} J(\eta(t)) \geq 0$. Since $J(\eta)$ is convex, $\frac{d}{d\eta} J(\eta)$ is non-decreasing. Consequently, with $\eta_0 = 0$, we know that $\eta^* \in]\eta_{T'-1}, \eta_{T'}]$, and therefore, $q_{\eta^*}(t) = 0$ for all $t \leq T'$, and $q_{\eta^*}(t) = \frac{\theta x Z(t) + \log(c\eta^*/\theta x)}{c-\theta x}$ for $t = T'+1, \dots, T$. Substituting, these values of $q_{\eta^*}(t)$ in the equation $T e^{c\rho} = \sum_{t=1}^T e^{c|q_{\eta^*}(t)|}$ and simplifying for η^* gives (83). The proof of the lemma is complete. \square

Proof of Lemma 6.4. The dual problem of (52) is given by

$$\inf_{\eta \geq 0} J(\eta) = \eta e^{c\rho} + \frac{1}{T} \sum_{t=1}^T \sup_{q(t) \in \mathbb{R}} e^{-\theta x(z(t)+q(t))} - \eta e^{c|q(t)|}, \quad (84)$$

Proof of Lemma 6.4 (i): Optimal solution. Since the maximization over $q(t)$ is separable over t we see that

$$J(\eta) = \eta e^{c\rho} + \frac{1}{T} \sum_{s=1}^T \sup_{q'(s) \in \mathbb{R}} e^{-\theta x(Z(s)+q'(s))} - \eta e^{c|q'(s)|},$$

where the sequence $(Z(s))_s$ is non-increasing. Consequently, we conclude from Lemma 8.2 that (84) admits a unique optimal solution η^* given by (83). Given η^* , the optimal solution $q_{\eta^*}(t)$ for each $t = 1, 2, \dots, T$ is obtained from (80). Substituting for $q_{\eta^*}(t)$ from (80) and simplifying, we easily verify that $z^*(t) = z(t) + q_{\eta^*}(t)$. Consequently, $\mathbf{Q}^* = \frac{1}{T} \sum_{t=1}^T \delta(z^*(t))$ is the optimal solution to (52). This proves assertion (i) of the lemma.

Proof of Lemma 6.4 (ii): Lower and upper bounds. To prove assertion (ii) of the lemma, we first see that since $\eta^* > 0$, the first-order optimality conditions imply $0 = \frac{d}{d\eta} J(\eta^*) = e^{c\rho} - \frac{1}{T} \sum_{t=1}^T e^{c|q_{\eta^*}(t)|}$. Therefore, for any $s = 1, 2, \dots, T$, we have

$$e^{c|q_{\eta^*}(s)|} \leq \sum_{t=1}^T e^{c|q_{\eta^*}(t)|} = T e^{c\rho}.$$

Taking $\log_e(\cdot)$ on both sides of the inequality, we obtain $|q_{\eta^*}(s)| \leq \rho + \frac{\log(T)}{c}$. Since $q_{\eta^*}(s) \leq 0$, we have $-\rho - \frac{\log(T)}{c} \leq q_{\eta^*}(s) \leq 0$ for all $s = 1, \dots, T$. Combining this with $z^*(s) = z(s) + q_{\eta^*}(s)$ and $\underline{z} \leq z(s) \leq \bar{z}$, assertion (ii) follows. \square

Proofs for regularity of the entropic risk.

Proof of Lemma 6.2. For any $x \in \Delta^n$, we shall first show that

$$\operatorname{argmax}_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x, \mathbb{P}) = \operatorname{argmax}_{\mathbb{P}' \in \mathcal{P}'_c} \mathcal{E}(x, \mathbb{P}'). \quad (85)$$

To show $\operatorname{argmax}_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x, \mathbb{P}) \subset \operatorname{argmax}_{\mathbb{P}' \in \mathcal{P}'_c} \mathcal{E}(x, \mathbb{P}')$, let $\mathbb{P}_x \in \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x, \mathbb{P})$, due to the first-order optimality conditions together with (54), we see that

$$\mathbb{P}_x = \operatorname{argmax}_{\mathbb{Q} \in \mathcal{P}_c} d\mathcal{E}_x(\mathbb{P}_x; \mathbb{Q}) = \operatorname{argmax}_{\mathbb{Q}' \in \mathcal{P}'_c} d\mathcal{E}_x(\mathbb{P}_x; \mathbb{Q}'), \text{ and thus, } \mathbb{P}_x \in \operatorname{argmax}_{\mathbb{P}' \in \mathcal{P}'_c} \mathcal{E}(x, \mathbb{P}').$$

Similarly, to show that $\operatorname{argmax}_{\mathbb{P}' \in \mathcal{P}'_c} \mathcal{E}(x, \mathbb{P}') \subset \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x, \mathbb{P})$, let $\mathbb{P}'_x \in \operatorname{argmax}_{\mathbb{P}' \in \mathcal{P}'_c} \mathcal{E}(x, \mathbb{P}')$, then from again the first-order optimality conditions together with (54), we have

$$\mathbb{P}'_x = \operatorname{argmax}_{\mathbb{Q}' \in \mathcal{P}'_c} d\mathcal{E}_x(\mathbb{P}'_x; \mathbb{Q}') = \operatorname{argmax}_{\mathbb{Q} \in \mathcal{P}_c} d\mathcal{E}_x(\mathbb{P}'_x; \mathbb{Q}), \text{ and thus, } \mathbb{P}'_x \in \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x, \mathbb{P}).$$

Now, for (x^*, \mathbb{P}^*) to be a saddle-point of (47), we have

$$x^* \in \operatorname{argmin}_{x \in \mathcal{X}} \mathcal{E}(x, \mathbb{P}^*) \quad \text{and} \quad \mathbb{P}^* \in \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}_c} \mathcal{E}(x^*, \mathbb{P}) = \operatorname{argmax}_{\mathbb{P}' \in \mathcal{P}'_c} \mathcal{E}(x^*, \mathbb{P}'),$$

where the last equality is due to (85). Thus, (x^*, \mathbb{P}^*) is also a saddle point of (49), and vice versa. The proof is now complete. \square

Proof of Lemma 6.3. For every $j = 1, 2, \dots, n$, since $\underline{\xi}_j \leq \xi_j \leq \bar{\xi}_j$, \mathbb{P}_j -almost surely for all $\mathbb{P}_j \in \mathcal{W}_c((\hat{\mathbb{P}}_j, \rho))$, we have

$$e^{-\theta_j x_j \bar{\xi}_j} \leq \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}] \leq e^{-\theta_j x_j \underline{\xi}_j} \quad \text{for all } \mathbb{P}_j \in \mathcal{W}_c((\hat{\mathbb{P}}_j, \rho)). \quad (86)$$

Proof of Lemma 6.3(i): Continuous derivatives. Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}'_c$, and $x, y \in \Delta^n$, we have

$$\begin{aligned} d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}) - d\mathcal{E}_y(\mathbb{P}; \mathbb{Q}) &= \sum_{j=1}^n \frac{1}{\theta_j} \left(\frac{\mathbb{E}_{\mathbb{Q}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi'_j}]} - \frac{\mathbb{E}_{\mathbb{Q}_j}[e^{-\theta_j y_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j y_j \xi'_j}]} \right) \\ &= \sum_{j=1}^n \frac{1}{\theta_j} \left(\frac{\mathbb{E}_{\mathbb{Q}_j}[e^{-\theta_j x_j \xi_j}] \cdot \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j y_j \xi'_j}] - \mathbb{E}_{\mathbb{Q}_j}[e^{-\theta_j y_j \xi_j}] \cdot \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi'_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi'_j}] \cdot \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j y_j \xi'_j}]} \right) \\ &= \sum_{j=1}^n \frac{1}{\theta_j} \left(\frac{\mathbb{E}_{\mathbb{Q}_j \times \mathbb{P}_j} \left[e^{-\theta_j (x_j \xi_j + y_j \xi'_j)} - e^{-\theta_j (y_j \xi_j + x_j \xi'_j)} \right]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi'_j}] \cdot \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j y_j \xi'_j}]} \right) \end{aligned}$$

Due to convexity of $(\cdot) \mapsto e^{(\cdot)}$, for any $a, b \in \mathbb{R}$, we have $e^b - e^a \leq |b - a| e^{\max\{a, b\}}$, which gives

$$\begin{aligned} &e^{-\theta_j (x_j \xi_j + y_j \xi'_j)} - e^{-\theta_j (y_j \xi_j + x_j \xi'_j)} \\ &\leq \theta_j |x_j \xi_j + y_j \xi'_j - y_j \xi_j - x_j \xi'_j| e^{-\theta_j \min\{x_j \xi_j + y_j \xi'_j, x_j \xi'_j + y_j \xi_j\}} \\ &\leq \theta_j |(x_j - y_j)(\xi_j - \xi'_j)| e^{-\theta_j (2\underline{\xi}_j)} \quad \text{since } x_j, y_j \in [0, 1] \text{ and } \xi_j, \xi'_j \in [\underline{\xi}_j, \bar{\xi}_j] \\ &\leq \theta_j |x_j - y_j| (\bar{\xi}_j - \underline{\xi}_j) e^{-\theta_j (2\underline{\xi}_j)}. \end{aligned}$$

This together with (86) finally gives

$$\begin{aligned} d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}) - d\mathcal{E}_y(\mathbb{P}; \mathbb{Q}) &\leq \sum_{j=1}^n \frac{1}{\theta_j} \frac{\theta_j |x_j - y_j| (\bar{\xi}_j - \underline{\xi}_j) e^{-\theta_j(2\underline{\xi}_j)}}{e^{-2\theta_j \bar{\xi}_j}} = \sum_{j=1}^n |x_j - y_j| (\bar{\xi}_j - \underline{\xi}_j) e^{2\theta_j(\bar{\xi}_j - \underline{\xi}_j)} \\ &\leq \|x - y\|_2 \sqrt{\sum_{j=1}^n (\bar{\xi}_j - \underline{\xi}_j)^2 e^{4\theta_j(\bar{\xi}_j - \underline{\xi}_j)}}, \end{aligned}$$

where the last inequality is due to Cauchy-Schwartz. This establishes assertion (i) of the lemma.

Proof of Lemma 6.3(ii): Smoothness. For every $x \in \Delta^n$, and $\mathbb{P}, \mathbb{Q} \in \mathcal{P}_c$, we recall that

$$d\mathcal{E}_x(\mathbb{P}; \mathbb{Q}) = \sum_{j=1}^n \frac{1}{\theta_j} \frac{\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}.$$

Then for any $\gamma \in [0, 1]$ and $\mathbb{P}_\gamma = \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$, we have from (14a) that $d\mathcal{E}_x(\mathbb{P}; \mathbb{P}_\gamma) = \gamma d\mathcal{E}_x(\mathbb{P}; \mathbb{Q})$.

Moreover, using the relations $\mathbb{E}_{\mathbb{P} - \mathbb{P}_\gamma}[\cdot] = -\gamma \mathbb{E}_{\mathbb{Q} - \mathbb{P}}[\cdot]$, we also verify that

$$d\mathcal{E}_x(\mathbb{P}_\gamma; \mathbb{P}) = \sum_{j=1}^n \frac{1}{\theta_j} \frac{-\gamma \mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j)}[e^{-\theta_j x_j \xi_j}]}.$$

Thus, we obtain

$$\begin{aligned} d\mathcal{E}_x(\mathbb{P}; \mathbb{P}_\gamma) + d\mathcal{E}_x(\mathbb{P}_\gamma; \mathbb{P}) &= \gamma \sum_{j=1}^n \frac{1}{\theta_j} \left(\frac{\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]} - \frac{\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j)}[e^{-\theta_j x_j \xi_j}]} \right) \\ &= \gamma \sum_{j=1}^n \frac{\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\theta_j} \left(\frac{\mathbb{E}_{\mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j)}[e^{-\theta_j x_j \xi_j}] - \mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}]}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}] \cdot \mathbb{E}_{\mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j)}[e^{-\theta_j x_j \xi_j}]} \right) \\ &= \gamma^2 \sum_{j=1}^n \frac{1}{\theta_j} \frac{(\mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}])^2}{\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}] \cdot \mathbb{E}_{\mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j)}[e^{-\theta_j x_j \xi_j}]} \end{aligned}$$

Since $\mathbb{P}_j, \mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j) \in \mathcal{W}_c((\hat{\mathbb{P}}_j, \rho))$, employing (86) gives

$$\begin{cases} \left(\mathbb{E}_{\mathbb{P}_j}[e^{-\theta_j x_j \xi_j}] \right)^{-1} \leq e^{\theta_j x_j \bar{\xi}_j} \\ \left(\mathbb{E}_{\mathbb{P}_j + \gamma(\mathbb{Q}_j - \mathbb{P}_j)}[e^{-\theta_j x_j \xi_j}] \right)^{-1} \leq e^{\theta_j x_j \bar{\xi}_j}, \quad \text{and} \\ \mathbb{E}_{\mathbb{Q}_j - \mathbb{P}_j}[e^{-\theta_j x_j \xi_j}] \leq e^{-\theta_j x_j \underline{\xi}_j} - e^{-\theta_j x_j \bar{\xi}_j}. \end{cases}$$

Putting things together, we finally obtain

$$d\mathcal{E}_x(\mathbb{P}; \mathbb{P}_\gamma) + d\mathcal{E}_x(\mathbb{P}_\gamma; \mathbb{P}) \leq \gamma^2 \sum_{j=1}^n \frac{1}{\theta_j} e^{2\theta_j x_j \bar{\xi}_j} (e^{-\theta_j x_j \underline{\xi}_j} - e^{-\theta_j x_j \bar{\xi}_j})^2 = \gamma^2 \sum_{j=1}^n \frac{1}{\theta_j} \left(e^{\theta_j(\bar{\xi}_j - \underline{\xi}_j)} - 1 \right)^2.$$

Thus, assertion (ii) of the lemma follows and the proof is now complete. \square

8.5. Proofs of Section 7 (variance risk)

In this part, we cover the technical proofs of the theoretical statements in Section 7. The first proof is concerned with the regularity of the min-variance portfolio selection problem.

Proofs for regularity conditions (Variance).

Proof of Lemma 7.2. For any $x \in \mathcal{X}$, since V_x is an RR measure, we simplify (57) to get

$$dV_x(\mathbb{P}; \mathbb{Q}) = x^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})x - 2(x^\top \mu_{\mathbb{P}})(x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}})) \quad (87a)$$

$$= x^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}}) \cdot x + (x^\top \mu_{\mathbb{P}})^2 + (x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 - (x^\top \mu_{\mathbb{Q}})^2. \quad (87b)$$

Proof of Lemma 7.2(i): Continuous derivatives. Firstly, we begin by showing that the inequality

$$x^\top Mx - y^\top My \leq 2 \|M\|_o \|x - y\|_2 \quad \text{holds for any } x, y \in \Delta^n \text{ and } M \in \mathbb{S}^{n \times n}, \quad (88)$$

where, $\|M\|_o := \max_{\|v\|_2=1} |\langle v, Mv \rangle|$, is the operator norm. This follows since

$$\begin{aligned} x^\top Mx - y^\top My &= (x + y)^\top M(x - y) \\ &\leq \|x - y\|_2 \|M(x + y)\|_o \quad \text{from the Cauchy-Schwartz inequality} \\ &\leq \|x - y\|_2 \|x + y\|_2 \|M\|_o \quad \text{from the definition of } \|\cdot\|_o \\ &\leq 2 \|M\|_o \|x - y\|_2 \quad \text{since } \|z\|_2 \leq \|z\|_1 = 1 \text{ for all } z \in \Delta^n. \end{aligned}$$

Now, for any $x, y \in \Delta^n$, we conclude from (87b) that

$$\begin{aligned} dV_x(\mathbb{P}; \mathbb{Q}) - dV_y(\mathbb{P}; \mathbb{Q}) &= x^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})x - y^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})y \\ &\quad + (x^\top \mu_{\mathbb{P}})^2 + (x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 - (x^\top \mu_{\mathbb{Q}})^2 \\ &\quad - (y^\top \mu_{\mathbb{P}})^2 - (y^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 + (y^\top \mu_{\mathbb{Q}})^2. \end{aligned}$$

Employing (88) for $M = (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})$, $\mu_{\mathbb{P}}\mu_{\mathbb{P}}^\top$, $\mu_{\mathbb{Q}}\mu_{\mathbb{Q}}^\top$, and $(\mu_{\mathbb{P}} - \mu_{\mathbb{Q}})(\mu_{\mathbb{P}} - \mu_{\mathbb{Q}})^\top$; we have the inequalities

$$\left\{ \begin{array}{l} x^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})x - y^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})y \leq 2 \|\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}}\|_o \|x - y\|_2 \\ (x^\top \mu_{\mathbb{P}})^2 - (y^\top \mu_{\mathbb{P}})^2 \leq 2 \|\mu_{\mathbb{P}}\|_2^2 \|x - y\|_2 \\ (y^\top \mu_{\mathbb{Q}})^2 - (y^\top \mu_{\mathbb{Q}})^2 \leq 2 \|\mu_{\mathbb{Q}}\|_2^2 \|x - y\|_2 \\ (x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 - (y^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 \leq 2 \|\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}\|_2^2 \|x - y\|_2, \end{array} \right. \quad (89)$$

where we have used the fact that $\|vv^\top\|_o = \|v\|_2^2$ for any v . These inequalities give us the upper bound

$$dV_x(\mathbb{P}; \mathbb{Q}) - dV_y(\mathbb{P}; \mathbb{Q}) \leq 2 \left(\|\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}}\|_o + \|\mu_{\mathbb{P}}\|_2^2 + \|\mu_{\mathbb{Q}}\|_2^2 + \|\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}\|_2^2 \right) \|x - y\|_2$$

Since $\|\mu_{\mathbb{P}}\|_2 \leq \|\mu_{\mathbb{P}} - \hat{\mu}\|_2 + \|\hat{\mu}\|_2$, and $\mathbb{P}, \hat{\mathbb{P}} \in \mathcal{P}$, we conclude from (58) that $\|\mu_{\mathbb{P}}\|_2 \leq B_\mu + \|\hat{\mu}\|_2$, (and similarly for $\mu_{\mathbb{Q}}$). This together with the condition $\|\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}}\|_o \leq B_\Sigma$ from (58), finally gives

$$dV_x(\mathbb{P}; \mathbb{Q}) - dV_y(\mathbb{P}; \mathbb{Q}) \leq 2 \left(B_\Sigma + 2(B_\mu + \|\hat{\mu}\|_2)^2 + B_\mu^2 \right) \|x - y\|_2$$

The constant C_1 in the assertion (i) of the Lemma is immediately picked as $C_1 = 2 \left(B_\Sigma + 2(B_\mu + \|\hat{\mu}\|_2)^2 + B_\mu^2 \right)$.

Proof of Lemma 7.2(ii): Smoothness. Consider any $x \in \mathcal{X}$, $\mathbb{P}, \mathbb{Q} \in \mathcal{W}_m(\hat{\mathbb{P}}, \rho)$, and let $\mathbb{P}_\gamma := \mathbb{P} + \gamma(\mathbb{Q} - \mathbb{P})$ for $\gamma \in [0, 1]$. Using $\Sigma_{\mathbb{P}_\gamma} = \Sigma_{\mathbb{P}} + \gamma(\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})$ and $\mu_{\mathbb{P}_\gamma} = \mu_{\mathbb{P}} + \gamma(\mu_{\mathbb{Q}} - \mu_{\mathbb{P}})$, we

conclude from (87a), that

$$\begin{aligned} dV_x(\mathbb{P}; \mathbb{P}_\gamma) &= x^\top (\Sigma_{\mathbb{P}_\gamma} - \Sigma_{\mathbb{P}})x + 2(x^\top \mu_{\mathbb{P}})(x^\top (\mu_{\mathbb{P}_\gamma} - \mu_{\mathbb{P}})) \\ &= \gamma \left(x^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})x - 2(x^\top \mu_{\mathbb{P}})(x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}})) \right). \end{aligned}$$

Similarly, we also have

$$\begin{aligned} dV_x(\mathbb{P}_\gamma; \mathbb{P}) &= x^\top (\Sigma_{\mathbb{P}} - \Sigma_{\mathbb{P}_\gamma})x + 2(x^\top \mu_{\mathbb{P}_\gamma})(x^\top (\mu_{\mathbb{P}} - \mu_{\mathbb{P}_\gamma})) \\ &= 2\gamma^2 (x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 + \gamma \left(2(x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))(x^\top \mu_{\mathbb{P}}) - x^\top (\Sigma_{\mathbb{Q}} - \Sigma_{\mathbb{P}})x \right). \end{aligned}$$

Combining the two equalities, we get

$$dV_x(\mathbb{P}; \mathbb{P}_\gamma) + dV_x(\mathbb{P}_\gamma; \mathbb{P}) = 2\gamma^2 (x^\top (\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}))^2 \leq 2\gamma^2 \|x\|_2^2 \|\mu_{\mathbb{Q}} - \mu_{\mathbb{P}}\|_2^2 \leq \gamma^2 (2B_\mu^2).$$

The last inequality follows from (58) and the fact that $\|x\|_2 \leq \|x\|_1 = 1$, for every $x \in \Delta^n$. Thus, the risk measure V_x is $(2B_\mu^2)$ -smooth, uniformly over $x \in \mathcal{X}$. This proves assertion (ii) of the Lemma and the proof is now complete. \square

Proofs for the case of unconstrained support, ($\Xi = \mathbb{R}^n$).

Proof of Lemma 7.3. Substituting $q = s\bar{q}$ for $s \geq 0$ and $\|\bar{q}\| = 1$ in (63), it is equivalently reformulated as

$$\begin{cases} \max_{s \geq 0, \bar{q} \in \mathbb{R}^n} & (s(x^\top \bar{q}) + x^\top (\xi - v))^2 - \eta s^m \\ \text{subject to} & \|\bar{q}\| = 1. \end{cases} \quad (90)$$

Observe that $(s(x^\top \bar{q}) + x^\top (\xi - v))^2 \leq (s|x^\top \bar{q}| + |x^\top (\xi - v)|)^2$, where, equality holds if and only if $\text{sgn}(x^\top \bar{q}) = \text{sgn}(x^\top (\xi - v))$. Moreover, applying the Holder's inequality $|x^\top \bar{q}| \leq \|\bar{q}\| \|x\|_* = \|x\|_*$ yields

$$(s(x^\top \bar{q}) + x^\top (\xi - v))^2 \leq (s\|x\|_* + |x^\top (\xi - v)|)^2. \quad (91)$$

The upper bound (91) is achieved if and only if $\bar{q} = \text{sgn}(x^\top (\xi - v))\bar{q}_x$, for any $\bar{q}_x \in \text{argmax}_{\{\|\bar{q}\| \leq 1\}} x^\top \bar{q}$. Thus, every such \bar{q} is an optimal solution to (90), irrespective of s . Simplifying the optimization over \bar{q} in (90), the problem reduces to

$$\max_{s \geq 0} (s\|x\|_* + |x^\top (\xi - v)|)^2 - \eta s^m. \quad (92)$$

Now, an optimal solution to (63) exists if and only if (92) admits an optimal solution s^* ; in which case, we have $q(\eta, \xi) = s^* \text{sgn}(x^\top (\xi - v))\bar{q}_x$.

Optimal solution to (92) under different settings. If $m < 2$, it is straightforward to see that the optimal value of (92) is unbounded irrespective of η . Consequently, no optimal solution exists for the maximization problem (63) in this setting. On the contrary, if $m > 2$, it is also easily seen that the objective function of (92) is coercive if and only if $\eta > 0$. Therefore, the maximal value of (92) (and Consequently (63)), is bounded and achieved.

For $m = 2$, the objective function of (92) can be simplified to

$$-s^2(\eta - \|x\|_*^2) + 2s\|x\|_* |x^\top (\xi - v)| + |x^\top (\xi - v)|^2.$$

It is clear that the optimal value of (92) is unbounded if $\eta < \|x\|_*^2$. Whereas, if $\eta > \|x\|_*^2$, it is bounded, in which case, equating the derivative w.r.t. s equal to 0, gives that the optimal value

of (92) (and consequently, also $J(\eta, \xi)$) is equal to $\frac{\eta}{\eta - \|x\|_*^2} |x^\top(\xi - v)|^2$, which is achieved at $s^* = \frac{\|x\|_* |x^\top(\xi - v)|}{\eta - \|x\|_*^2}$. Finally, if $\eta = \|x\|_*^2$, the optimal value of (92) is unbounded if $x^\top(\xi_i - v) \neq 0$; otherwise, it is bounded and equal to 0 which is achieved for any $s \geq 0$. \square

Proof of Lemma 7.4. On the one hand, if $|x^\top(\xi_i - v)| = 0$ for all $i = 1, 2, \dots, N$, we now see that the dual problem (62) reduces to

$$\inf_{\eta \geq \|x\|_*^2} \eta \rho^2,$$

which admits the optimal solution $\eta^* = \|x\|_*^2$ with an optimal value $\|x\|_*^2 \rho^2$. On the other hand, if $|x^\top(\xi_i - v)| > 0$ for at least some $i \in \{1, 2, \dots, N\}$, the dual problem (62) reduces to

$$\inf_{\eta > \|x\|_*^2} \eta \rho^2 + \frac{\eta}{\eta - \|x\|_*^2} \left(\frac{1}{N} \sum_{i=1}^N |x^\top(\xi_i - v)|^2 \right). \quad (93)$$

It is easily verified that the first-order optimality conditions for (93) are satisfied at

$$\eta_x := \|x\|_*^2 + \frac{\|x\|_*}{\rho} \sqrt{\frac{1}{N} \sum_{i=1}^N |x^\top(\xi_i - v)|^2}.$$

Thus, η_x is the unique optimal solution to (93), and consequently, to the dual problem (62). Moreover, due to strong duality of (62), we also know that $\mathbf{Q}_x(\xi) = \frac{1}{N} \sum_{i=1}^N \delta(\xi - (\xi_i + q'_i))$ is an optimal solution to the linear worst case distribution problem (61) for any collection $q'_i \in q(\eta, \xi_i)$, $i = 1, 2, \dots, N$. \square

Proof of Proposition 7.5. We first see that if $x^* \in \operatorname{argmin}_{x \in \mathcal{X}} (\rho \|x\|_* + \sqrt{\langle x, Vx \rangle})$, then the first-order necessary optimality conditions imply that there exists a sub-gradient g of the function $(\rho \|x\|_* + \sqrt{\langle x, Vx \rangle})$ at x^* for which the inclusion $x^* \in \operatorname{argmin}_{y \in \mathcal{X}} \langle g, y \rangle$ holds. A quick look reveals that g must be of the form $\rho \bar{q}_{x^*} + (1/\sqrt{\langle x^*, Vx^* \rangle}) Vx^*$, where $\bar{q}_{x^*} \in \operatorname{argmax}_{\{\|\bar{q}\| \leq 1\}} \langle x^*, \bar{q} \rangle$. Therefore, there exists some $\bar{q}_{x^*} \in \operatorname{argmax}_{\{\|\bar{q}\| \leq 1\}} \langle x^*, \bar{q} \rangle$ such that x^* satisfies the inclusion

$$x^* \in \operatorname{argmin}_{y \in \mathcal{X}} \left(\rho \langle y, \bar{q}_{x^*} \rangle + \frac{\langle x^*, Vy \rangle}{\sqrt{\langle x^*, Vx^* \rangle}} \right). \quad (94)$$

Selecting such a \bar{q}_{x^*} to define \mathbb{P}^* in (67), we also verify that

$$\mu_{\mathbb{P}^*} = \frac{1}{N} \sum_{i=1}^N (\xi_i + q_i^*) = \hat{\mu} + \frac{\rho \frac{1}{N} \sum_{i=1}^N \langle x^*, \xi_i - \hat{\mu} \rangle}{\sqrt{\langle x^*, (\hat{\Sigma} - \hat{\mu} \hat{\mu}^\top) x^* \rangle}} \bar{q}_{x^*} = \hat{\mu}.$$

Now, we establish that the pair (x^*, \mathbb{P}^*) is a saddle point by proving that both x^* and \mathbb{P}^* are optimal solutions to their respective problems while the other is held fixed.

Optimality of x^ for the minimization condition.* We begin by showing that the inclusion $x^* \in \operatorname{argmin}_{y \in \mathcal{X}} V(y, \mathbb{P}^*)$ holds, by showing that the corresponding first-order optimality condition

$$x^* \in \operatorname{argmin}_{y \in \mathcal{X}} \langle \nabla_1 V(x^*, \mathbb{P}^*), y \rangle = \operatorname{argmin}_{y \in \mathcal{X}} \left\langle (\Sigma_{\mathbb{P}^*} - \hat{\mu} \hat{\mu}^\top) x^*, y \right\rangle, \quad (95)$$

is satisfied. This condition is also sufficient for optimality since $V(y, \mathbb{P}^*)$ is convex in y . Simplifying the cost function in (95) yields

$$\begin{aligned} & \left\langle (\Sigma_{\mathbb{P}^*} - \widehat{\mu}\widehat{\mu}^\top)x^*, y \right\rangle \\ &= \frac{1}{N} \sum_{i=1}^N \langle x^*, \xi_i + q_i^* \rangle \langle y, \xi_i + q_i^* \rangle - \left\langle x^*, (\widehat{\mu}\widehat{\mu}^\top)y \right\rangle \\ &= \frac{1}{N} \sum_{i=1}^N \left\langle x^*, (\xi_i \xi_i^\top - \widehat{\mu}\widehat{\mu}^\top)y \right\rangle + \langle x^*, \xi_i \rangle \langle y, q_i^* \rangle + \langle x^*, q_i^* \rangle \langle y, \xi_i \rangle + \langle x^*, q_i^* \rangle \langle y, q_i^* \rangle. \end{aligned}$$

Letting $V := \widehat{\Sigma} - \widehat{\mu}\widehat{\mu}^\top$, we see that $V = \frac{1}{N} \sum_{i=1}^N (\xi_i - \widehat{\mu})(\xi_i - \widehat{\mu})^\top$, and $\frac{1}{N} \sum_{i=1}^N |\langle x^*, \xi_i - \widehat{\mu} \rangle|^2 = \langle x^*, Vx^* \rangle$, from which we simplify

$$\left\{ \begin{aligned} \frac{1}{N} \sum_{i=1}^N \langle x^*, \xi_i \rangle \langle y, q_i^* \rangle &= \rho \langle y, \bar{q}_{x^*} \rangle \frac{\frac{1}{N} \sum_{i=1}^N \langle x^*, \xi_i - \widehat{\mu} \rangle \langle \xi_i, x^* \rangle}{\sqrt{\langle x^*, Vx^* \rangle}} &= \rho \langle y, \bar{q}_{x^*} \rangle \sqrt{\langle x^*, Vx^* \rangle}, \\ \frac{1}{N} \sum_{i=1}^N \langle x^*, q_i^* \rangle \langle y, \xi_i \rangle &= \rho \|x^*\|_* \frac{\frac{1}{N} \sum_{i=1}^N \langle x^*, \xi_i - \widehat{\mu} \rangle \langle \xi_i, y \rangle}{\sqrt{\langle x^*, Vx^* \rangle}} &= \rho \|x^*\|_* \frac{\langle x^*, Vy \rangle}{\sqrt{\langle x^*, Vx^* \rangle}}, \\ \frac{1}{N} \sum_{i=1}^N \langle x^*, q_i^* \rangle \langle y, q_i^* \rangle &= \rho^2 \|x^*\|_* \langle y, \bar{q}_{x^*} \rangle \frac{\frac{1}{N} \sum_{i=1}^N |\langle x^*, \xi_i - \widehat{\mu} \rangle|^2}{\langle x^*, Vx^* \rangle} &= \rho^2 \|x^*\|_* \langle y, \bar{q}_{x^*} \rangle. \end{aligned} \right.$$

Employing the above relations, we see that the objective function in (94) simplifies to

$$\begin{aligned} \left\langle (\Sigma_{\mathbb{P}^*} - \widehat{\mu}\widehat{\mu}^\top)x^*, y \right\rangle &= \langle x^*, Vy \rangle + \rho \|x^*\|_* \frac{\langle x^*, Vy \rangle}{\sqrt{\langle x^*, Vx^* \rangle}} + \rho \langle y, \bar{q}_{x^*} \rangle \sqrt{\langle x^*, Vx^* \rangle} + \rho^2 \|x^*\|_* \langle y, \bar{q}_{x^*} \rangle \\ &= \left(\sqrt{\langle x^*, Vx^* \rangle} + \rho \|x^*\|_* \right) \left(\frac{\langle x^*, Vy \rangle}{\sqrt{\langle x^*, Vx^* \rangle}} + \rho \langle y, \bar{q}_{x^*} \rangle \right). \end{aligned}$$

In view of the inclusion (94), the sufficient optimality condition (95) follows immediately. Consequently, the inclusion $x^* \in \arg\min_{x \in \mathcal{X}} V(x, \mathbb{P}^*)$ also holds.

Optimality of \mathbb{P}^ for the maximization condition.* Similar to the proof of the minimization condition, we establish the maximization condition of the saddle point by showing that the first-order optimality conditions are satisfied. We first recall from (65) (with $v = \mu_{\mathbb{P}^*} = \widehat{\mu}$) that

$$\eta_{x^*} = \|x^*\|_*^2 + \frac{\|x^*\|_*}{\rho} \sqrt{\langle x^*, (\widehat{\Sigma} - \widehat{\mu}\widehat{\mu}^\top)x^* \rangle}.$$

On the one hand, if $\langle x^*, (\widehat{\Sigma} - \widehat{\mu}\widehat{\mu}^\top)x^* \rangle = 0$, we have $\eta_{x^*} = \|x^*\|_*^2$, in which case, we conclude from assertion (ii-c) of Lemma 7.3 that $0 \in q(\eta_{x^*}, \xi_i, \widehat{\mu})$ for all $i = 1, 2, \dots, N$. Moreover, $\langle x^*, (\widehat{\Sigma} - \widehat{\mu}\widehat{\mu}^\top)x^* \rangle = 0$ also implies that $\langle x^*, \xi_i - \widehat{\mu} \rangle = 0$, and hence $q_i^* = 0$ for all $i = 1, 2, \dots, N$. Thus, we have $q_i^* \in q(\eta_{x^*}, \xi_i, \mu_{\mathbb{P}^*})$ for all $i = 1, 2, \dots, N$.

On the other hand, if $\langle x^*, (\widehat{\Sigma} - \widehat{\mu}\widehat{\mu}^\top)x^* \rangle \neq 0$, we have $\eta_{x^*} > \|x^*\|_*^2$, in which case, we see

$$q_i^* = \frac{\rho \langle x^*, \xi_i - \widehat{\mu} \rangle}{\sqrt{\langle x^*, (\widehat{\Sigma} - \widehat{\mu}\widehat{\mu}^\top)x^* \rangle}} \bar{q}_{x^*} = \frac{\|x^*\|_* \langle x^*, \xi_i - \widehat{\mu} \rangle}{\eta_{x^*} - \|x^*\|_*^2} \bar{q}_{x^*} \in q(\eta_{x^*}, \xi_i, \widehat{\mu}) \quad \text{for all } i = 1, 2, \dots, N,$$

where the last inclusion follows from assertion (ii-b) of Lemma 7.3 since $\bar{q}_{x^*} \in \operatorname{argmax}_{\|\bar{q}\| \leq 1} \langle x^*, \bar{q} \rangle$. Finally, noting that $\hat{\mu} = \mu_{\mathbb{P}^*}$, we have $q_i^* \in q(\eta_{x^*}, \xi_i, \mu_{\mathbb{P}^*})$ for all $i \leq N$. Since $\mathbb{P}^* = \frac{1}{N} \sum_{i=1}^N \delta(\xi - \xi_i - q_i^*)$, we conclude from Lemma 7.4 that $\mathbb{P}^* \in \operatorname{argmax}_{\mathbb{Q} \in \mathcal{P}} dV_{x^*}(\mathbb{P}^*; \mathbb{Q})$. Thus, \mathbb{P}^* satisfies the first-order optimality conditions which are also sufficient. Hence, in view of Remark 3.5, we conclude that the inclusion $\mathbb{P}^* \in \operatorname{argmax}_{\mathbb{P} \in \mathcal{P}} V(x^*, \mathbb{P})$ also holds. Therefore, (x^*, \mathbb{P}^*) is indeed a saddle point of (56). The proof is now complete. \square

Proofs for the case of ellipsoidal support ($\Xi = \mathcal{E}_M$).

Proof of Lemma 7.6. Letting $q' := q + \xi$, the maximization problem (63) can be equivalently written in terms of q' as

$$\sup_{\{q': \langle q', Mq' \rangle \leq 1\}} \left\langle q', (xx^\top - \eta \mathbb{I}_n)q' \right\rangle + 2 \left\langle q', \eta \xi - (xx^\top)v \right\rangle + ((x^\top v)^2 - \eta \|\xi\|_2^2). \quad (96)$$

Even though $M \succeq 0$, which makes the feasible set of (96) convex; it is to be observed that the objective function is concave if and only if $\eta \geq \|x\|_2^2$. Therefore, (96) is not a convex problem in general. However, we observe that (96) is a quadratically constrained quadratic program, that is strictly feasible. For such problems, the *S-procedure* guarantees an equivalent reformulation as a tractable SDP

$$\begin{cases} \min_{\lambda \in [0, +\infty), \theta \in \mathbb{R}} & -\theta \\ \text{subject to} & \begin{bmatrix} \eta \mathbb{I}_n - xx^\top + \lambda M & (xx^\top)v - \eta \xi \\ (xx^\top)v^\top - \eta \xi^\top & \eta \|\xi\|_2^2 - (x^\top v)^2 - \lambda - \theta \end{bmatrix} \succeq 0, \end{cases} \quad (97)$$

with the optimal values of (97) and (96) being equal. Moreover, if (θ^*, λ^*) is a solution to the SDP (97), we also conclude from the S-procedure that $q' = (\eta \mathbb{I}_n - xx^\top + \lambda^* M)^{-1}(\eta \xi - xx^\top v)$, is an optimal solution to the maximization problem (96). Consequently, for every $\eta \geq 0$, we conclude

$$q(\eta, \xi) = -\xi + (\eta \mathbb{I}_n - xx^\top + \lambda^* M)^{-1}(\eta \xi - xx^\top v),$$

is a solution to the maximization problem (96). Substituting for each $i = 1, 2, \dots, N$, the maximization problem over q_i in (62) with its equivalent SDP (97), we immediately arrive at (68). Now, suppose $(\eta^*, \lambda^*, \theta^*)$ is a solution to the SDP (68). Then, the pair (\mathbb{Q}_x, η_x) as given by (69) is a solution to the FW problem (61) and its dual (62), respectively. This concludes the proof. \square

REFERENCES

- [1] A. BECK AND M. TEBoulLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM journal on imaging sciences, 2 (2009), pp. 183–202.
- [2] D. P. BERTSEKAS, *Control of uncertain systems with a set-membership description of the uncertainty.*, PhD thesis, Massachusetts Institute of Technology, 1971.
- [3] D. BERTSIMAS, D. B. BROWN, AND C. CARAMANIS, *Theory and applications of robust optimization*, SIAM review, 53 (2011), pp. 464–501.
- [4] J. BLANCHET, L. CHEN, AND X. Y. ZHOU, *Distributionally robust mean-variance portfolio selection with Wasserstein distances*, Management Science, 68 (2022), pp. 6382–6410.
- [5] J. BLANCHET, K. MURTHY, AND F. ZHANG, *Optimal transport-based distributionally robust optimization: Structural properties and iterative schemes*, Mathematics of Operations Research, 47 (2022), pp. 1500–1529.
- [6] S. P. BOYD AND L. VANDENBERGHE, *Convex optimization*, Cambridge university press, 2004.

- [7] J. CAI, J. LI, AND T. MAO, *Distributionally robust optimization under distorted expectations*, Available at SSRN 3566708, (2020).
- [8] Y. CHEN AND Y. HU, *Multivariate coherent risk measures induced by multivariate convex risk measures*, Positivity, 24 (2020), pp. 711–727.
- [9] K. L. CLARKSON, *Coresets, sparse greedy approximation, and the frank-wolfe algorithm*, ACM Transactions on Algorithms (TALG), 6 (2010), pp. 1–30.
- [10] E. DELAGE AND Y. YE, *Distributionally robust optimization under moment uncertainty with application to data-driven problems*, Operations research, 58 (2010), pp. 595–612.
- [11] J. C. DUCHI AND H. NAMKOONG, *Learning models with uniform performance via distributionally robust optimization*, The Annals of Statistics, 49 (2021), pp. 1378–1406.
- [12] E. ERDOĞAN AND G. IYENGAR, *Ambiguous chance constrained problems and robust optimization*, Mathematical Programming, 107 (2006), pp. 37–61.
- [13] M. FRANK AND P. WOLFE, *An algorithm for quadratic programming*, Naval research logistics quarterly, 3 (1956), pp. 95–110.
- [14] R. GAO, *Finite-sample guarantees for Wasserstein distributionally robust optimization: Breaking the curse of dimensionality*, Operations Research, (2022).
- [15] R. GAO, X. CHEN, AND A. J. KLEYWEGT, *Wasserstein distributionally robust optimization and variation regularization*, Operations Research, (2022).
- [16] R. GAO AND A. KLEYWEGT, *Distributionally robust stochastic optimization with Wasserstein distance*, Mathematics of Operations Research, (2022).
- [17] J. GOH AND M. SIM, *Distributionally robust optimization and its tractable approximations*, Operations research, 58 (2010), pp. 902–917.
- [18] E. Y. HAMEDANI, A. JALILZADEH, N. S. AYBAT, AND U. V. SHANBHAG, *Iteration complexity of randomized primal-dual methods for convex-concave saddle point problems*, arXiv preprint arXiv:1806.04118, 5 (2018).
- [19] M. JAGGI, *Revisiting frank-wolfe: Projection-free sparse convex optimization*, in International Conference on Machine Learning, PMLR, 2013, pp. 427–435.
- [20] R. JIANG AND Y. GUAN, *Data-driven chance constrained stochastic program*, Mathematical Programming, 158 (2016), pp. 291–327.
- [21] C. KENT, J. BLANCHET, AND P. GLYNN, *Frank-wolfe methods in probability space*, preprint available at arXiv:2105.05352, (2021).
- [22] G. M. KORPELEVICH, *The extragradient method for finding saddle points and other problems*, Matecon, 12 (1976), pp. 747–756.
- [23] D. KUHN, P. MOHAJERIN ESFAHANI, V. A. NGUYEN, AND S. A. S., *Wasserstein distributionally robust optimization: Theory and applications in machine learning*, in Operations research & management science in the age of analytics, Informs, 2019, pp. 130–166.
- [24] N. LANZETTI, S. BOLOGNANI, AND F. DÖRFLER, *First-order conditions for optimization in the Wasserstein space*, arXiv preprint arXiv:2209.12197, (2022).
- [25] L. LIU, Y. ZHANG, Z. YANG, R. BABANEZHAD, AND Z. WANG, *Infinite-dimensional optimization for zero-sum games via variational transport*, in International Conference on Machine Learning, PMLR, 2021, pp. 7033–7044.
- [26] H. M. MARKOWITZ, *Portfolio selection, the journal of finance. 7 (1)*, 1952.
- [27] P. MOHAJERIN ESFAHANI AND D. KUHN, *Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations*, Mathematical Programming, 171 (2018), pp. 115–166.
- [28] A. NEMIROVSKI, *Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems*, SIAM Journal on Optimization, 15 (2004), pp. 229–251.
- [29] A. NEMIROVSKIJ AND D. YUDIN, *Problem complexity and method efficiency in optimization*, New York, (1983).
- [30] Y. NESTEROV, *A method of solving a convex programming problem with convergence rate $o(1/k^{**2})$* , Doklady Akademii Nauk SSSR, 269 (1983), p. 543.

- [31] Y. NESTEROV, *Introductory lectures on convex optimization: A basic course*, vol. 87, Springer Science & Business Media, 2003.
- [32] Y. NESTEROV, *Smooth minimization of non-smooth functions*, *Mathematical programming*, 103 (2005), pp. 127–152.
- [33] D. H. NGUYEN AND T. SAKURAI, *A particle-based algorithm for distributional optimization on \textit{Constrained Domains} via variational transport and mirror descent*, arXiv preprint arXiv:2208.00587, (2022).
- [34] V. A. NGUYEN, S. S. ABADEH, D. FILIPOVIĆ, AND D. KUHN, *Mean-covariance robust risk measurement*, preprint available at arXiv:2112.09959, (2021).
- [35] V. A. NGUYEN, S. SHAFIEEZADEH-ABADEH, D. KUHN, AND P. MOHAJERIN ESFAHANI, *Bridging bayesian and minimax mean square error estimation via wasserstein distributionally robust optimization*, *Mathematics of Operations Research*, 48 (2023), pp. 1–37.
- [36] F. PEDREGOSA, G. NEGIAR, A. ASKARI, AND M. JAGGI, *Linearly convergent frank-wolfe with backtracking line-search*, in *International conference on artificial intelligence and statistics*, PMLR, 2020, pp. 1–10.
- [37] G. PFLUG AND D. WOZABAL, *Ambiguity in portfolio selection*, *Quantitative Finance*, 7 (2007), pp. 435–442.
- [38] K. POSTEK, D. DEN HERTOOG, AND B. MELENBERG, *Computationally tractable counterparts of distributionally robust constraints on risk measures*, *SIAM Review*, 58 (2016), pp. 603–650.
- [39] H. RAHIMIAN AND S. MEHROTRA, *Distributionally robust optimization: A review*, preprint available at arXiv:1908.05659, (2019).
- [40] N. RUJEERAPAIBOON, D. KUHN, AND W. WIESEMANN, *Robust growth-optimal portfolios*, *Management Science*, 62 (2016), pp. 2090–2109.
- [41] H. SCARF, *A min max solution of an inventory problem*, *Studies in the mathematical theory of inventory and production*, (1958).
- [42] S. SHAFIEEZADEH-ABADEH, L. AOLARITEI, F. DÖRFLER, AND D. KUHN, *New perspectives on regularization and computation in optimal transport-based distributionally robust optimization*, preprint available at arXiv:2303.03900, (2023).
- [43] S. SHAFIEEZADEH-ABADEH, P. MOHAJERIN ESFAHANI, AND D. KUHN, *Distributionally robust logistic regression*, *Advances in Neural Information Processing Systems*, 28 (2015).
- [44] M. SHERIFF AND P. MOHAJERIN ESFAHANI, *Distributionally Robust Optimization*. <https://github.com/MRayyans/DRO>, 2024.
- [45] J. E. SMITH AND R. L. WINKLER, *The optimizer’s curse: Skepticism and postdecision surprise in decision analysis*, *Management Science*, 52 (2006), pp. 311–322.
- [46] H. SONG, X. ZENG, Y. CHEN, AND Y. HU, *Multivariate shortfall and divergence risk statistics*, *Entropy*, 21 (2019), p. 1031.
- [47] F. UHLIG, *A recurring theorem about pairs of quadratic forms and extensions: A survey*, *Linear algebra and its applications*, 25 (1979), pp. 219–237.
- [48] A. M. V. DEM’YANOV, F. RUBINOV, *Minimization of functionals in normed spaces*, *SIAM Journal on Control*, 6 (1968), pp. 73–88.
- [49] A. M. V. DEMYANOV, F. RUBINOV, *Approximate methods in optimization problems*, *Modern Analytic and Computational Methods in Science and Mathematics*, (1970).
- [50] J. WANG, R. GAO, AND Y. XIE, *Sinkhorn distributionally robust optimization*, preprint available at arXiv:2109.11926, (2021).
- [51] W. WIESEMANN, D. KUHN, AND M. SIM, *Distributionally robust convex optimization*, *Operations Research*, 62 (2014), pp. 1358–1376.
- [52] J. ZHEN, D. KUHN, AND W. WIESEMANN, *Mathematical foundations of robust and distributionally robust optimization*, preprint available at arXiv:2105.00760, (2021).
- [53] S. ZYMLER, D. KUHN, AND B. RUSTEM, *Worst-case value at risk of nonlinear portfolios*, *Management Science*, 59 (2013), pp. 172–188.