Robustifying Conditional Portfolio Decisions
via Optimal Transport

VIET ANH NGUYEN, FAN ZHANG, JOSÉ BLANCHET, ERICK DELAGE, YINYU YE

ABSTRACT. We propose a data-driven portfolio selection model that integrates side information, conditional estimation and robustness using the framework of distributionally robust optimization. Conditioning on the observed side information, the portfolio manager solves an allocation problem that minimizes the worst-case conditional risk-return trade-off, subject to all possible perturbations of the covariate-return probability distribution in an optimal transport ambiguity set. Despite the non-linearity of the objective function in the probability measure, we show that the distributionally robust portfolio allocation with side information problem can be reformulated as a finite-dimensional optimization problem. If portfolio decisions are made based on either the mean-variance or the mean-Conditional Value-at-Risk criterion, the resulting reformulation can be further simplified to second-order or semi-definite cone programs. Empirical studies in the US and Chinese equity markets demonstrate the advantage of our integrative framework against other benchmarks.

1. INTRODUCTION

We study distributionally robust portfolio decision rules which are informed by conditioning on observed side information. In the presence of contextual information which might be relevant for predicting future returns, our objective is to direct the statistical task (i.e., conditional estimation) towards the downstream task of selecting a portfolio which is robust to distributional shifts (including on local conditioning). Our ultimate goal is to provide a tractable, non-parametric, data-driven approach which bypasses the need for computing global conditional expectations. At the same time, our framework is flexible to accommodate a wide range of portfolio optimization selection criteria, as well as extend to other diverse applications. In order to explain our methodology and contribution in detail, let us consider a conventional single-period portfolio optimization selection problem of the form

$$\min_{\alpha \in \mathcal{A}} \ \mathcal{R}_{\mathbb{P}}[Y^{\top}\alpha] - \eta \cdot \mathbb{E}_{\mathbb{P}}[Y^{\top}\alpha],$$

where the real-valued vector α denotes the portfolio allocation choice within a feasible region $\mathcal{A}$, and the random vector Y denotes the assets' future return under a probability measure $\mathbb{P}$. The return of the portfolio is denoted by $Y^{\top}\alpha$, while $\mathcal{R}_{\mathbb{P}}[\cdot]$ captures the risk associated with α, and $\eta \geq 0$ is a parameter that weights the preference between the portfolio return and the associated risk.

Different choices of the measure of risk $\mathcal{R}_{\mathbb{P}}[\cdot]$ lead to different portfolio optimization models, including the mean-variance model [43], the mean-standard deviation model [37], the mean-Value-at-Risk model [6], and the mean-Conditional Value-at-Risk (CVaR) model [56]. The measure of risk $\mathcal{R}_{\mathbb{P}}[\cdot]$ often is chosen to belong to certain classes in order to induce desirable properties. For instance, the class of coherent risk measures [2] or the class of convex risk measures [28] are examples of risk measure classes which are often used in portfolio optimization [42]. In our setting, we consider a family of risk measures which admits a stochastic optimization representation of the form $\mathcal{R}_{\mathbb{P}}[Y^{\top}\alpha] = \min_{\beta \in \mathcal{B}} \mathbb{E}_{\mathbb{P}}[r(Y^{\top}\alpha, \beta)]$ for some auxiliary function r and statistical variable β. Hence, we focus on a generic portfolio allocation that minimizes an expected loss of

the form

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \ \mathbb{E}_{\mathbb{P}}[\ell(Y, \alpha, \beta)], \qquad \ell(Y, \alpha, \beta) \triangleq r(Y^\top \alpha, \beta) - \eta \cdot Y^\top \alpha. \tag{1}$$

In problem (1), optimizing over β characterizes the improvement of risk estimation, which is entangled with the procedure of finding the optimal portfolio allocation α. For example, consider when $\mathcal{R}_{\mathbb{P}}$ is the variance of the portfolio return, then $\mathcal{R}_{\mathbb{P}}[Y^\top \alpha] = \min_{\beta \in \mathbb{R}} \ \mathbb{E}_{\mathbb{P}}[(Y^\top \alpha - \beta)^2]$. This further implies that the optimal solution of β coincides with the expected return $\mathbb{E}_{\mathbb{P}}[Y^\top \alpha]$, and one thus can view β as an auxiliary statistical variable representing the estimation of the expected portfolio return.

Despite the simplicity and popularity of the portfolio optimization models (1), they are challenging to solve because the distribution $\mathbb{P}$ of the random stock returns is unknown to the portfolio managers. To overcome this issue, problem (1) is usually solved by resorting to a plug-in estimator $\widehat{\mathbb{P}}$ of $\mathbb{P}$, and this estimator is usually inferred from the available data. Unfortunately, solving (1) using an estimated distribution $\widehat{\mathbb{P}}$ may amplify the statistical estimation error. The corresponding portfolio allocation is vulnerable to the estimation errors of the input: *even small changes in the input parameters can result in large changes in the weights of the optimal portfolio* [17, 15]. Consequently, in terms of out-of-sample performance, the advantage of deploying the optimal portfolio weights is overwhelmed by the offset of estimation error [21]. To mitigate the estimation error, one may resort, at least, to the following two strategies:

(1) reduce the estimation bias by incorporating side information into the portfolio allocation problems;
(2) diminish the impact of estimation variance by employing robust portfolio optimization models.

Both of the above-mentioned strategies have led to successful stories. Starting from the capital asset-pricing model [59, 9], numerous studies exploit side information to explain and/or predict the cross-sectional variation of the stock returns. The predictive side information may include macro-economic factors [27], the company's financial statements [25] and historical trading data [32]. Alternatively, robust optimization formulations are applied to the portfolio allocation tasks to alleviate the impact of statistical error due to input uncertainty [51, 19, 66].

In this paper, we endeavor to unify both strategies of leveraging side information and enhancing robustness. In order to incorporate the side information into portfolio allocation, we use a random vector X to denote the side information that is correlated with the stock return Y, and consider the following conditional stochastic optimization (CSO) problem:

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \ \mathbb{E}_{\mathbb{P}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}], \tag{2}$$

where the set $\mathcal{N}$ represents our prior (or most current) knowledge about the covariate X. In problem (2), the distribution $\mathbb{P}$ is now lifted to a joint distribution between the covariate X and the return Y. Problem (2) is thus a joint portfolio allocation (over α) and statistical estimation (over β) problem, that minimizes the conditional expected loss of the portfolio, given that the outcome of X belongs to a set $\mathcal{N}$. To promote robust solutions and to hedge against the estimation error of $\mathbb{P}$, we will apply the distributionally robust framework. Instead of assessing the expected loss with respect to a single distribution, the distributionally robust formulation minimizes the expected conditional loss uniformly over a distributional ambiguity set that represents the portfolio managers' ambiguity regarding the underlying distribution $\mathbb{P}$.

We propose a data-driven portfolio optimization framework that *simultaneously exploits side information and robustifies over sampling error* and *model misspecification*. We formulate the distributionally robust portfolio optimization problem with side information as

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \ \sup_{\mathbb{Q} \in \mathbb{B}, \mathbb{Q}(X \in \mathcal{N}) \geq \varepsilon} \ \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}]. \tag{3}$$

A typical choice of the ambiguity set $\mathbb{B}$ is a neighborhood around a nominal distribution $\widehat{\mathbb{P}}$. In a data-driven setting, a popular choice of $\widehat{\mathbb{P}}$ is the empirical distribution, defined as a uniform distribution supported on the available data. As such, the data influences the optimization problem (3) via the channel of the ambiguity set prescription. The constraint $\mathbb{Q}(X \in \mathcal{N}) \geq \varepsilon$, in which ε is an input parameter, is added to (3) to control the

probability mass requirement on the set $\mathcal{N}$ and to avoid conditioning on sets of measure zero. In the limit, we also will study the case $\mathbb{Q}(X \in \mathcal{N}) > 0$ with strict inequality.

Intuitively, one can describe problem (3) as a zero-sum game between two players. The outer player is the statistical portfolio manager who optimizes over the allocation α and the estimator β in order to minimize the conditional expected loss, and the inner player can be regarded as a fictitious adversary whose goal is to increase the resulting loss by choosing an adversarial distribution.

Our formulation (3) has several benefits. First, our model is an end-to-end framework that directly generates investment decisions using data. This approach overcomes the sequential pipeline of feature extraction, return prediction and portfolio optimization, which is conventionally employed by portfolio managers. Therefore, our model, which consolidates return prediction and portfolio optimization, directly minimizes the conditional expected loss instead of the prediction error of return. Second, the robustness of the random return is directed by the portfolio choice. Indeed, the adversary who solves the supremum problem in (3) will aim to maximize the portfolio loss, and not the statistical loss. Thirdly, our proposed model does not require imposing any parametric family for prediction, and thus it reduces the risk of model misspecification.

1.1. **Main Contributions**

We study a robust portfolio optimization strategy that integrates conditional estimation and decision-making. This strategy leverages the side information to improve the ex-ante return prediction. Additionally, it also robustifies over estimation error in order to improve the out-of-sample performance of the portfolio allocation. Our contributions are summarized as follows.

(1) We present a comprehensive modelling framework for distributionally robust optimization for conditional decision-making, which provides flexibilities for selecting different conditional set and adjusting probability mass requirement. We study specifically the case where the set $\mathbb{B}$ will be prescribed using the notions of optimal transport, and we show that the qualitative behavior of the distributionally robust portfolio optimization problem with side information depends on whether the set $\mathcal{N}$ is a singleton, and whether the probability mass requirement ε is zero.

(2) We derive finite-dimensional reformulations[1] of the distributionally robust portfolio optimization problem with side information. Further, we also provide tractable reformulations for the distributionally robust mean-variance problem and mean-CVaR problem as convex conic (second-order cone or semi-definite cone) problems. This result is surprising given that the conditional expectation is a nonlinear function of the probability measure, and that minimizing the worst-case conditional expected loss is equivalent to a min-max optimization problem with a fractional objective function.

(3) We conduct extensive numerical experiments in both the US and Chinese equity markets. The results demonstrate the advantage of the distributionally robust portfolio optimization formulation, which has a higher Sharpe ratio and a better quantile performance for the portfolio return when compared with other benchmarks.

The distributionally robust optimization problem with side information as presented in (3) is generic, and our contributions are related to the vast literature in estimation-informed decision-making tasks. Moreover, by an appropriate choice of the loss function ℓ, the sets $(\mathcal{A}, \mathcal{B}, \mathcal{N})$ and the ambiguity set $\mathbb{B}$, the formulation (3) can be applied to many other problems in the field of operations research and management science such as supply chain management, transportation and energy planning.

1.2. **Related Literature**

We start this section by discussing existing work that exploits side information to enhance portfolio optimization. For applications of fundamental analysis, Fama and French leveraged market capitalization and price-earning ratio to construct the celebrated factor models [25, 26], which was later applied to portfolio

---

[1]To be precise, finite-dimensional optimization problems refer to problems that have a finite number of variables and constraints.

optimization [14]; Pástor and Stambaugh [49] reported the evidence of equity risk premium due to illiquidity. For applications of technical indicators, examples include momentum factor [32, 47], implied volatility [22], short interest [55], and investor sentiment [4]. However, when comparing with our proposed approach, these methods are vulnerable to estimation errors.

There is a rich literature on distributionally robust portfolio optimization, with diverse choices of distributional ambiguity set or different objective functions. When the objective function involves Value-at-Risk or CVaR, previous studies employed the moment-based ambiguity set [30, 65, 31] and Wasserstein ambiguity set [44] to model the worst-case risk. For distributionally robust portfolio optimizations that maximize piecewise affine loss functions, see [46] for moment-based ambiguity set, and [16] for event-wise ambiguity set. For the distributionally robust mean-variance problem, Lobo and Boyd [40] and Delage and Ye [19] are among the earliest to consider the worst-case mean variance model under moment ambiguity set. Other types of ambiguity sets, such as Wasserstein ambiguity set [51, 10], and optimal transport ambiguity set [13], are also studied afterwards. Due to the resemblance between Wasserstein distributionally robustification and norm regularization [11], the weight-constrained portfolio optimization problem [20] is also closely related to this stream of research aiming to boost the robustness of the optimal portfolio. Nevertheless, none of the aforementioned works considers how to incorporate side information and conditional estimation.

Our paper is also related to the topic of decision-making with side information, which is usually modelled by a conditional stochastic optimization (CSO) problem (2). Ban and Rudin [5] applied weighted sample average approximation (SAA) with weights learned from kernel regression to learn the optimal decision in inventory management with covariate information. Bertsimas and Kallus [7] applied weighted SAA to solve the general CSO problem, in which the weight is learned from nonparametric estimator such as k-nearest neighbors, kernel regressions, and random forests. Bertsimas and McCord [8] generalized the approach of [7] to multistage decision-making problems. Kannan et al. [35] developed a residual based SAA method, where average residuals during training of the learning model are added on to a point prediction of response, which is used as a response sample within the SAA. Decision tree based approaches are also developed to solve CSO. Athey et al. [3] proposed the forest-based estimates for local parameter estimation, generalizing the original random forest algorithm with observations of side information; Kallus and Mao [34] generalized the idea to solve stochastic optimization problems with side information, in which they seek to minimize the expected loss instead of the prediction accuracy. Conditional chance constrained programming was also studied recently [54].

Recently, distributionally robust optimization (DRO) formulations are integrated into the CSO problem to improve the out-of-sample performance of the solutions. For example, Kannan et al. [36] constructed a Wasserstein based uncertainty set over the empirical distribution of the residuals, robustifying their previous work [35]; Esteban-Pérez and Morales [24] leveraged some probability trimming methods and a partial mass transportation problem to model the distributionally robust CSO problem.

Despite some common characteristics in terms of methodology, such as the idea of using the theory of optimal transport to robustify the CSO problem, our paper is an independent contribution in comparison to [36] and [24].[2] There are two key differences between our model and that considered in [36]. Firstly, we integrate prediction and decision-making into a single optimization problem, while in [36] the prediction and the decision-making procedures are separated to simplify the problem. Secondly, the conditional side information is required to be a singleton set $X = x_0$ in [36], but our approach allows for a more flexible modelling of side information by conditioning on a more general event $X \in \mathcal{N}$. Our work is also distinct from Esteban-Pérez and Morales [24]. In contrast to [24], which employs probability trimmings to relax problem (3), we seek to provide an exact and tractable reformulation of the problem. Moreover, our approach tackles the optimal transport ambiguity set directly, and the results can be explained by picturing an adversary who perturbs the sample points in an intuitive manner.

---

[2]This claim is supported by the fact that this paper extends our previous work [48], which predates the publication dates of both [36] and [24].

Finally, this paper is a complete and comprehensive extension to our previous work [48], with a twofold improvement. On the one hand, this paper naturally unifies prediction and decision-making into a single DRO problem, while [48] solely considers the conditional estimation problem. On the other hand, this paper applies optimal transport distance to construct a more versatile class of distributional ambiguity set, whereas the ambiguity set in [48] is based on the type-∞ Wasserstein distance. Even though the type-∞ Wasserstein distance is also motivated by the theory of optimal transport, it behaves qualitatively different from a type-p ($p < \infty$) Wasserstein distance. The pathway to resolve the distributionally robust optimization problems in this paper hence differs significantly from the techniques employed in [48]. Finally, this paper also enriches the emerging field of Wasserstein DRO, which has gained momentum recently thanks to its applicability in a wide spectrum of practical problems [11, 29, 63, 23, 44, 39].

Notations. For any integer $M \geq 1$, we denote by $[M]$ the set of integers $\{1, \ldots, M\}$. For any $x \in \mathbb{R}$, we denote by $(x)^+ = \max\{x, 0\}$. We write $\mathbb{R}_+ \triangleq [0, \infty)$. Let $\|\cdot\|_p$ denote the l_p-norm for $p \in [1, \infty]$. For any metric space $(\mathcal{S}, d_{\mathcal{S}})$ equipped with a Borel sigma-algebra $\mathcal{F}_{\mathcal{S}}$, $\mathcal{M}(\mathcal{S})$ is the space of all probability measures defined on $(\mathcal{S}, \mathcal{F}_{\mathcal{S}})$. For any $s \in \mathcal{S}$, the Dirac's delta measure corresponding to s is denoted by δ_s, i.e., for all $E \in \mathcal{F}_{\mathcal{S}}$, $\delta_s(E) = 1$ if $s \in E$ and $\delta_s(E) = 0$ if $s \notin E$. For any subset $E \subseteq \mathcal{S}$, let ∂E denote the boundary of E, which is the closure of E minus the interior of E. The cone of $m \times m$ real positive semi-definite matrices is denoted by $\mathbb{S}_+^m$. For $m \times m$ real symmetric matrices A and B, we write $A \succeq B$ if and only if $A - B \in \mathbb{S}_+^m$. For a probability measure $\mathbb{P}$, let $\mathbb{E}_{\mathbb{P}}[\cdot]$ denote the expectation under measure $\mathbb{P}$.

2. Problem Setup

We delineate in this section the details on our robustification of the conditional portfolio allocation problem. To this end, we will construct the ambiguity set using the optimal transport cost.

Definition 2.1 (Optimal transport cost). *Let $\mathbb{D}$ be a nonnegative and continuous function on $\Xi \times \Xi$. The optimal transport cost between two distributions $\mathbb{Q}_1$ and $\mathbb{Q}_2$ supported on Ξ is defined as*

$$\mathbb{W}(\mathbb{Q}_1, \mathbb{Q}_2) \triangleq \min \left\{ \mathbb{E}_\pi[\mathbb{D}(\xi_1, \xi_2)] : \pi \in \Pi(\mathbb{Q}_1, \mathbb{Q}_2) \right\},$$

where $\Pi(\mathbb{Q}_1, \mathbb{Q}_2)$ is the set of all probability measures on $\Xi \times \Xi$ with marginals $\mathbb{Q}_1$ and $\mathbb{Q}_2$, respectively.

Intuitively, the optimal transport cost $\mathbb{W}(\mathbb{Q}_1, \mathbb{Q}_2)$ computes the cheapest cost of transporting the mass from an initial distribution $\mathbb{Q}_1$ to a target distribution $\mathbb{Q}_2$, given a cost function $\mathbb{D}(\xi_1, \xi_2)$ prescribing the cost of transporting a unit of mass from position ξ_1 to position ξ_2. The existence of an optimal joint distribution $\pi^\star$ that attains the minimal value in the definition is guaranteed by [62, Theorem 4.1]. The class of optimal transport cost is rich enough to encompass common probability metrics. In particular, if $\mathbb{D}$ is a metric on Ξ, then $\mathbb{W}$ coincides with the type-1 Wasserstein distance. In this paper, we do not restrict the ground transportation cost $\mathbb{D}$ to be a metric, and as a consequence, $\mathbb{W}$ is not necessarily a proper distance on the space of probability measure. Nevertheless, we will usually refer to $\mathbb{W}$ as an optimal transport distance with a slight abuse of terminology.

We now consider the covariate space $\mathcal{X}$, equipped with a ground cost $\mathbb{D}_{\mathcal{X}}$, and the return space $\mathcal{Y}$, equipped with a ground cost $\mathbb{D}_{\mathcal{Y}}$. The joint sample space $\mathcal{X} \times \mathcal{Y}$ endowed with a ground cost $\mathbb{D}$ constructed from $\mathbb{D}_{\mathcal{X}}$ and $\mathbb{D}_{\mathcal{Y}}$. We suppose that the portfolio manager possesses N historical data samples, where each sample is a pair of covariate-return $(\widehat{x}_i, \widehat{y}_i) \in \mathcal{X} \times \mathcal{Y}$. Using the optimal transport cost in Definition 2.1, we define the *joint ambiguity set* of the distributions of (X, Y) as

$$\mathbb{B}_\rho = \left\{ \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) : \mathbb{W}(\mathbb{Q}, \widehat{\mathbb{P}}) \leq \rho \right\}.$$

The set $\mathbb{B}_\rho$ contains all distributions that are of optimal transport distance less than or equal to ρ from the empirical distribution

$$\widehat{\mathbb{P}} \triangleq \frac{1}{N} \sum_{i \in [N]} \delta_{(\widehat{x}_i, \widehat{y}_i)},$$

defined as a uniform distribution supported on the training data $(\widehat{x}_i, \widehat{y}_i)_{i=1}^N$. In this form, $\mathbb{B}_\rho$ is a non-parametric ambiguity set. We propose to solve the distributionally robust conditional portfolio allocation problem

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]. \tag{4}$$

In (4), the portfolio manager optimizes over the allocation $\alpha \in \mathcal{A}$ and the estimate $\beta \in \mathcal{B}$, while the fictitious adversary chooses the distribution $\mathbb{Q}$. The side information input is specifically modelled as a neighborhood around a covariate $x_0 \in \mathcal{X}$, and the size of this neighborhood is controlled by a parameter $\gamma \in \mathbb{R}_+$. This neighborhood is prescribed using the ground cost $\mathbb{D}_\mathcal{X}$ on $\mathcal{X}$ as

$$\mathcal{N}_\gamma(x_0) \triangleq \{x \in \mathcal{X} : \mathbb{D}_\mathcal{X}(x, x_0) \leq \gamma\}.$$

The set $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$ is referred to as the fiber set. When $\mathcal{N}_\gamma(x_0)$ collapses into a singleton, it represents the event $X = x_0$, and we refer to this case as conditioning on a singular fiber set. In this perspective, our formulation is a generalization of the conventional conditional decision-making problem in the literature, which usually focus on conditioning on $X = x_0$. The constraint $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon$ indicates a minimum amount $\varepsilon \in [0, 1]$ of probability mass to be assigned to the fiber set $\mathcal{N}_\gamma(x_0)$. The feasible set $\mathcal{A}$ and $\mathcal{B}$ are assumed to be simple, in the sense that they are representable using second-order cone constraints. Overall, problem (4) involves three parameters: the ambiguity size $\rho \in \mathbb{R}_+$, the fiber size $\gamma \in \mathbb{R}_+$, and the probability mass requirement $\varepsilon \in [0, 1]$.

It is now instructive to discuss the difficulty level of problem (4). By the definition of the conditional expectation, problem (4) is equivalent to

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \frac{\mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) \mathbb{1}(X \in \mathcal{N}_\gamma(x_0))]}{\mathbb{E}_\mathbb{Q}[\mathbb{1}(X \in \mathcal{N}_\gamma(x_0))]}.$$

In this form, one can observe that the objective function is a fractional, and thus is a non-linear, function of the probability measure $\mathbb{Q}$. Problem (4) hence belongs to the class of distributionally robust fractional optimization problems [33, 64]. This type of problems is fundamentally different, and also fundamentally more difficult to solve, compared to existing models in the field of distributionally robust optimization that simply minimizes the worst-case expected loss [53].

It is important to emphasize that the auxiliary variable $\beta \in \mathcal{B}$ is intentionally regrouped to the feasible set of the outer minimization problem of formulation (4). This regrouping of portfolio allocation variables $\alpha \in \mathcal{A}$ and the estimation variable $\beta \in \mathcal{B}$ is a modelling choice. In the expected loss minimization perspective, this regrouping is without any loss of optimality under the conditions such as the compactness of $\mathcal{A}$ and $\mathcal{Y}$, as we illustrate in the following technical lemmas.

Lemma 2.2 (Mean-variance loss function). *The robustified conditional mean-variance portfolio allocation problem is*

$$\min_{\alpha \in \mathcal{A}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \mathrm{Variance}_\mathbb{Q}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)] - \eta \cdot \mathbb{E}_\mathbb{Q}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)].$$

If $\mathcal{A}$ and $\mathcal{Y}$ are compact, then the above optimization problem is equivalent to (4) with ℓ representing the mean-variance loss function of the form

$$\ell(y, \alpha, \beta) = (y^\top \alpha - \beta)^2 - \eta \cdot y^\top \alpha. \tag{5}$$

and the set $\mathcal{B}$ defined as

$$\mathcal{B} = \Big[\inf_{\alpha \in \mathcal{A}, \ y \in \mathcal{Y}} y^\top \alpha, \ \sup_{\alpha \in \mathcal{A}, \ y \in \mathcal{Y}} y^\top \alpha \Big].$$

Lemma 2.3 (Mean-CVaR loss function). *The robustified conditional mean-CVaR portfolio allocation problem with risk tolerance $\tau \in (0, 1)$ is*

$$\min_{\alpha \in \mathcal{A}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \mathrm{CVaR}_\mathbb{Q}^{1-\tau}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)] - \eta \cdot \mathbb{E}_\mathbb{Q}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)].$$

If $\mathcal{A}$ and $\mathcal{Y}$ are compact, the above optimization problem is equivalent to problem (4) *with ℓ representing the mean-CVaR loss function of the form*

$$\ell(y, \alpha, \beta) = -\eta y^\top \alpha + \beta + \frac{1}{\tau}(-y^\top \alpha - \beta)^+ = \max\left\{-\eta y^\top \alpha + \beta, -(\eta + \frac{1}{\tau})y^\top \alpha + (1 - \frac{1}{\tau})\beta\right\} \qquad (6)$$

and $\mathcal{B}$ as defined in Lemma 2.2.

Theoretically, the assumptions on the compactness of the sets $\mathcal{A}$ and $\mathcal{Y}$ may be restrictive. However, in practice and especially in the portfolio decision setting, the feasible allocation set $\mathcal{A}$ is usually compact. Moreover, we empirically observe that restricting $\mathcal{Y}$ to be inside a ball of sufficiently large diameter does not alter the numerical solution. Alternatively, weak duality also implies that problem (4) is a conservative formulation of the robustified risk measure minimization problem, and thus the optimal solution of (4) can still be considered as a robust, or more risk-averse, portfolio allocation. The proofs of Lemmas 2.2 and 2.3 are relegated to the appendix.

Throughout this paper, we make the following regularity assumption.

Assumption 2.4 (Regularity conditions). *The following assumptions hold.*

(i) *Separable ground cost: The joint space $\mathcal{X} \times \mathcal{Y}$ is endowed with a separable cost function $\mathbb{D}$ of the form*

$$\mathbb{D}\big((x, y), (x', y')\big) = \mathbb{D}_{\mathcal{X}}(x, x') + \mathbb{D}_{\mathcal{Y}}(y, y') \quad \forall (x, y), (x', y') \in \mathcal{X} \times \mathcal{Y}.$$

The individual ground transport costs $\mathbb{D}_{\mathcal{X}}$ and $\mathbb{D}_{\mathcal{Y}}$ are symmetric, non-negative and continuous on $\mathcal{X} \times \mathcal{X}$ and $\mathcal{Y} \times \mathcal{Y}$, respectively. Moreover, $\mathbb{D}\big((x, y), (x', y')\big) = 0$ if and only if $(x, y) = (x', y')$.

(ii) *Projection: For any $i \in [N]$, there exist a unique $\widehat{x}_i^p \in \mathcal{N}_\gamma(x_0)$ such that*

$$0 \leq \kappa_i \triangleq \min_{x \in \mathcal{N}_\gamma(x_0)} \mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i) = \mathbb{D}_{\mathcal{X}}(\widehat{x}_i^p, \widehat{x}_i). \qquad (7)$$

(iii) *Vicinity: For any $(x, y) \in \partial \mathcal{N}_\gamma(x_0) \times \mathcal{Y}$ and for any radius $r > 0$, the set*

$$\{(x', y') \in (\mathcal{X} \times \mathcal{Y}) \backslash (\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) : \mathbb{D}_{\mathcal{X}}(x', x) + \mathbb{D}_{\mathcal{Y}}(y', y) \leq r\}$$

is non-empty.

The conditions on $\mathbb{D}_{\mathcal{X}}$ and $\mathbb{D}_{\mathcal{Y}}$ in Assumption 2.4(i) are trivially satisfied if these individual ground costs are chosen as continuous functions of norms on $\mathcal{X}$ and $\mathcal{Y}$, respectively. Assumption 2.4(ii) asserts the existence of the projection of any training sample point $(\widehat{x}_i, \widehat{y}_i)$ onto the set $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$. It is easy to see that $\kappa_i = 0$ and $\widehat{x}_i^p = \widehat{x}_i$ whenever $\widehat{x}_i \in \mathcal{N}_\gamma(x_0)$ thanks to the choice of $\mathbb{D}$ in Assumption 2.4(i). Assumption 2.4(iii) indicates that any points on the boundary of the fiber set can be shifted outside the fiber with arbitrary small cost.

Notice that for a fixed amount of fiber probability ε, if the transportation budget ρ is small, the feasible set of the inner supremum problem in (4) may be empty. We define the minimum value of the radius ρ so that this feasible set is non-empty as

$$\rho_{\min}(x_0, \gamma, \varepsilon) \triangleq \inf\left\{W(\mathbb{Q}, \widehat{\mathbb{P}}) : \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}), \ \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon\right\}. \qquad (8)$$

The value κ_i defined in Assumption 2.4(ii) signifies the unit cost of moving a point mass from the observation $(\widehat{x}_i, \widehat{y}_i)$ to the fiber set $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$. The magnitude of κ_i depends on γ, however, this dependence is made implicit. Using this definition of κ, the next proposition asserts that the value of $\rho_{\min}(x_0, \gamma, \varepsilon)$ can be computed by solving a finite dimensional optimization problem.

Proposition 2.5 (Minimum radius). *The value $\rho_{\min}(x_0, \gamma, \varepsilon)$ equals to the optimal value of a linear program*

$$\rho_{\min}(x_0, \gamma, \varepsilon) = \min\left\{N^{-1} \sum_{i \in [N]} \kappa_i v_i : v \in [0, 1]^N, \ \sum_{i \in [N]} v_i \geq N\varepsilon\right\}, \qquad (9)$$

where κ are defined as in (7). *Furthermore, there exists a measure $\mathbb{Q} \in \mathbb{B}_\rho$ such that $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon$ if and only if $\rho \geq \rho_{\min}(x_0, \gamma, \varepsilon)$.*

Notice that the minimization problem (9) can be formulated as a fractional knapsack problem, and it can be solved in time $O(N \log N)$ using a greedy heuristics [18], see also [38, Proposition 17.1].

Proof of Proposition 2.5. Using the definition of the optimal transport cost, we can compute $\rho_{\min}(x_0, \varepsilon)$ as

$$\rho_{\min}(x_0, \gamma, \varepsilon) = \begin{cases} \inf & \int_{(\mathcal{X} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{Y})} [\mathbb{D}_{\mathcal{X}}(x, x') + \mathbb{D}_{\mathcal{Y}}(y, y')] \, \pi \big((\mathrm{d}x \times \mathrm{d}y), (\mathrm{d}x' \times \mathrm{d}y') \big) \\ \text{s.t.} & \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}), \ \pi \in \Pi(\mathbb{Q}, \widehat{\mathbb{P}}) \\ & \int_{(\mathcal{X} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{Y})} \mathbf{1}_{\mathcal{N}_\gamma(x_0)}(x) \, \pi \big((\mathrm{d}x \times \mathrm{d}y), (\mathrm{d}x' \times \mathrm{d}y') \big) \geq \varepsilon. \end{cases}$$

Because $\widehat{\mathbb{P}}$ is an empirical measure, any joint probability measure $\pi \in \Pi(\mathbb{Q}, \widehat{\mathbb{P}})$ can be written as $\pi = N^{-1} \sum_{i \in [N]} \pi_i \otimes \delta_{(\widehat{x}_i, \widehat{y}_i)}$ using the collection of probability measures $\{\pi_i\}_{i \in [N]}$, and $\otimes$ denotes the Kronecker product of two probability measures. One thus can reformulate $\rho_{\min}(x_0, \gamma, \varepsilon)$ as

$$\rho_{\min}(x_0, \gamma, \varepsilon) = \begin{cases} \inf & N^{-1} \sum_{i \in [N]} \int_{\mathcal{X} \times \mathcal{Y}} [\mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y, \widehat{y}_i)] \, \pi_i(\mathrm{d}x \times \mathrm{d}y) \\ \text{s.t.} & \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) \ \forall i \in [N], \ \sum_{i \in [N]} \pi_i(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq N\varepsilon. \end{cases} \tag{10}$$

Let $\{\pi_i^\star\}_{i \in [N]}$ be the optimal solution of the above optimization problem. We now show that $\pi_i^\star$ should be of the form

$$\pi_i^\star = \beta_i \delta_{(\widehat{x}_i^p, \widehat{y}_i)} + (1 - \beta_i) \delta_{(\widehat{x}_i, \widehat{y}_i)}$$

for some $\upsilon_i \in [0, 1]$ and where $\widehat{x}_i^p$ is the projection of $\widehat{x}_i$ onto $\mathcal{N}_\gamma(x_0)$ defined in (7). Suppose otherwise, then denote $P_i = \pi_i^\star(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})$. Consider now

$$\pi_i' = P_i \delta_{(\widehat{x}_i^p, \widehat{y}_i)} + (1 - P_i) \delta_{(\widehat{x}_i, \widehat{y}_i)}.$$

It is trivial that $\pi_i'(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq \sum_{i \in [N]} P_i = \sum_{i \in [N]} \pi_i^\star(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq N\varepsilon$. Moreover, we have

$$\sum_{i \in [N]} \int_{\mathcal{X} \times \mathcal{Y}} [\mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y, \widehat{y}_i)] \, \pi_i'(\mathrm{d}x \times \mathrm{d}y) = \sum_{i \in [N]} [\mathbb{D}_{\mathcal{X}}(\widehat{x}_i^p, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(\widehat{y}_i, \widehat{y}_i)] P_i$$

$$\leq \sum_{i \in [N]} \int_{\mathcal{X} \times \mathcal{Y}} [\mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y, \widehat{y}_i)] \, \pi_i^\star(\mathrm{d}x \times \mathrm{d}y),$$

which implies that $\{\pi_i'\}_{i \in [N]}$ is at least as good as $\{\pi_i^\star\}_{i \in [N]}$ in the optimization problem (10).

Next, by restricting the decision variables to $\pi_i = \upsilon_i \delta_{(\widehat{x}_i^p, \widehat{y}_i)} + (1 - \upsilon_i) \delta_{(\widehat{x}_i, \widehat{y}_i)}$, one now can consider the equivalent reformulation

$$\rho_{\min}(x_0, \gamma, \varepsilon) = \begin{cases} \inf & N^{-1} \sum_{i \in [N]} \kappa_i \upsilon_i \\ \text{s.t.} & \upsilon \in [0, 1]^N, \ \sum_{i \in [N]} \upsilon_i \geq N\varepsilon. \end{cases}$$

Because the feasible set is compact and the objective function is continuous, the minimization operator is justified thanks to Weierstrass' maximum value theorem [1, Corollary 2.35].

It remains to show the existence of a measure $\mathbb{Q}^\star \in \mathbb{B}_\rho$ with $\mathbb{Q}^\star(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon$ if and only if $\rho \geq \rho_{\min}(x_0, \gamma, \varepsilon)$. The definition of $\rho_{\min}(x_0, \gamma, \varepsilon)$ immediately proves the "only if" part, because (8) implies that

$$\mathbb{B}_\rho \cap \{\mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) : \ \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon\} = \emptyset \quad \forall \rho < \rho_{\min}(x_0, \gamma, \varepsilon).$$

Now we prove the "if" part. Given a minimizer $\{\upsilon_i^\star\}_{i \in [N]}$ that solves (9), there exists a collection of probability measures $\{\pi_i^\star\}_{i \in [N]}$ defined as $\pi_i^\star \triangleq \upsilon_i^\star \delta_{(\widehat{x}_i^p, \widehat{y}_i)} + (1 - \upsilon_i^\star) \delta_{(\widehat{x}_i, \widehat{y}_i)}$ such that

$$\int_{\mathcal{X} \times \mathcal{Y}} [\mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y, \widehat{y}_i)] \, \pi_i^\star(\mathrm{d}x \times \mathrm{d}y) = \kappa_i \upsilon_i^\star.$$

The probability measure $\mathbb{Q}^\star \triangleq N^{-1} \sum_{i \in [N]} \pi_i^\star$ satisfies

$$\mathbb{W}(\mathbb{Q}^\star, \widehat{\mathbb{P}}) \leq \frac{1}{N} \sum_{i \in [N]} \int_{\mathcal{X} \times \mathcal{Y}} [\mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y, \widehat{y}_i)] \, \pi_i^\star(\mathrm{d}x \times \mathrm{d}y) \leq \frac{1}{N} \sum_{i \in [N]} \kappa_i v_i^\star = \rho_{\min}(x_0, \gamma, \varepsilon) \leq \rho,$$

where the last equality is from the optimality of $v^\star$. This implies that $\mathbb{Q}^\star \in \mathbb{B}_\rho$ and completes the proof. $\square$

3. Tractable Reformulations for Singular Fiber Set

We study in this section the case where the radius of the fiber set is $\gamma = 0$, which implies that $\mathcal{N}_\gamma(x_0) = \{x_0\}$. In this case, we simply recover the conventional portfolio allocation problem conditional on $X = x_0$. A schematic representation of the sample space is presented in Figure 1. Interestingly, the qualitative behavior of the robustified conditional portfolio allocation problem (4) depends on whether $\varepsilon = 0$ or $\varepsilon > 0$. We will explore these two cases in the subsequent subsections.
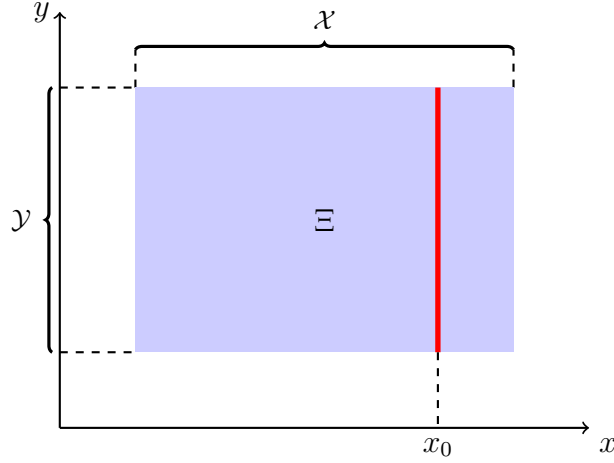

FIGURE 1. Generic setting for conditional decision-making with a singular fiber set. Shaded area represents the support $\Xi \in \mathcal{X} \times \mathcal{Y}$ of the joint covariate-return random variables. Bold red line indicates the fiber set $\{x_0\} \times \mathcal{Y}$.

3.1. Null fiber probability $\varepsilon = 0$

We consider now the situation when $\gamma = \varepsilon = 0$. Notice that the probability mass constraint in this case should be taken as a strict inequality of the form $\mathbb{Q}(X = x_0) > 0$ to avoid conditioning on sets of measure zero. This is equivalent to viewing the mass constraint in the limit as ε tends to zero. If we use the type-∞ Wasserstein distance to dictate the set $\mathbb{B}_\rho$, then the results from [48] can be utilized in order the compute the worst-case conditional expected loss. However, in this paper, we use the optimal transport of Definition 2.1 to prescribe $\mathbb{B}_\rho$, and the worst-case expected loss becomes uninformative as is shown in the following result.

Proposition 3.1 (Uninformative solution when $\varepsilon = \gamma = 0$). *Suppose that $\varepsilon = \gamma = 0$. For any $x_0 \in \mathcal{X}$ and $\rho \in \mathbb{R}_{++}$, the worst-case conditional expected loss becomes*

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X=x_0)>0} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X = x_0] = \sup_{y \in \mathcal{Y}} \ell(y, \alpha, \beta).$$

Proof of Proposition 3.1. Because $\widehat{\mathbb{P}}$ is an empirical measure, any joint probability measure $\pi \in \Pi(\mathbb{Q}, \widehat{\mathbb{P}})$ can be written as $\pi = \frac{1}{N} \sum_{i \in [N]} \pi_i \otimes \delta_{(\widehat{x}_i, \widehat{y}_i)}$, where each π_i is a probability measure on $\mathcal{X} \times \mathcal{Y}$. Thus by the

definition of the optimal transport cost, we find

$$\mathbb{B}_\rho = \left\{ \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) : \begin{array}{l} \exists \ \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) \ \forall i \in [N] \text{ such that } \mathbb{Q} = N^{-1} \sum_{i \in [N]} \pi_i \\ \frac{1}{N} \sum_{i \in [N]} \int_\Xi [\mathbb{D}_\mathcal{X}(x, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(y, \widehat{y}_i)] \pi_i(\mathrm{d}x \times \mathrm{d}y) \leq \rho \end{array} \right\}.$$

Fix any arbitrary $y_0 \in \mathcal{Y}$. For any $i \in [N]$, let $(x_i', y_i') \in \mathcal{X} \times \mathcal{Y}$ be such that $x_i' \neq x_0$ and

$$\mathbb{D}_\mathcal{X}(x_i', \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(y_i', \widehat{y}_i) \leq \frac{\rho}{2N}.$$

Consider the following set of probability measures $\{\pi_i\}_{i \in [N]}$ defined through

$$\forall i \in [N-1] : \quad \pi_i = \begin{cases} \delta_{(x_i', y_i')} & \text{if } i \in [N-1], \\ \upsilon \delta_{(x_0, y_0)} + (1-\upsilon)\delta_{(x_N', y_N')} & \text{if } i = N, \end{cases}$$

for some $\upsilon \in (0, 1]$ satisfying $\upsilon[\mathbb{D}_\mathcal{X}(x_0, \widehat{x}_N) + \mathbb{D}_\mathcal{Y}(y_0, \widehat{y}_N)] \leq \rho/(2N)$. Given this specific construction of $\{\pi_i\}_{i \in [N]}$, we can verify that

$$\frac{1}{N} \sum_{i \in [N]} \int_{\mathcal{X} \times \mathcal{Y}} [\mathbb{D}_\mathcal{X}(x, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(y, \widehat{y}_i)] \pi_i(\mathrm{d}x \times \mathrm{d}y) \leq \rho, \quad \text{and} \quad \sum_{i \in [N]} \pi_i(X = x_0) = \upsilon > 0.$$

As such, the measure $\mathbb{Q}' = N^{-1} \sum_{i \in [N]} \pi_i$ satisfies $\mathbb{Q}' \in \mathbb{B}_\rho$ and $\mathbb{Q}'(X = x_0) > 0$. We thus have

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X=x_0)>0} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta)|X = x_0] \geq \mathbb{E}_{\mathbb{Q}'}[\ell(Y, \alpha, \beta)|X = x_0] = \ell(y_0, \alpha, \beta).$$

Because the choice of y_0 is arbitrary, we find

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X=x_0)>0} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta)|X = x_0] \geq \sup_{y \in \mathcal{Y}} \ell(y, \alpha, \beta).$$

Moreover, because the distribution of Y given $X = x_0$ is supported on $\mathcal{Y}$, we have

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X=x_0)>0} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta)|X = x_0] \leq \sup_{\mathbb{Q} \in \mathbb{B}_\rho} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta)|X = x_0] \leq \sup_{y \in \mathcal{Y}} \ell(y, \alpha, \beta).$$

This observation establishes the postulated equality and completes the proof. □

The result in Proposition 3.1 can be justified heuristically as follows. Consider an adversary who can move sample points on $\mathcal{X} \times \mathcal{Y}$ to maximize the loss subject to the optimal transport distance budget constraint. If $\widehat{x}_i = x_0$, then the adversary can slightly perturb $\widehat{x}_i$ within a infinitesimal distance so that the sample no longer belongs to the fiber set. Because of the continuous assumption on the ground metric, this perturbation cost an infinitesimally small amount of energy. The adversary can repeat until there is no sample point lying on the fiber set. Now, the adversary can pick any sample and slice out a tiny amount of mass and move that slice to any location on $\{x_0\} \times \mathcal{Y}$. Because the ground metric is continuous by Assumption 2.4(i) and because the slice can be chosen arbitrarily small, this would cost an infinitesimally small amount of energy. As the thin slice can be put on any point on $\{x_0\} \times \mathcal{Y}$, the resulting distribution will generate the robust conditional expected loss.

Proposition 3.1 reveals that the conditioning problem with $\varepsilon = \gamma = 0$ results in a robust optimization formulation, which can be overly conservative. This result is also negative in the sense that the worst-case conditional expected loss depends only on the support $\mathcal{Y}$, and it does *not* depend on the data $(\widehat{x}_i, \widehat{y}_i)$ that were collected. Thus, in this case, the notion of data-driven decision-making becomes obsolete, and thus we do not pursue the reformulation any further.

3.2. Strictly positive fiber probability $\varepsilon > 0$

We now study the situation with a singular fiber set $\gamma = 0$, but the probability mass requirement is set to a strict positive value $\varepsilon \in (0, 1]$. The robust conditional portfolio allocation problem (4) can now be rewritten

explicitly as

$$\min_{\alpha\in\mathcal{A},\ \beta\in\mathcal{B}}\ \sup_{\mathbb{Q}\in\mathbb{B}_\rho,\mathbb{Q}(X=x_0)\geq\varepsilon}\ \mathbb{E}_\mathbb{Q}[\ell(Y,\alpha,\beta)|X=x_0]. \tag{11}$$

The next theorem asserts that the worst-case conditional expected loss admits a finite-dimensional reformulation.

Theorem 3.2 (Worst-case conditional expected loss when $\gamma=0$). *Suppose that $\varepsilon\in(0,1]$ and $\rho>\rho_{\min}(x_0,0,\varepsilon)$.* *For any feasible solution (α,β), we have*

$$\sup_{\mathbb{Q}\in\mathbb{B}_\rho,\mathbb{Q}(X=x_0)\geq\varepsilon}\mathbb{E}_\mathbb{Q}[\ell(Y,\alpha,\beta)|X=x_0]$$

$$=\inf_{\substack{\lambda_1\in\mathbb{R}_+\\\lambda_2\in\mathbb{R}}}\left\{\rho\lambda_1+\varepsilon\lambda_2+\frac{1}{N}\sum_{i\in[N]}\left(\sup_{y_i\in\mathcal{Y}}\left\{\varepsilon^{-1}\ell(y_i,\alpha,\beta)-[\mathbb{D}_\mathcal{X}(x_0,\widehat{x}_i)+\mathbb{D}_\mathcal{Y}(y_i,\widehat{y}_i)]\lambda_1-\lambda_2\right\}\right)^+\right\}.$$

Proof of Theorem 3.2. By applying Proposition A.4, we have

$$\sup_{\mathbb{Q}\in\mathbb{B}_\rho,\mathbb{Q}(X=x_0)\geq\varepsilon}\mathbb{E}_\mathbb{Q}[\ell(Y,\alpha,\beta)|X=x_0]$$

$$=\left\{\begin{array}{ll}\displaystyle\sup_{\upsilon\in[0,1]^N,\sum_{i\in[N]}\upsilon_i=N\varepsilon}&\displaystyle\sup\quad(N\varepsilon)^{-1}\sum_{i\in[N]}\upsilon_i\mathbb{E}_{\mu_0^i}[\ell(Y,\alpha,\beta)]\\&\text{s.t.}\quad\mu_0^i\in\mathcal{M}(\mathcal{Y})\ \forall i\in[N]\\&\qquad\displaystyle\sum_{i\in[N]}\upsilon_i\big(\mathbb{D}_\mathcal{X}(x_0,\widehat{x}_i)+\mathbb{E}_{\mu_0^i}[\mathbb{D}_\mathcal{Y}(Y,\widehat{y}_i)]\big)\leq N\rho.\end{array}\right.$$

For any υ satisfying $\sum_{i\in[N]}\upsilon_i\kappa_i>N\rho$, the inner supremum subproblem is infeasible, thus, without loss of optimality, we can add the constraint $\sum_{i\in[N]}\upsilon_i\kappa_i\leq N\rho$ into the outer supremum. Denote temporarily by $\mathcal{U}$ the set

$$\mathcal{U}=\left\{\upsilon\in[0,1]^N:\sum_{i\in[N]}\upsilon_i=N\varepsilon,\ \sum_{i\in[N]}\upsilon_i\kappa_i\leq N\rho\right\}.$$

For any $\upsilon\in\mathcal{U}$, strong duality holds because $\mu_0^i:=\delta_{\widehat{y}_i}$ constitutes a Slater point. The duality result for moment problem [58, Proposition 3.4] implies that the inner supremum problem is equivalent to

$$\begin{array}{ll}\inf&\lambda_1\big(N\rho-\displaystyle\sum_{i\in[N]}\upsilon_i\mathbb{D}_\mathcal{X}(x_0,\widehat{x}_i)\big)+\displaystyle\sum_{i\in[N]}\theta_i\\\text{s.t.}&\lambda_1\in\mathbb{R}_+,\ \theta\in\mathbb{R}^N\\&\upsilon_i\mathbb{D}_\mathcal{Y}(y_i,\widehat{y}_i)\lambda_1+\theta_i\geq(N\varepsilon)^{-1}\upsilon_i\ell(y_i,\alpha,\beta)\quad\forall y_i\in\mathcal{Y},\ \forall i\in[N].\end{array}$$

By rescaling $\theta_i\leftarrow\upsilon_i\theta_i$, we have the equivalent form

$$\sup_{\mathbb{Q}\in\mathbb{B}_\rho,\mathbb{Q}(X=x_0)\geq\varepsilon}\mathbb{E}_\mathbb{Q}[\ell(Y,\alpha,\beta)|X=x_0]$$

$$=\left\{\begin{array}{ll}\displaystyle\sup_{\upsilon\in[0,1]^N,\sum_{i\in[N]}\upsilon_i=N\varepsilon}&\inf\quad\lambda_1\big(N\rho-\displaystyle\sum_{i\in[N]}\upsilon_i\mathbb{D}_\mathcal{X}(x_0,\widehat{x}_i)\big)+\displaystyle\sum_{i\in[N]}\upsilon_i\theta_i\\&\text{s.t.}\quad\lambda_1\in\mathbb{R}_+,\ \theta\in\mathbb{R}^N\\&\qquad\mathbb{D}_\mathcal{Y}(y_i,\widehat{y}_i)\lambda_1+\theta_i\geq(N\varepsilon)^{-1}\ell(y_i,\alpha,\beta)\quad\forall y_i\in\mathcal{Y},\ \forall i\in[N]\end{array}\right.$$

$$=\left\{\begin{array}{ll}\inf&\displaystyle\sup_{\upsilon\in[0,1]^N,\sum_{i\in[N]}\upsilon_i=N\varepsilon}\lambda_1\big(N\rho-\displaystyle\sum_{i\in[N]}\upsilon_i\mathbb{D}_\mathcal{X}(x_0,\widehat{x}_i)\big)+\displaystyle\sum_{i\in[N]}\upsilon_i\theta_i\\\text{s.t.}&\lambda_1\in\mathbb{R}_+,\ \theta\in\mathbb{R}^N\\&\mathbb{D}_\mathcal{Y}(y_i,\widehat{y}_i)\lambda_1+\theta_i\geq(N\varepsilon)^{-1}\ell(y_i,\alpha,\beta)\quad\forall y_i\in\mathcal{Y},\ \forall i\in[N].\end{array}\right.$$

$$= \begin{cases} \inf & N\rho\lambda_1 + N\varepsilon\lambda_2 + \sum_{i\in[N]} \vartheta_i \\ \text{s.t.} & \lambda_1 \in \mathbb{R}_+, \ \lambda_2 \in \mathbb{R}, \ \theta \in \mathbb{R}^N, \ \vartheta \in \mathbb{R}^N_+ \\ & \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\lambda_1 + \nu_i \geq (N\varepsilon)^{-1}\ell(y_i, \alpha, \beta) \quad \forall y_i \in \mathcal{Y}, \ \forall i \in [N] \\ & \lambda_2 + \vartheta_i \geq \theta_i - \mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i)\lambda_1 \qquad \forall i \in [N], \end{cases}$$

where the second equality follows from interchanging the min-max operators using Sion's minimax theorem [60]. By eliminating θ, we have

$$\sup_{\mathbb{Q}\in\mathbb{B}_\rho, \mathbb{Q}(X=x_0)\geq\varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X = x_0]$$

$$= \begin{cases} \inf & N\rho\lambda_1 + N\varepsilon\lambda_2 + \sum_{i\in[N]} \vartheta_i \\ \text{s.t.} & \lambda_1 \in \mathbb{R}_+, \ \lambda_2 \in \mathbb{R}, \ \vartheta \in \mathbb{R}^N_+ \\ & \vartheta_i \geq \sup_{y_i\in\mathcal{Y}} \left\{ (N\varepsilon)^{-1}\ell(y_i, \alpha, \beta) - [\mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)]\lambda_1 - \lambda_2 \right\} \quad \forall i \in [N], \end{cases}$$

$$= \begin{cases} \inf & \rho\lambda_1 + \varepsilon\lambda_2 + \dfrac{1}{N}\sum_{i\in[N]} \vartheta_i \\ \text{s.t.} & \lambda_1 \in \mathbb{R}_+, \ \lambda_2 \in \mathbb{R}, \ \vartheta \in \mathbb{R}^N_+ \\ & \vartheta_i \geq \sup_{y_i\in\mathcal{Y}} \left\{ \varepsilon^{-1}\ell(y_i, \alpha, \beta) - [\mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)]\lambda_1 - \lambda_2 \right\} \quad \forall i \in [N], \end{cases}$$

where the second equality follows by rescaling the dual variables. Eliminating ϑ leads to the desired result. $\square$

It is instructive to discuss the insights that leads to the result presented in Theorem 3.2. Notice that the supremum problem in the statement of Theorem 3.2 has a fractional objective function. The first step of the proof establishes that without any loss of optimality, the *inequality* constraint $\mathbb{Q}(X = x_0) \geq \varepsilon$ can be reduced to an *equality* constraint of the form $\mathbb{Q}(X = x_0) = \varepsilon$ (see Proposition A.4 for a formal statement of this result). Leveraging this result, it suffices to solve an equivalent problem which is linear in the probability measure of the form

$$\sup_{\mathbb{Q}\in\mathbb{B}_\rho, \mathbb{Q}(X=x_0)=\varepsilon} \frac{1}{\varepsilon}\mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)\mathbb{1}(X = x_0)].$$

At this point, standard duality techniques can be applied to reformulate the original problem into a finite-dimensional optimization problem.

We acknowledge that a similar duality result has been proposed in [24, Theorem 1] using the trimming approach. Notice that the trimming procedure in [24] is rather restrictive: it was designed so that any feasible distribution $\mathbb{Q}$ should satisfy $\mathbb{Q}(X = x_0) = \varepsilon$, and thus the fractional objective function becomes a linear function of $\mathbb{Q}$. In direct comparison, our approach can be considered to be less stringent: we only impose an inequality $\mathbb{Q}(X = x_0) \geq \varepsilon$, and thus our adversary has a bigger feasible set and is more powerful.

By combining the infimum reformulation of Theorem 3.2 with the outer infimum problem of problem (11), the portfolio allocation problem (11) is thus reformulatable as a finite-dimensional optimization problem. More specifically, problem (11) becomes

$$\begin{aligned} \inf \quad & \rho\lambda_1 + \varepsilon\lambda_2 + \frac{1}{N}\sum_{i\in[N]} \theta_i \\ \text{s.t.} \quad & \alpha \in \mathcal{A}, \ \beta \in \mathcal{B}, \ \lambda_1 \in \mathbb{R}_+, \ \lambda_2 \in \mathbb{R}, \ \theta \in \mathbb{R}^N_+ \\ & \theta_i \geq \sup_{y_i\in\mathcal{Y}} \left\{ \varepsilon^{-1}\ell(y_i, \alpha, \beta) - \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\lambda_1 \right\} - \mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i)\lambda_1 - \lambda_2 \quad \forall i \in [N]. \end{aligned} \tag{12}$$

In Propositions 3.3 and 3.4, we provide a second-order cone reformulation of problem (12) tailored for the mean-variance and mean-CVaR objective functions for a special instance with $\mathcal{X} = \mathbb{R}^n$, $\mathcal{Y} = \mathbb{R}^m$, $\mathbb{D}_{\mathcal{X}}(x, \widehat{x}) = \|x - \widehat{x}\|^2$ and $\mathbb{D}_{\mathcal{Y}}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$. Notice that $\mathbb{D}_{\mathcal{X}}$ is constructed from an arbitrary norm on $\mathbb{R}^n$ while $\mathbb{D}_{\mathcal{Y}}$ is constructed as the *squared* Euclidean norm on $\mathbb{R}^m$.

Proposition 3.3 (Mean-variance loss function)**.** *Suppose that ℓ is the mean-variance loss function of the form (5), $\gamma = 0$, $\varepsilon \in (0, 1]$ and $\rho > \rho_{\min}(x_0, 0, \varepsilon)$. Suppose in addition that $\mathcal{X} = \mathbb{R}^n$, $\mathcal{Y}_0 = \mathbb{R}^m$, $\mathbb{D}_{\mathcal{X}}(x, \widehat{x}) = \|x-$*

$\widehat{x}\|^2$ and $\mathbb{D}_{\mathcal{Y}}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$. *The distributionally robust portfolio allocation model with side information* (11) *is equivalent to the second-order cone program*

$$\inf \quad \rho\lambda_1 + \varepsilon\lambda_2 + \frac{1}{N}\sum_{i\in[N]}\theta_i$$

$$\text{s.t.} \quad \alpha \in \mathcal{A}, \ \beta \in \mathcal{B}, \ \lambda_1 \in \mathbb{R}_+, \ \lambda_2 \in \mathbb{R}, \ \theta \in \mathbb{R}_+^N, \ z \in \mathbb{R}_+^N, \ w \in \mathbb{R}$$

$$0 \le w \le 1$$

$$\left\| \begin{bmatrix} 2\alpha \\ 1 - w - \varepsilon\lambda_1 \end{bmatrix} \right\|_2 \le 1 - w + \varepsilon\lambda_1$$

$$\left.\begin{array}{l} z_i = \varepsilon\theta_i + \varepsilon\|x_0 - \widehat{x}_i\|^2\lambda_1 + \varepsilon\lambda_2 + \frac{1}{4}\eta^2 + \eta\beta \\ \left\| \begin{bmatrix} 2\widehat{y}_i^\top\alpha - 2\beta - \eta \\ z_i - w \end{bmatrix} \right\|_2 \le z_i + w \end{array}\right\} \quad \forall i \in [N].$$

Proof of Proposition 3.3. By exploiting the quadratic form of ℓ, the supremum problem in the last constraint of (12) is a quadratic optimization problem that admits a closed form expression as

$$\sup_{y_i \in \mathcal{Y}} \left\{ \varepsilon^{-1}\ell(y_i, \alpha, \beta) - \lambda_1 \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i) \right\}$$

$$= \sup_{y_i \in \mathcal{Y}} \left\{ \varepsilon^{-1}(y_i^\top\alpha - \beta)^2 - \varepsilon^{-1}\eta \cdot y_i^\top\alpha - \lambda_1 \cdot \|y_i - \widehat{y}_i\|_2^2 \right\}$$

$$= -\frac{1}{4}\varepsilon^{-1}\eta^2 - \varepsilon^{-1}\eta\beta + \sup_{y_i \in \mathcal{Y}} \left\{ \varepsilon^{-1}(y_i^\top\alpha - \beta - \frac{1}{2}\eta)^2 - \lambda_1 \cdot \|y_i - \widehat{y}_i\|_2^2 \right\}$$

Let $\Delta_i = y_i^\top\alpha - \widehat{y}_i^\top\alpha$, then we have

$$\sup_{y_i \in \mathcal{Y}} \left\{ \varepsilon^{-1}\ell(y_i, \alpha, \beta) - \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\lambda_1 \right\}$$

$$= -\frac{1}{4}\varepsilon^{-1}\eta^2 - \varepsilon^{-1}\eta\beta + \sup_{\Delta_i \in \mathbb{R}} \left\{ \varepsilon^{-1}(\Delta_i + \widehat{y}_i^\top\alpha - \beta - \frac{1}{2}\eta)^2 - \lambda_1 \cdot \inf_{y_i \in \mathbb{R}^m : y_i^\top\alpha - \widehat{y}_i^\top\alpha = \Delta_i} \|y_i - \widehat{y}_i\|_2^2 \right\}$$

$$= -\frac{1}{4}\varepsilon^{-1}\eta^2 - \varepsilon^{-1}\eta\beta + \sup_{\Delta_i \in \mathbb{R}} \left\{ \varepsilon^{-1}(\Delta_i + \widehat{y}_i^\top\alpha - \beta - \frac{1}{2}\eta)^2 - \lambda_1\Delta_i^2\|\alpha\|_2^{-2} \right\},$$

where the last equality is obtained by applying the Cauchy-Schwarz inequality, i.e., $\Delta_i = (y_i - \widehat{y}_i)^T\alpha \le \|y_i - \widehat{y}_i\|_2\|\alpha\|_2$, which implies that the minimum for $\|y_i - \widehat{y}_i\|_2^2$ is $\Delta_i\|\alpha\|_2^{-2}$. By combining cases, we find

$$\sup_{y_i \in \mathcal{Y}} \left\{ \varepsilon^{-1}\ell(y_i, \alpha, \beta) - \lambda_1 \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i) \right\}$$

$$= \begin{cases} \varepsilon^{-1} \cdot \frac{(\widehat{y}_i^\top\alpha - \beta - \frac{1}{2}\eta)^2}{1 - \varepsilon^{-1}\|\alpha\|_2^2/\lambda_1} - \frac{1}{4}\varepsilon^{-1}\eta^2 - \varepsilon^{-1}\eta\beta & \text{if } \lambda_1 > \varepsilon^{-1}\|\alpha\|_2^2, \\ -\frac{1}{4}\varepsilon^{-1}\eta^2 - \varepsilon^{-1}\eta\beta & \text{if } \lambda_1 = \varepsilon^{-1}\|\alpha\|_2^2 \text{ and } \widehat{y}_i^\top\alpha = \beta + \eta/2, \\ +\infty & \text{otherwise.} \end{cases}$$

Consider the case of $\lambda_1 > \varepsilon^{-1}\|\alpha\|^2$ first, in which the last constraint of (12) is satisfied if and only if there exist ancillary variables $w \in (0, 1)$ and $z \in \mathbb{R}_+^N$, such that

$$1 - \varepsilon^{-1}\|\alpha\|_2^2/\lambda_1 \ge w \quad \Longleftrightarrow \quad \varepsilon(1 - w)\lambda_1 \ge \|\alpha\|_2^2$$

and

$$\theta_i \ge \varepsilon^{-1} \cdot \frac{(\widehat{y}_i^\top\alpha - \beta - \frac{1}{2}\eta)^2}{1 - \varepsilon^{-1}\|\alpha\|_2^2/\lambda_1} - \frac{1}{4}\varepsilon^{-1}\eta^2 - \varepsilon^{-1}\eta\beta \Longleftrightarrow \begin{cases} z_i = \varepsilon\theta_i + \varepsilon\|x_0 - \widehat{x}_i\|^2\lambda_1 + \varepsilon\lambda_2 + \frac{1}{4}\eta^2 + \eta\beta \\ z_i w \ge \left(\widehat{y}_i^\top\alpha - \beta - \frac{1}{2}\eta\right)^2 \end{cases} \quad \forall i \in [N].$$

Using the equivalent reformulation between hyperbolic constraint and second-order cone constraint [41], we have

$$\varepsilon(1 - w)\lambda_1 \ge \|\alpha\|_2^2 \Longleftrightarrow \left\| \begin{bmatrix} 2\alpha \\ 1 - w - \varepsilon\lambda_1 \end{bmatrix} \right\|_2 \le 1 - w + \varepsilon\lambda_1$$

and

$$z_i w \geq \left(\widehat{y}_i^\top \alpha - \beta - \frac{1}{2}\eta \right)^2 \iff \left\| \begin{bmatrix} 2\widehat{y}_i^\top \alpha - 2\beta - \eta \\ z_i - w \end{bmatrix} \right\|_2 \leq z_i + w \quad \forall i \in [N].$$

Now we consider the case of $\lambda_1 = \varepsilon^{-1}\|\alpha\|_2^2$, where the last constraint of (12) is equivalent to the cone constraints when $w = 0$. Finally notice that $w = 1$ recovers the original constraints when $\alpha = 0$. Combining all of the above cases and using them to replace the last constraint of (12) completes the proof. $\quad\square$

Proposition 3.4 (Mean-CVaR loss function). *Suppose that ℓ is the mean-CVaR loss function of the form* (6), *$\gamma = 0$, $\varepsilon \in (0,1]$ and $\rho > \rho_{\min}(x_0, 0, \varepsilon)$. Suppose in addition that $\mathcal{X} = \mathbb{R}^n$, $\mathcal{Y} = \mathbb{R}^m$, $\mathbb{D}_{\mathcal{X}}(x, \widehat{x}) = \|x - \widehat{x}\|^2$ and $\mathbb{D}_{\mathcal{Y}}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$. The distributionally robust portfolio allocation model with side information* (11) *is equivalent to the second-order cone program*

$$\inf \quad \rho\lambda_1 + \varepsilon\lambda_2 + \frac{1}{N}\sum_{i\in[N]} \theta_i$$

$$\text{s.t.} \quad \alpha \in \mathcal{A}, \ \beta \in \mathcal{B}, \ \lambda_1 \in \mathbb{R}_+, \ \lambda_2 \in \mathbb{R}, \ \theta \in \mathbb{R}_+^N, \ z \in \mathbb{R}_+^N, \ \tilde{z} \in \mathbb{R}_+^N$$

$$\left. \begin{array}{l} z_i = \theta_i + \lambda_1\|x_0 - \widehat{x}_i\|^2 + \lambda_2 + \varepsilon^{-1}\eta\widehat{y}_i^\top\alpha - \varepsilon^{-1}\beta \\[2mm] \tilde{z}_i = \theta_i + \lambda_1\|x_0 - \widehat{x}_i\|^2 + \lambda_2 + \varepsilon^{-1}(\eta + \frac{1}{\tau})\widehat{y}_i^\top\alpha - \varepsilon^{-1}(1 - \frac{1}{\tau})\beta \\[2mm] \left\| \begin{bmatrix} \varepsilon^{-1}\eta\alpha \\ z_i - \lambda_1 \end{bmatrix} \right\|_2 \leq z_i + \lambda_1, \quad \left\| \begin{bmatrix} \varepsilon^{-1}(\eta + \tau^{-1})\alpha \\ \tilde{z}_i - \lambda_1 \end{bmatrix} \right\|_2 \leq \tilde{z}_i + \lambda_1 \end{array} \right\} \forall i \in [N].$$

Proof of Proposition 3.4. Notice that ℓ is a pointwise maximum of two linear functions of y_i. In this case, the supremum problem in the last constraint of (12) can be written as

$$\sup_{y_i \in \mathcal{Y}} \left\{ \varepsilon^{-1}\ell(y_i, \alpha, \beta) - \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\lambda_1 \right\}$$

$$= \max \left\{ \begin{array}{l} \sup_{y_i \in \mathbb{R}^m} \left\{ -\varepsilon^{-1}\eta y_i^\top\alpha + \varepsilon^{-1}\beta - \lambda_1 \cdot \|y_i - \widehat{y}_i\|_2^2 \right\}, \\[2mm] \sup_{y_i \in \mathbb{R}^m} \left\{ -\varepsilon^{-1}(\eta + \frac{1}{\tau})y_i^\top\alpha + \varepsilon^{-1}(1 - \frac{1}{\tau})\beta - \lambda_1 \cdot \|y_i - \widehat{y}_i\|_2^2 \right\} \end{array} \right\}$$

$$= \max \left\{ -\varepsilon^{-1}\eta\widehat{y}_i^\top\alpha + \varepsilon^{-1}\beta + \frac{\varepsilon^{-2}\eta^2}{4\lambda_1}\|\alpha\|_2^2, \ -\varepsilon^{-1}(\eta + \frac{1}{\tau})\widehat{y}_i^\top\alpha + \varepsilon^{-1}(1 - \frac{1}{\tau})\beta + \frac{\varepsilon^{-2}(\eta + \tau^{-1})^2}{4\lambda_1}\|\alpha\|_2^2 \right\}.$$

Therefore, using additional variables $z_i \in \mathbb{R}_+$ and $\tilde{z}_i \in \mathbb{R}_+$, we can reformulate the last constraint of (12) as the following set of constraints

$$\left. \begin{array}{l} z_i = \theta_i + \lambda_1\|x_0 - \widehat{x}_i\|^2 + \lambda_2 + \varepsilon^{-1}\eta\widehat{y}_i^\top\alpha - \varepsilon^{-1}\beta \\[1mm] 4\lambda_1 z_i \geq \varepsilon^{-2}\eta^2\|\alpha\|_2^2 \\[1mm] \tilde{z}_i = \theta_i + \lambda_1\|x_0 - \widehat{x}_i\|^2 + \lambda_2 + \varepsilon^{-1}(\eta + \frac{1}{\tau})\widehat{y}_i^\top\alpha - \varepsilon^{-1}(1 - \frac{1}{\tau})\beta \\[1mm] 4\lambda_1\tilde{z}_i \geq \varepsilon^{-2}(\eta + \tau^{-1})^2\|\alpha\|_2^2 \end{array} \right\} \quad \forall i \in [N].$$

Using the equivalent reformulation between hyperbolic constraint and second-order cone constraint [41], we have for each $i \in [N]$

$$4\lambda_1 z_i \geq \varepsilon^{-2}\eta^2\|\alpha\|_2^2 \iff \left\| \begin{bmatrix} \varepsilon^{-1}\eta\alpha \\ z_i - \lambda_1 \end{bmatrix} \right\|_2 \leq z_i + \lambda_1,$$

and

$$4\lambda_1\tilde{z}_i \geq \varepsilon^{-2}(\eta + \tau^{-1})^2\|\alpha\|_2^2 \iff \left\| \begin{bmatrix} \varepsilon^{-1}(\eta + \tau^{-1})\alpha \\ \tilde{z}_i - \lambda_1 \end{bmatrix} \right\|_2 \leq \tilde{z}_i + \lambda_1,$$

which finishes the proof. $\quad\square$

Both Proposition 3.3 and 3.4 leverage the fact that $\mathcal{Y} = \mathbb{R}^m$ in order to simplify the semi-infinite constraints into second-order cone constraints. In case the set $\mathcal{Y}$ is an ellipsoid of the form $\mathcal{Y} = \{y \in \mathbb{R}^m : y^\top Q y + 2q^\top y +$

$q_0 \leq 0\}$ with a non-empty interior, then semi-definite cone constraint counterparts are also available by employing the S-lemma [52].

4. TRACTABLE REFORMULATIONS FOR NONSINGULAR CONDITIONING SET

We now focus on the portfolio allocation conditional on $X \in \mathcal{N}_\gamma(x_0)$ for some fiber radius $\gamma > 0$. A fiber set of this form was first used in [48] in the conditional estimation setting with the aim to hedge against noisy covariate information x_0 and also to improve the statistical performance of the solution approach. The results from [48] rely heavily from the specification of the ambiguity set using the type-∞ Wasserstein distance, and these results are not transferrable to the ambiguity set under investigation in this paper.

As a parallel counterpart to Section 3, we will study two separate cases depending on whether the probability requirement ε is zero or strictly positive. We start by discussing the behavior of the optimization problem when $\varepsilon = 0$.

4.1. **Null fiber probability** $\varepsilon = 0$

Similar to Section 3.1, we consider the probability mass constraint of the form $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) > 0$ with a strict inequality to avoid conditioning on sets of measure zero. The complication of conditioning on sets of measure zero, as highlighted in Section 3.1 for $\varepsilon = 0$, still arises in this case when the radius ρ is big enough. To illustrate this problem, we define the following quantity

$$\rho_{\max}(x_0, \gamma) \triangleq \inf\{\mathrm{W}(\mathbb{Q}, \widehat{\mathbb{P}}) : \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}), \ \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) = 0\}.$$

Intuitively, $\rho_{\max}(x_0, \gamma)$ indicates the minimum budget required to transport all the training samples out of the fiber $\mathcal{N}_\gamma(x_0)$. If the empirical distribution $\widehat{\mathbb{P}}$ satisfies $\widehat{\mathbb{P}}(X \in \mathcal{N}_\gamma(x_0)) = 0$, which means that there is no training samples $(\widehat{x}_i, \widehat{y}_i)$ falling inside the fiber set $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$, then it is trivial that $\rho_{\max}(x_0, \gamma) = 0$. If $\widehat{\mathbb{P}}(X \in \mathcal{N}_\gamma(x_0)) > 0$, then the value of $\rho_{\max}(x_0, \gamma)$ is known in closed form. To this end, define the following index sets

$$\mathcal{I}_1 = \{i \in [N] : \widehat{x}_i \in \mathcal{N}_\gamma(x_0)\}, \qquad \mathcal{I}_2 = \{i \in [N] : \widehat{x}_i \notin \mathcal{N}_\gamma(x_0)\}. \tag{13}$$

The sets $\mathcal{I}_1$ and $\mathcal{I}_2$ divides the training samples into two mutually exclusive sets dependent on whether the training samples fall inside or outside the fiber. For any $i \in [N]$, let d_i be the distance from $\widehat{x}_i$ to the boundary $\partial \mathcal{N}_\gamma(x_0)$ of the set $\mathcal{N}_\gamma(x_0)$, that is,

$$\forall i \in [N] : \qquad d_i = \min_{x \in \partial \mathcal{N}_\gamma(x_0)} \mathbb{D}_\mathcal{X}(x, \widehat{x}_i). \tag{14}$$

Note that the distance d_i defined above is closely related to the values of κ_i that is defined in (7). Indeed, if $\widehat{x}_i \notin \mathcal{N}_\gamma(x_0)$ then d_i is equal to κ_i. However, if $\widehat{x}_i$ is in the interior of the set $\mathcal{N}_\gamma(x_0)$ then $d_i > 0$ while $\kappa_i = 0$. Using the definition of the set $\mathcal{I}_1$ and the boundary projection distance d_i, the maximum radius $\rho_{\max}(x_0, \gamma)$ can be computed in closed form, as asserted by the next proposition.

Proposition 4.1 (Expression for $\rho_{\max}(x_0, \gamma)$). *We have* $\rho_{\max}(x_0, \gamma) = N^{-1} \sum_{i \in \mathcal{I}_1} d_i$.

Proof of Proposition 4.1. Without any loss of optimality, we can substitute the constraint $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) = 0$ by the inequality constraint $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \leq 0$ to obtain

$$\rho_{\max}(x_0, \gamma) = \inf\{\mathrm{W}(\mathbb{Q}, \widehat{\mathbb{P}}) : \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}), \ \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \leq 0\}$$

$$= \begin{cases} \inf & \mathbb{E}_\pi[\mathbb{D}(\xi, \xi')] \\ \mathrm{s.\,t.} & \pi \in \mathcal{M}((\mathcal{X} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{Y})) \\ & \mathbb{E}_\pi[\mathbb{1}_{\mathcal{N}_\gamma(x_0) \times \mathcal{Y}}(\xi)] \leq 0 \\ & \mathbb{E}_\pi[\mathbb{1}_{(\widehat{x}_i, \widehat{y}_i)}(\xi')] = \frac{1}{N} \quad \forall i \in [N], \end{cases}$$

where ξ represents the joint random vector (X, Y). By a weak duality result [61, Section 2.2], we have

$$\rho_{\max}(x_0, \gamma) \geq \begin{cases} \sup & \frac{1}{N} \sum_{i=1}^{N} b_i \\ \text{s.t.} & b \in \mathbb{R}^N, \ \zeta \in \mathbb{R}_+ \\ & \sum_{i=1}^{N} b_i \mathbb{1}_{(\widehat{x}_i, \widehat{y}_i)}(x', y') - \zeta \mathbb{1}_{\mathcal{N}_\gamma(x_0)}(x) \leq \mathbb{D}\big((x, y), (x', y')\big) \quad \forall (x, y), (x', y') \in \mathcal{X} \times \mathcal{Y} \end{cases}$$

$$= \begin{cases} \sup & \frac{1}{N} \sum_{i=1}^{N} b_i \\ \text{s.t.} & b \in \mathbb{R}^N, \ \zeta \in \mathbb{R}_+ \\ & b_i - \zeta \mathbb{1}_{\mathcal{N}_\gamma(x_0)}(x_i) \leq \mathbb{D}\big((x_i, y_i), (\widehat{x}_i, \widehat{y}_i)\big) \quad \forall (x_i, y_i) \in \mathcal{X} \times \mathcal{Y}, \ \forall i \in [N] \end{cases}$$

$$= \begin{cases} \sup_{\zeta \in \mathbb{R}_+} \frac{1}{N} \sum_{i \in [N]} \inf_{(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}} \big\{ \mathbb{D}\big((x_i, y_i), (\widehat{x}_i, \widehat{y}_i)\big) + \zeta \mathbb{1}_{\mathcal{N}_\gamma(x_0)}(x_i) \big\} \end{cases}$$

$$= \sup_{\zeta \in \mathbb{R}_+} \frac{1}{N} \sum_{i \in [N]} \inf_{x_i \in \mathcal{X}} \big\{ \mathbb{D}_\mathcal{X}(x_i, \widehat{x}_i) + \zeta \mathbb{1}_{\mathcal{N}_\gamma(x_0)}(x_i) \big\},$$

where the last equality follows from the decomposition of $\mathbb{D}$ and the fact that the minimizer in y_i is $\widehat{y}_i$. For any $\zeta \geq 0$, we have for any $i \in [N]$

$$\inf_{x_i \in \mathcal{X}} \big\{ \mathbb{D}_\mathcal{X}(x_i, \widehat{x}_i) + \zeta \mathbb{1}_{\mathcal{N}_\gamma(x_0)}(x_i) \big\} = \begin{cases} \min\{\zeta, d_i\} & \text{if } \widehat{x}_i \in \mathcal{N}_\gamma(x_0), \\ 0 & \text{if } \widehat{x}_i \notin \mathcal{N}_\gamma(x_0), \end{cases}$$

which implies that

$$\rho_{\max}(x_0, \gamma) \geq \sup_{\zeta \in \mathbb{R}_+} \frac{1}{N} \sum_{i \in \mathcal{I}_1} \min(\zeta, d_i) = \frac{1}{N} \sum_{i \in \mathcal{I}_1} d_i.$$

In the last step, we show that the above inequality is tight. For any value ρ such that $\rho > \frac{1}{N} \sum_{i \in \mathcal{I}_1} d_i$, Assumption 2.4(iii) and the continuity of $\mathbb{D}_\mathcal{X}$ and $\mathbb{D}_\mathcal{Y}$ imply that there exists $(x_i', y_i') \in (\mathcal{X} \times \mathcal{Y}) \backslash (\mathcal{N}_\gamma(x_0) \times \mathcal{Y})$ such that

$$\mathbb{D}_\mathcal{X}(x_i', \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(y_i', \widehat{y}_i) \leq d_i + \frac{1}{N}(\rho - \frac{1}{N} \sum_{i \in \mathcal{I}_1} d_i).$$

The distribution

$$\mathbb{Q}' = \frac{1}{N} \Big(\sum_{i \in \mathcal{I}_1} \delta_{(x_i', y_i')} + \sum_{i \in \mathcal{I}_2} \delta_{(\widehat{x}_i, \widehat{y}_i)} \Big)$$

thus satisfies $\mathbb{W}(\mathbb{Q}', \widehat{\mathbb{P}}) \leq \rho$ and $\mathbb{Q}'(X \in \mathcal{N}_\gamma(x_0)) \leq 0$. This implies that $\rho_{\max}(x_0, \gamma) = N^{-1} \sum_{i \in \mathcal{I}_1} d_i$ and completes the proof. $\qquad \square$

The result of Proposition 4.1 is also intuitive: if $\widehat{x}_i$ falls in the fiber $\mathcal{N}_\gamma(x_0)$ then $i \in \mathcal{I}_1$, and we will need a sample-wise budget of d_i/N to transport $\widehat{x}_i$ out of $\mathcal{N}_\gamma(x_0)$. The value $\rho_{\max}(x_0, \gamma)$ is thus obtained by summing all d_i/N over the set $\mathcal{I}_1$. If $\mathcal{I}_1 = \emptyset$ then there is no training sample inside the fiber set, and thus $\rho_{\max}(x_0, \gamma) = 0$.

The computation of $\rho_{\max}(x_0, \gamma)$ provides a natural upper bound on the radius ρ. Indeed, if the radius ρ prescribing the ambiguity set is bigger than $\rho_{\max}(x_0, \gamma)$, then we recover the robust worst-case conditional loss. This robust loss is uninformative because it depends only on the support $\mathcal{Y}$ and it is *independent* of the training samples. This negative result is a reminiscence of Proposition 4.2 and highlights once again the sophistication of the distributionally robust conditional decision-making problem.

Proposition 4.2 (Uninformative solution when ρ is sufficiently large)**.** *When $\gamma \in \mathbb{R}_{++}$, $\varepsilon = 0$, and $\rho > \rho_{\max}(x_0, \gamma)$, we have*

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) > 0} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)] = \sup_{y \in \mathcal{Y}} \ell(y, \alpha, \beta).$$

Proof of Proposition 4.2. For simplicity, we assume that $\rho_{\max}(x_0, \gamma) > 0$, which implies that $\mathcal{I}_1 \neq \emptyset$. Let's consider some index $j \in \mathcal{I}_1$. Assumption 2.4(iii) and the continuity of $\mathbb{D}_\mathcal{X}$ and $\mathbb{D}_\mathcal{Y}$ imply that there exists $x_i' \in \mathcal{X} \backslash \mathcal{N}_\gamma(x_0)$ such that

$$\mathbb{D}_\mathcal{X}(x_i', \widehat{x}_i) \leq d_i + \frac{1}{2N}(\rho - \rho_{\max}(x_0, \gamma)).$$

Fix an arbitrary value $y_0 \in \mathcal{Y}$. Consider the distribution

$$\mathbb{Q}' = \frac{1}{N}\Big(\sum_{i \in \mathcal{I}_1 \backslash \{j\}} \delta_{(x_i', \widehat{y}_i)} + \sum_{i \in \mathcal{I}_2} \delta_{(\widehat{x}_i, \widehat{y}_i)} + \upsilon \delta_{(\widehat{x}_j, y_0)} + (1 - \upsilon) \delta_{(\widehat{x}_j, \widehat{y}_j)} \Big)$$

in which $\upsilon \in \mathbb{R}_{++}$ is chosen so that

$$\upsilon \mathbb{D}_\mathcal{Y}(y_0, \widehat{y}_j) \leq \frac{1}{2}(\rho - \rho_{\max}(x_0, \gamma)).$$

It is easy to verify that $\mathbb{Q}'(X \in \mathcal{N}_\gamma(x_0)) = \upsilon/N > 0$, $\mathbb{W}(\mathbb{Q}', \widehat{\mathbb{P}}) \leq \rho$ and that

$$\mathbb{E}_{\mathbb{Q}'}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)] = \ell(y_0, \alpha, \beta).$$

Because y_0 has been chosen arbitrarily, a similar argument as in the proof of Proposition 3.1 leads to the necessary result.

Note that alternatively, the case where $\rho_{\max}(x_0, \gamma) = 0$ implies that all samples $\{\widehat{x}_i\}_{i \in [N]}$ lie in the set $(\mathcal{X} \backslash \mathcal{N}_\gamma(x_0)) \cup \partial\mathcal{N}_\gamma(x_0)$. If $\mathcal{I}_1 \neq \emptyset$, then the same argument as before applies. Otherwise, fixing j to be such that $\widehat{x}_j \notin \mathcal{N}_\gamma(x_0)$, the same results is obtained with

$$\mathbb{Q}' = \frac{1}{N}\Big(\sum_{i \in \mathcal{I}_2 \backslash \{j\}} \delta_{(\widehat{x}_i, \widehat{y}_i)} + \upsilon \delta_{(\widehat{x}_j^*, y_0)} + (1 - \upsilon) \delta_{(\widehat{x}_j, \widehat{y}_j)} \Big)$$

where $\widehat{x}_j^\star$ is the projection of $\widehat{x}_j$ onto $\partial\mathcal{N}_\gamma(x_0)$. The proof is complete. $\square$

We now provide a heuristic justification for the result of Proposition 4.2. To form the worst-case distribution, the adversary first moves all the samples with index in $\mathcal{I}_1$ out of the fiber, and this would cost an amount of energy $\rho_{\max}(x_0, \gamma)$. After that, the adversary can pick any sample, slice out an infinitesimally small amount of mass, and then move that slice to any point in the fiber $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$. Because $\rho > \rho_{\max}(x_0, \gamma)$ and because the ground metric is continuous, this new arrangement of the samples is feasible, and constitutes the worst-case distribution that maximizes the conditional expected loss.

Consider now the case in which the ambiguity size ρ is strictly smaller than the maximum value $\rho_{\max}(x_0, \gamma)$. It can be shown that in this situation, any distribution $\mathbb{Q}$ that is feasible in the supremum problem of (11) should satisfy $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \underline{\varepsilon}$ for some strictly positive lower bound $\underline{\varepsilon}$. Further, this value $\underline{\varepsilon}$ can be quantified by solving a linear optimization problem.

Proposition 4.3 (Strictly positive probability requirement equivalence)**.** *Suppose that $\gamma \in \mathbb{R}_{++}$, $\varepsilon = 0$ and $\rho < \rho_{\max}(x_0, \gamma)$. Then there exists an $\underline{\varepsilon} > 0$, such that the distributionally robust portfolio allocation model with side information* (11) *is equivalent to*

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \underline{\varepsilon}} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)].$$

In particular, this equivalence holds for:

$$\underline{\varepsilon} = \min\left\{ \frac{1}{N} \sum_{i \in \mathcal{I}_1} p_i : p \in [0, 1]^N, \ \frac{1}{N} \sum_{i \in \mathcal{I}_1} d_i(1 - p_i) \leq \rho \right\}.$$

Proof of Proposition 4.3. We start by proving by contradiction that, when $\rho < \rho_{\max}(x_0, \gamma)$, we necessarily have that $\inf_{\mathbb{Q} \in \mathbb{B}_\rho} \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) > 0$. Let us assume that:

$$\inf_{\mathbb{Q} \in \mathbb{B}_\rho} \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) = 0 \,.$$

This implies that for all $\varepsilon > 0$, there exists a $\mathbb{Q} \in \mathbb{B}_\rho$ such that $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \leq \varepsilon$. Based on the following representation of $\mathbb{B}_\rho$:

$$\mathbb{B}_\rho = \left\{ \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) : \begin{array}{l} \exists \, \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) \,\, \forall i \in [N] \text{ such that } \mathbb{Q} = N^{-1} \sum_{i \in [N]} \pi_i \\ \frac{1}{N} \sum_{i \in [N]} \mathbb{E}_{\pi_i}[\mathbb{D}_\mathcal{X}(X, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq \rho \end{array} \right\}, \quad (15)$$

it must be that there exists an assignment for $\{\pi_i\}_{i=1}^N$ that satisfies:

$$\frac{1}{N} \sum_{i \in [N]} \mathbb{E}_{\pi_i}[\mathbb{D}_\mathcal{X}(X, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq \rho$$

and

$$\frac{1}{N} \sum_{i \in [N]} \pi_i(X \in \mathcal{N}_\gamma(x_0)) \leq \varepsilon.$$

We can work out the following steps:

$$\rho \geq \frac{1}{N} \sum_{i \in [N]} \mathbb{E}_{\pi_i}[\mathbb{D}_\mathcal{X}(X, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \geq \frac{1}{N} \sum_{i \in [N]} \mathbb{E}_{\pi_i}[\mathbb{D}_\mathcal{X}(X, \widehat{x}_i) | X \notin \mathcal{N}_\gamma(x_0)] \pi_i(X \notin \mathcal{N}_\gamma(x_0))$$

$$\geq \frac{1}{N} \sum_{i \in [N]} d_i (1 - \pi_i(X \in \mathcal{N}_\gamma(x_0))) \geq \frac{1}{N} \sum_{i \in [N]} d_i - \left(\max_i d_i \right) \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \rho_{\max} - \left(\max_i d_i \right) \varepsilon \,.$$

Given that this is the case for all ε, we can conclude that $\rho \geq \rho_{\max}$, which contradicts our assumption about the ambiguity size ρ.

Next, we turn to how to evaluate $\underline{\varepsilon} := \inf_{\mathbb{Q} \in \mathbb{B}_\rho} \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0))$. Given any $\bar{\mathbb{Q}} \in \mathbb{B}_\rho$, associated to some $\{\bar{\pi}_i\}_{i \in [N]}$ based on the representation in (15), one can construct a new measure $\bar{\mathbb{Q}}^\star$ of the form:

$$\mathbb{Q}_p^\star = \frac{1}{N} \sum_{i \in \mathcal{I}_1} \left((1 - p_i) \delta_{(\widehat{x}_i^\star, \widehat{y}_i)} + p_i \delta_{(\widehat{x}_i, \widehat{y}_i)} \right) + \frac{1}{N} \sum_{i \in \mathcal{I}_2} \delta_{(\widehat{x}_i, \widehat{y}_i)},$$

where $\widehat{x}_i^\star \in \mathcal{X} \backslash \mathcal{N}_\gamma(x_0)$ is a point arbitrarily close to the projection of $\widehat{x}_i$ onto $\partial \mathcal{N}_\gamma(x_0)$, based on $\bar{\mathbb{Q}}^\star \triangleq \mathbb{Q}_{\bar{p}}^\star$ with $\bar{p}_i \triangleq \bar{\pi}_i(X \in \mathcal{N}_\gamma(x_0))$. It is easy to see how $\bar{\mathbb{Q}}^\star \in \mathbb{B}_\rho$ and achieves:

$$\bar{\mathbb{Q}}^\star(X \in \mathcal{N}_\gamma(x_0)) \leq \bar{\mathbb{Q}}(X \in \mathcal{N}_\gamma(x_0)) \,.$$

Hence, using similar arguments as in the proof of Proposition 4.2, we have that

$$\underline{\varepsilon} = \begin{cases} \inf & \mathbb{Q}_p^\star(X \in \mathcal{N}_\gamma(x_0)) \\ \text{s.t.} & p \in [0, 1]^N, \,\, \widehat{x}_i^\star \in \mathcal{X} \backslash \mathcal{N}_\gamma(x_0) \,\, \forall i \in \mathcal{I}_1 \\ & \mathbb{Q}_p^\star = \frac{1}{N} \sum_{i \in \mathcal{I}_1} \left((1 - p_i) \delta_{(\widehat{x}_i^\star, \widehat{y}_i)} + p_i \delta_{(\widehat{x}_i, \widehat{y}_i)} \right) + \frac{1}{N} \sum_{i \in \mathcal{I}_2} \delta_{(\widehat{x}_i, \widehat{y}_i)} \\ & \frac{1}{N} \sum_{i \in \mathcal{I}_1} d_i (1 - p_i) \leq \rho \end{cases}$$

$$= \min \left\{ \frac{1}{N} \sum_{i \in \mathcal{I}_1} p_i : p \in [0, 1]^N, \,\, \frac{1}{N} \sum_{i \in \mathcal{I}_1} d_i (1 - p_i) \leq \rho \right\} \,.$$

This completes the proof. $\qquad \square$

The linear program that defines $\underline{\varepsilon}$ in Proposition 4.3 can be further shown to be a fractional knapsack problem, and it can be solved efficiently to optimality using a greedy heuristics [38, Proposition 17.1]. The conditions of Proposition 4.3 is satisfied only when $\rho_{\max}(x_0, \gamma) > 0$, which further implies that $\mathcal{I}_1 \neq \emptyset$ and there exists at least one training sample in the fiber set $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$. With a ambiguity size ρ which is strictly smaller than $\rho_{\max}(x_0, \gamma)$, it is not possible for the adversary to remove all the samples out of the fiber set

$\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$. The lower bound value $\underline{\varepsilon}$ here represents the lowest possible amount of probability mass that is left on the fiber.

The result of Proposition 4.3 indicates that the portfolio allocation problem when $\rho < \rho_{\max}(x_0, \gamma)$ can be obtained by solving the general problem with a strictly positive probability mass requirement. This general case is our subject of study in the next subsection.

4.2. Strictly positive fiber probability $\varepsilon > 0$

We now consider the last, and also the most general, case with a nonsingular fiber set $\gamma > 0$ and a strictly positive fiber probability $\varepsilon > 0$. More specifically, we aim to solve the portfolio allocation of the form

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathcal{B}} \ \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]. \tag{16}$$

As the first step towards solving (16), we present in the next theorem the finite dimensional reformulation of the worst-case conditional expected loss.

Theorem 4.4 (Worst-case conditional expected loss for $\varepsilon > 0$). *Suppose that $\gamma \in \mathbb{R}_{++}$, $\varepsilon \in \mathbb{R}_{++}$, and $\rho > \rho_{\min}(x_0, \gamma, \varepsilon)$. For any feasible solution (α, β), we have*

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]$$

$$= \begin{cases} \inf & \phi + (N\varepsilon)^{-1}\nu^+ - N^{-1}\nu^- \\ \text{s.t.} & \lambda \in \mathbb{R}_+^N, \ s \in \mathbb{R}^N, \ \nu^+ \in \mathbb{R}_+, \ \nu^- \in \mathbb{R}_+, \ \phi \in \mathbb{R}, \ \varphi \in \mathbb{R}_+, \ \psi \in \mathbb{R}_+^N \\ & \phi - d_i\varphi + \psi_i - s_i \geq 0 & \forall i \in \mathcal{I}_1 \\ & \phi + d_i\varphi + \psi_i - s_i \geq 0 & \forall i \in \mathcal{I}_2 \\ & \nu^+ - \nu^- + (\sum_{i \in \mathcal{I}_1} d_i - N\rho)\varphi - \sum_{i \in [N]} \psi_i \geq 0 \\ & \varphi - \lambda_i \geq 0 & \forall i \in [N] \\ & s_i \geq \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\} & \forall i \in [N], \end{cases}$$

where the parameters d_i are defined as in (14).

Proof of Theorem 4.4. By exploiting the definition of the optimal transport cost and the fact that any joint probability measure $\pi \in \Pi(\mathbb{Q}, \widehat{\mathbb{P}})$ can be written as $\pi = N^{-1} \sum_{i \in [N]} \pi_i \otimes \delta_{(\widehat{x}_i, \widehat{y}_i)}$, where each π_i is a probability measure on $\mathcal{X} \times \mathcal{Y}$, we have

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)] = \begin{cases} \sup & \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)] \\ \text{s.t.} & \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) \quad \forall i \in [N] \\ & \sum_{i \in [N]} \pi_i(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq N\varepsilon \\ & \mathbb{Q} = N^{-1} \sum_{i \in [N]} \pi_i, \ \sum_{i \in [N]} \mathbb{W}(\pi_i, \delta_{(\widehat{x}_i, \widehat{y}_i)}) \leq N\rho \end{cases}$$

$$= \begin{cases} \sup & \dfrac{\sum_{i \in [N]} \mathbb{E}_{\pi_i}[\ell(Y, \alpha, \beta) \mathbb{1}_{\mathcal{N}_\gamma(x_0)}(X)]}{\sum_{i \in [N]} \pi_i(X \in \mathcal{N}_\gamma(x_0))} \\ \text{s.t.} & \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) \quad \forall i \in [N] \\ & \sum_{i \in [N]} \pi_i(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq N\varepsilon \\ & \sum_{i \in [N]} \mathbb{W}(\pi_i, \delta_{(\widehat{x}_i, \widehat{y}_i)}) \leq N\rho. \end{cases}$$

For any $i \in \mathcal{I}$, let $\widehat{x}_i^\star = \arg \min_{x \in \partial \mathcal{N}_\gamma(x_0)} \mathbb{D}_{\mathcal{X}}(x, \widehat{x}_i)$ be again the projection of $\widehat{x}_i$ onto $\partial \mathcal{N}_\gamma(x_0)$. Moreover, for any $i \in \mathcal{I}_1$, let (x_i', y_i') be a point on $(\mathcal{X} \times \mathcal{Y}) \backslash (\mathcal{N}_\gamma(x_0) \times \mathcal{Y})$. Using a continuity and greedy argument, we can choose (x_i', y_i') sufficiently close to $(\widehat{x}_i^\star, \widehat{y}_i)$ so that it suffices to consider the conditional distribution π_i of the form

$$\pi_i(\mathrm{d}x \times \mathrm{d}y) = \begin{cases} p_i \delta_{\widehat{x}_i}(\mathrm{d}x) \mu_0^i(\mathrm{d}y) + (1 - p_i)\delta_{(x_i', y_i')}(\mathrm{d}x \times \mathrm{d}y) & \text{if } i \in \mathcal{I}_1, \\ p_i \delta_{\widehat{x}_i^\star}(\mathrm{d}x) \mu_0^i(\mathrm{d}y) + (1 - p_i)\delta_{(\widehat{x}_i, \widehat{y}_i)}(\mathrm{d}x \times \mathrm{d}y) & \text{if } i \in \mathcal{I}_2, \end{cases} \tag{17}$$

for a parameter $p_i \in [0,1]$ and a conditional measure $\mu_0^i \in \mathcal{M}(\mathcal{Y})$. Intuitively, p_i represents the portion of the sample point $(\widehat{x}_i, \widehat{y}_i)$ that is transported to the fiber $\mathcal{N}_\gamma(x_0) \times \mathcal{Y}$, and μ_0^i is the conditional distribution of Y given $X \in \mathcal{N}_\gamma(x_0)$ that is obtained by transporting the sample point $(\widehat{x}_i, \widehat{y}_i)$. Using this representation, we can rewrite $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) = N^{-1} \sum_{i \in [N]} p_i$. By exploiting the definition of d_i, the worst-case conditional expected loss can be re-expressed as

$$
\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]
$$

$$
= \begin{cases}
\sup & \dfrac{\sum_{i \in [N]} p_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)]}{\sum_{i \in [N]} p_i} \\
\text{s.t.} & p_i \in [0,1], \ \mu_0^i \in \mathcal{M}(\mathcal{Y}) \quad \forall i \in [N] \\
& \displaystyle\sum_{i \in \mathcal{I}_1} \left(p_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] + (1 - p_i) d_i \right) + \sum_{i \in \mathcal{I}_2} \left(p_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] + p_i d_i \right) \leq N\rho \\
& \sum_{i \in \mathcal{I}} p_i \geq N\varepsilon.
\end{cases}
$$

For any $p = (p_i)_{i=1,\dots,N} \in [0,1]^N$ such that $\sum_{i \in [N]} p_i \geq N\varepsilon$, define the function

$$
f(p) \triangleq \begin{cases}
\sup & \displaystyle\sup_{\mu_0^i : \mathbb{E}_{\mu_0^i}[p_i \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq q_i \ \forall i} \sum_{i \in [N]} \mathbb{E}_{\mu_0^i}[p_i \ell(Y, \alpha, \beta)] \\
\text{s.t.} & q \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{I}_1} (q_i + (1 - p_i) d_i) + \sum_{i \in \mathcal{I}_2} (q_i + p_i d_i) \leq N\rho.
\end{cases}
$$

Using the general Charnes-Cooper variable transformation

$$
v_i = \frac{p_i}{\sum_{i \in [N]} p_i}, \quad t = \frac{1}{\sum_{i \in [N]} p_i},
$$

the equivalent characterization [57, Lemma 1] implies that

$$
\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)] = \sup_{p \in [0,1]^N, \sum_{i \in [N]} p_i \geq \varepsilon} \frac{f(p)}{\sum_{i \in [N]} p_i}
$$

$$
= \begin{cases}
\sup_{v,t} & t f(v/t) \\
\text{s.t.} & v \in [0,1]^N, \ t \in \mathbb{R} \\
& v_i \leq t \quad \forall i \in [N], \ \sum_{i \in [N]} v_i = 1 \\
& N^{-1} \leq t \leq (N\varepsilon)^{-1},
\end{cases} \tag{18}
$$

where the fact that $t \geq N^{-1}$ is implied by $v_i \leq t$ and $\sum_{i \in [N]} v_i = 1$ yet makes explicit the fact that $t > 0$.

For any feasible value of v and t, its corresponding objective value can be evaluated as

$$
t f(v/t) = \begin{cases}
\sup & \displaystyle\sup_{\mu_0^i : \mathbb{E}_{\mu_0^i}[v_i \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq q_i t \ \forall i} \sum_{i \in [N]} \mathbb{E}_{\mu_0^i}[v_i \ell(Y, \alpha, \beta)] \\
\text{s.t.} & q \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{I}_1} \left(q_i + (1 - \frac{v_i}{t}) d_i \right) + \sum_{i \in \mathcal{I}_2} \left(q_i + \frac{v_i}{t} d_i \right) \leq N\rho
\end{cases}
$$

$$
= \begin{cases}
\sup & \displaystyle\sup_{\mu_0^i : \mathbb{E}_{\mu_0^i}[v_i \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq q_i t \ \forall i} \sum_{i \in [N]} \mathbb{E}_{\mu_0^i}[v_i \ell(Y, \alpha, \beta)] \\
\text{s.t.} & q \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{I}_1} (q_i t + (t - v_i) d_i) + \sum_{i \in \mathcal{I}_2} (q_i t + v_i d_i) \leq N\rho t
\end{cases}
$$

$$
= \begin{cases}
\sup & \displaystyle\sup_{\mu_0^i : \mathbb{E}_{\mu_0^i}[v_i \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq \theta_i \ \forall i} \sum_{i \in [N]} \mathbb{E}_{\mu_0^i}[v_i \ell(Y, \alpha, \beta)] \\
\text{s.t.} & \theta \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{I}_1} (\theta_i + (t - v_i) d_i) + \sum_{i \in \mathcal{I}_2} (\theta_i + v_i d_i) \leq N\rho t,
\end{cases}
$$

where the last equality follows from the change of variable $\theta_i \leftarrow q_i t$. The inner supremum problems over μ_0^i are separable, and can be written using standard conic duality results [12, Theorem 1] as

$$
\sup_{\mu_0^i : \mathbb{E}_{\mu_0^i}[v_i \mathbb{D}_{\mathcal{Y}}(Y, \widehat{y}_i)] \leq \theta_i \ \forall i} \sum_{i \in [N]} \mathbb{E}_{\mu_0^i}[v_i \ell(Y, \alpha, \beta)] = \inf_{\lambda \in \mathbb{R}_+^N} \sum_{i \in [N]} \theta_i \lambda_i + v_i \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\}.
$$

Thus we find for any feasible (v, t)

$$
tf(v/t) = \begin{cases}
\sup & \inf_{\lambda \in \mathbb{R}_+^N} \sum_{i \in [N]} \theta_i \lambda_i + v_i \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\} \\
\text{s.t.} & \theta \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{I}_1}(\theta_i + (t - v_i)d_i) + \sum_{i \in \mathcal{I}_2}(\theta_i + v_i d_i) \leq N\rho t,
\end{cases}
$$
$$
= \begin{cases}
\inf_{\lambda \in \mathbb{R}_+^N} & \sup \quad \sum_{i \in [N]} \theta_i \lambda_i + v_i \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\} \\
& \text{s.t.} \quad \theta \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{I}_1}(\theta_i + (t - v_i)d_i) + \sum_{i \in \mathcal{I}_2}(\theta_i + v_i d_i) \leq N\rho t
\end{cases} \tag{19}
$$

where the second equality holds thanks to Sion's minimax theorem because the feasible set for θ is compact, and the function

$$
(\theta, \lambda) \mapsto \sum_{i \in [N]} \theta_i \lambda_i + v_i \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\}
$$

is convex in λ, affine in θ and jointly continuous in both θ and λ. Let $\mathcal{Q}(v, t)$ be the feasible set for the variable θ of the form

$$
\mathcal{Q}(v, t) = \left\{ \theta \in \mathbb{R}_+^N : \sum_{i \in \mathcal{I}_1}(\theta_i + (t - v_i)d_i) + \sum_{i \in \mathcal{I}_2}(\theta_i + v_i d_i) \leq N\rho t \right\}.
$$

By rejoining (18), (19) and definition of $\mathcal{Q}(v, t)$ above, we have

$$
\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)]
$$
$$
= \sup_{(v,t) \in \Upsilon} \inf_{\lambda \in \mathbb{R}_+^N} \sup_{\theta \in \mathcal{Q}(v,t)} \sum_{i \in [N]} \theta_i \lambda_i + v_i \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\},
$$

where the feasible set Υ of (v, t) in the above problem is understood to be the feasible set of (18), that is,

$$
\Upsilon \triangleq \left\{ (v, t) \in [0, 1]^N \times \mathbb{R} : v_i \leq t \quad \forall i \in [N], \ \sum_{i \in [N]} v_i = 1, \ N^{-1} \leq t \leq (N\varepsilon)^{-1} \right\}.
$$

Since Υ is compact, Sion's minimax theorem [60] applies and we obtain

$$
\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)]
$$
$$
= \inf_{\lambda \in \mathbb{R}_+^N} \sup_{(v,t) \in \Upsilon} \sup_{\theta \in \mathcal{Q}(v,t)} \sum_{i \in [N]} \theta_i \lambda_i + v_i \sup_{y_i \in \mathcal{Y}} \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\}
$$
$$
= \inf_{\lambda \in \mathbb{R}_+^N} \sup_{(v,t) \in \Upsilon} \sup_{\theta \in \mathcal{Q}(v,t)} \sum_{i \in [N]} \theta_i \lambda_i + v_i \inf_{s_i \in \mathcal{S}_i(\alpha, \beta, \lambda_i)} s_i
$$
$$
= \inf_{\lambda \in \mathbb{R}_+^N, s_i \in \mathcal{S}_i(\alpha, \beta, \lambda_i) \forall i} \sup_{(v,t) \in \Upsilon} \sup_{\theta \in \mathcal{Q}(v,t)} \sum_{i \in [N]} \theta_i \lambda_i + v_i s_i \tag{20}
$$

where $\mathcal{S}_i(\alpha, \beta, \lambda_i) \triangleq \{s \in \mathbb{R} : s \geq \ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i) \ \forall y_i \in \mathcal{Y}\}$, and where we exploited the fact that the supremum over $y_i \in \mathcal{Y}$ is not affected by the choice of v, t, and θ, and that each $v_i \geq 0$ to get rid of the apparent bilinearity in v_i and y_i. Given that the inner two layers of supremum represent a linear program, linear programming duality can be applied to obtain:

$$
\begin{aligned}
\inf \quad & \phi + (N\varepsilon)^{-1}\nu^+ - N^{-1}\nu^- \\
\text{s.t.} \quad & \nu^+ \in \mathbb{R}_+, \ \nu^- \in \mathbb{R}_+, \ \phi \in \mathbb{R}, \ \varphi \in \mathbb{R}_+, \ \psi \in \mathbb{R}_+^N \\
& \phi - d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_1 \\
& \phi + d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_2 \\
& \nu^+ - \nu^- + (\textstyle\sum_{i \in \mathcal{I}_1} d_i - N\rho)\varphi - \textstyle\sum_{i \in [N]} \psi_i \geq 0 \\
& \varphi - \lambda_i \geq 0 && \forall i \in [N].
\end{aligned}
\tag{21}
$$

One can confirm that strong linear programming duality necessarily applied given that $\rho > \rho_{\min}(x_0, \gamma, \varepsilon)$ implies the existence of a feasible $\mathbb{Q} \in \mathbb{B}_\rho$ such that $\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon$. This probability measure can be used to create a feasible triplet (υ, t, θ) for the supremum problem:

$$
t \triangleq 1/(N\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0))), \quad v_i \triangleq \pi_i(X \in \mathcal{N}_\gamma(x_0))/(N\mathbb{Q}(X \in \mathcal{N}_\gamma(x_0))), \quad \theta_i \triangleq 0,
$$

with $\{\pi_i\}_{i \in [N]}$ as the set of probability measures that certify that $\mathbb{Q} \in \mathbb{B}_\rho$. Rejoining the infimum operation in (21) in (20) completes the proof. $\qquad\square$

Under the condition of Theorem 4.4, the portfolio allocation problem with side information (16) is equivalent to the following finite-dimensional optimization problem

$$
\begin{aligned}
\min \quad & \phi + (N\varepsilon)^{-1}\nu^+ - N^{-1}\nu^- \\
\text{s.t.} \quad & \alpha \in \mathcal{A}, \ \beta \in \mathcal{B}, \ \lambda \in \mathbb{R}_+^N, \ s \in \mathbb{R}^N, \ \nu^+ \in \mathbb{R}_+, \ \nu^- \in \mathbb{R}_+, \ \phi \in \mathbb{R}, \ \varphi \in \mathbb{R}_+, \ \psi \in \mathbb{R}_+^N \\
& \phi - d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_1 \\
& \phi + d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_2 \\
& \nu^+ - \nu^- + (\textstyle\sum_{i \in \mathcal{I}_1} d_i - N\rho)\varphi - \textstyle\sum_{i \in [N]} \psi_i \geq 0 \\
& \varphi - \lambda_i \geq 0 && \forall i \in [N] \\
& s_i \geq \sup_{y_i \in \mathcal{Y}} \ \{\ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_\mathcal{Y}(y_i, \widehat{y}_i)\} && \forall i \in [N].
\end{aligned}
\tag{22}
$$

The last constraint of problem (22) is a semi-infinite constraint which can be reformulated into a set of semi-definite constraints under specific situations. These results are highlighted in Proposition 4.5 and 4.6.

Proposition 4.5 (Mean-variance loss function). *Suppose that ℓ is the mean-variance loss function of the form (5), $\gamma \in \mathbb{R}_{++}$, $\varepsilon \in (0, 1]$ and $\rho > \rho_{\min}(x_0, \gamma, \varepsilon)$. Suppose in addition that $\mathcal{Y} = \mathbb{R}^m$ and $\mathbb{D}_\mathcal{Y}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$. The distributionally robust portfolio allocation model with side information (22) is equivalent to the semi-definite optimization problem*

$$
\begin{aligned}
\min \quad & \phi + (N\varepsilon)^{-1}\nu^+ - N^{-1}\nu^- \\
\text{s.t.} \quad & \alpha \in \mathcal{A}, \ \beta \in \mathcal{B}, \ \lambda \in \mathbb{R}_+^N, \ s \in \mathbb{R}^N, \ t \in \mathbb{R}_+, \ A_i \in \mathbb{S}_+^m \quad \forall i \in [N] \\
& \nu^+ \in \mathbb{R}_+, \ \nu^- \in \mathbb{R}_+, \ \phi \in \mathbb{R}, \ \varphi \in \mathbb{R}_+, \ \psi \in \mathbb{R}_+^N \\
& \phi - d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_1 \\
& \phi + d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_2 \\
& \nu^+ - \nu^- + (\textstyle\sum_{i \in \mathcal{I}_1} d_i - N\rho)\varphi - \textstyle\sum_{i \in [N]} \psi_i \geq 0 \\
& \varphi - \lambda_i \geq 0 && \forall i \in [N] \\
& \begin{bmatrix} \lambda_i I - A_i & \alpha \\ \alpha^\top & 1 \end{bmatrix} \succeq 0, \ \begin{bmatrix} A_i & (\beta + \eta/2)\alpha - \lambda_i\widehat{y}_i \\ (\beta + \eta/2)\alpha^\top - \lambda_i\widehat{y}_i^\top & s_i + \lambda_i\|\widehat{y}_i\|_2^2 - t \end{bmatrix} \succeq 0 \quad \forall i \in [N] \\
& \begin{bmatrix} t & \beta \\ \beta & 1 \end{bmatrix} \succeq 0,
\end{aligned}
$$

where the parameters d_i are defined as in (14).

Proof of Proposition 4.5. For each $i \in [N]$, by the definition of the loss function ℓ in (5) and by the choice of the transport cost $\mathbb{D}_\mathcal{Y}$, the constraint

$$
\sup_{y_i \in \mathcal{Y}} \ \ell(y_i, \alpha, \beta) - \lambda_i \mathbb{D}_\mathcal{Y}(y_i, \widehat{y}_i) \leq s_i
$$

is equivalent to

$$y_i^\top(\alpha\alpha^\top - \lambda_i I)y_i + \left(2\lambda_i\widehat{y}_i - (2\beta + \eta)\alpha\right)^\top y_i + \beta^2 - \lambda_i\|\widehat{y}_i\|_2^2 - s_i \leq 0 \quad \forall y_i \in \mathcal{Y}.$$

This constraint can be further linearized in α and β using the fact that it is equivalent to:

$$\exists A_i \in \mathbb{S}_+^m, t \in \mathbb{R}_+, A_i \succeq \alpha\alpha^\top - \lambda_i I, \; t \geq \beta^2, \; y_i^\top A_i y_i + \left(2\lambda_i\widehat{y}_i - (2\beta + \eta)\alpha\right)^\top y_i + t - \lambda_i\|\widehat{y}_i\|_2^2 - s_i \leq 0 \quad \forall y_i \in \mathcal{Y}.$$

Because $\mathcal{Y} = \mathbb{R}^m$, the set of constraints indexed by y_i can be further written as the semi-definite constraint

$$\begin{bmatrix} A_i & (\beta + \eta/2)\alpha - \lambda_i\widehat{y}_i \\ (\beta + \eta/2)\alpha^\top - \lambda_i\widehat{y}_i^\top & s_i + \lambda\|\widehat{y}_i\|_2^2 - t \end{bmatrix} \succeq 0$$

while using Schur complement, we also obtain two additional linear matrix inequalities to capture the two other constraints on A_i and t:

$$\begin{bmatrix} \lambda_i I - A_i & \alpha \\ \alpha^\top & 1 \end{bmatrix} \succeq 0, \quad \begin{bmatrix} t & \beta \\ \beta & 1 \end{bmatrix} \succeq 0.$$

This completes the proof. $\qquad\square$

Proposition 4.6 (Mean-CVaR loss function). *Suppose that ℓ is the mean-CVaR loss function of the form* (6), *$\gamma \in \mathbb{R}_{++}$, $\varepsilon \in (0, 1]$ and $\rho > \rho_{\min}(x_0, \gamma, \varepsilon)$. Suppose in addition that $\mathcal{Y} = \mathbb{R}^m$ and $\mathbb{D}_{\mathcal{Y}}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$. The distributionally robust portfolio allocation model with side information* (22) *is equivalent to the semi-definite optimization problem*

$$
\begin{aligned}
\min \quad & \phi + (N\varepsilon)^{-1}\nu^+ - N^{-1}\nu^- \\
\text{s.t.} \quad & \alpha \in \mathcal{A}, \; \beta \in \mathcal{B}, \; \lambda \in \mathbb{R}_+^N, \; s \in \mathbb{R}^N \\
& \nu^+ \in \mathbb{R}_+, \; \nu^- \in \mathbb{R}_+, \; \phi \in \mathbb{R}, \; \varphi \in \mathbb{R}_+, \; \psi \in \mathbb{R}_+^N \\
& \phi - d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_1 \\
& \phi + d_i\varphi + \psi_i - s_i \geq 0 && \forall i \in \mathcal{I}_2 \\
& \nu^+ - \nu^- + \left(\sum_{i \in \mathcal{I}_1} d_i - N\rho\right)\varphi - \sum_{i \in [N]}\psi_i \geq 0 \\
& \varphi - \lambda_i \geq 0 && \forall i \in [N] \\
& \begin{bmatrix} \lambda_i I & \frac{\eta}{2}\alpha - \lambda_i\widehat{y}_i \\ \frac{\eta}{2}\alpha^\top - \lambda_i\widehat{y}_i^\top & s_i + \lambda_i\|\widehat{y}_i\|_2^2 - \beta \end{bmatrix} \succeq 0, \; \begin{bmatrix} \lambda_i I & \frac{\eta\tau+1}{2\tau}\alpha - \lambda_i\widehat{y}_i \\ \frac{\tau+1}{2\tau}\alpha^\top - \lambda_i\widehat{y}_i^\top & s_i + \lambda_i\|\widehat{y}_i\|_2^2 - (1 - 1/\tau)\beta \end{bmatrix} \succeq 0 \;\; \forall i \in [N],
\end{aligned}
$$

where the parameters d_i are defined as in (14).

Proof of Proposition 4.6. Notice that ℓ is a pointwise maximum of two linear functions of y_i. In this case, the epigraphical formulation of the constraint

$$\sup_{y_i \in \mathcal{Y}} \left\{\ell(y_i, \alpha, \beta) - \lambda_i\mathbb{D}_{\mathcal{Y}}(y_i, \widehat{y}_i)\right\} \leq s_i$$

is equivalent to the following set of two semi-infinite constraint

$$\begin{cases} -\eta y_i^\top\alpha + \beta - \lambda_i\|y_i - \widehat{y}_i\|_2^2 \leq s_i & \forall y_i \in \mathbb{R}^m \\ -(\eta + \frac{1}{\tau})y_i^\top\alpha + (1 - \frac{1}{\tau})\beta - \lambda_i\|y_i - \widehat{y}_i\|_2^2 \leq s_i & \forall y_i \in \mathbb{R}^m. \end{cases}$$

The proof now follows from the same line of argument as the proof of Proposition 4.5. $\qquad\square$

Both Proposition 4.5 and 4.6 leverage the fact that $\mathcal{Y} = \mathbb{R}^m$ in order to reformulate the problem using semi-definite constraint. In case that $\mathcal{Y}$ is an ellipsoid of the form $\mathcal{Y} = \{y \in \mathbb{R}^m : y^\top Q y + 2q^\top y + q_0 \leq 0\}$, then the S-lemma [52] can also be applied to devise a similar semi-definite optimization problem. These details are omitted.

Despite having a fractional objective function, the results in this section assert that distributionally robust portfolio allocation with side information problem can be solved efficiently using convex conic programming solvers. This is in stark contrast with existing results in distributionally robust fractional programming where

only *non*convex reformulations are available, and the optimal solution is obtained by solving a sequence of convex optimization problems after bisection [33, 64].

5. Numerical Experiment

In this section, we compare the empirical performance of our proposed conditional portfolio allocation model against several benchmark models. To this end, we conduct the numerical experiment using real historical data collected from (i) the US stock market, and (ii) the Chinese stock market. This choice of markets will help demonstrate the efficacy of our method when it is deployed on both a developed (US) and an emerging (Chinese) market.

We now formally describe how we model the return and the side information in the experiment.

Return: We use the historical S&P 500 constituents data[3] for US market and the China A-shares data[4] for the Chinese market in 2010-2016. To ensure that there is enough historical data to get a reliable comparison among different optimization models, we only consider the stocks listed on the market before 2010. We use the 1-day stock percentage return as our response variable Y. The sampling frequency of all the data is 1 sample per day.

Side information: We use the Fama-French three-factor model [25] to construct the covariate $X = (X_1, X_2, X_3) \in \mathbb{R}^3$. The first covariate X_1 is the market return (Rm-Rf) factor, X_2 is the "small minus big" (SMB) factor and X_3 is the "high minus low" (HML) factor. Each factor is defined as the most recently observed return of an associated portfolio:

- Rm-Rf factor: the portfolio weight on each stock is proportional to its market capitalization (MC);
- SMB factor: longs the stocks with MC in the lowest 50% quantile and shorts the stocks with MC in the highest 50% quantile;
- HML factor: longs the stocks with P/E ratio in the highest 30% quantile and shorts the stocks with P/E ratio in the lowest 30% quantile.

For the US stock market, we download the factor data from the data library of Kenneth French's homepage.[5] For the Chinese market, we use raw data including price earning ratio (P/E ratio) and market capitalization to compute the covariate X.

Models: We first determine the feasible region of the portfolio optimization problem. If a stock in the sampled stock pool is not tradable at day t, then we set its weight to be zero. In addition, as the Chinese stock market has strong regulations on short selling, we enforce the investment weight vector α to be non-negative. To summarize, we set the feasible region of the stock weight to be a subset of the the non-negative simplex

$$\mathcal{A} = \left\{ \alpha \in \mathbb{R}_+^m \mid \sum_{i \in [m]} \alpha_i = 1, \ \alpha_j = 0 \text{ if stock } j \text{ is not tradable} \right\}.$$

We will compare the performance of the following six portfolio allocation models:

(i) the Equal Weighted model (EW): in this case, the portfolio allocation is given by the minimizer of the optimization problem

$$\min_{\alpha \in \mathcal{A}} \ \|\alpha\|_2. \tag{EW}$$

If all stocks are tradable, then the EW allocation coincides with the $1/m$-portfolio, which is renowned for its robustness [21]. This portfolio is parameter-free, and was previously shown to be the limit of distributionally robust portfolio allocations as the ambiguity size increases [50].

---

[3]The daily sampled S&P 500 data is downloaded from the Wharton Research Data Services: `https://wrds-www.wharton.upenn.edu`.

[4]China A-shares are the stock shares of mainland China-based companies that trade on the two major Chinese stock exchanges, the Shanghai Stock Exchange (SSE) and the Shenzhen Stock Exchange (SZSE). The Chinese stock market data is downloaded using `https://github.com/waditu/tushare`.

[5]Available for download from `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library`.

(ii) the unconditional Mean Variance model (MV) [43]: the portfolio allocation solves the following optimization problem

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathbb{R}} \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(Y, \alpha, \beta)] \tag{MV}$$

with the loss function ℓ prescribed in (5), and the distribution $\widehat{\mathbb{P}}$ is the empirical distribution supported on the available return data.

(iii) the Distributionally Robust unconditional Mean Variance model (DRMV): the portfolio allocation solves

$$\min_{\alpha \in \mathcal{A}} \ \sqrt{\alpha^\top \mathrm{Variance}_{\widehat{\mathbb{P}}}(Y)\alpha} - \eta \cdot \alpha^\top \mathbb{E}_{\widehat{\mathbb{P}}}[Y] + \sqrt{(1+\eta^2)\rho}\|\alpha\|_2. \tag{DRMV}$$

The above optimization problem is the reformulation of the distributionally robust mean variance portfolio allocation with a Wasserstein ambiguity set on the distributions of the asset returns Y [10]. The tuning parameter for this method is $\rho \in \{0.1, 0.2, 0.5\}$.

(iv) the Conditional Mean Variance model (CMV) with the loss function ℓ prescribed in (5), and the allocation is the solution of

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathbb{R}} \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)]. \tag{CMV}$$

The parameter γ is set to the a-quantile of the empirical distribution of the distance between x_0 and the the training covariate vectors[6], where the quantile value is in the range $a \in \{10\%, 20\%, 50\%\}$. Notice that by the choice of γ, the empirical distribution satisfies $\widehat{\mathbb{P}}(X \in \mathcal{N}_\gamma(x_0)) > 0$ and the conditional expectation is well-defined.

(v) the Distributionally Robust Conditional Mean-Variance model (DRCMV) with type-∞ Wasserstein ambiguity set [48], in which the portfolio allocation solves

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathbb{R}} \sup_{\mathbb{Q} \in \mathbb{B}_\rho^\infty, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) > 0} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)]. \tag{DRCMV}$$

The parameter γ is set to the a-quantile of the empirical distribution of the distance between x_0 and the training covariate vectors, where the quantile value is in the range $\{10\%, 20\%, 50\%\}$. The radius ρ is set to the b-quantile of the empirical distribution of the distance between x_0 and the training covariate vectors, where the quantile value is in the range $b \in \{5\%, 10\%, 25\%\}$. Using the result from [48, Proposition 2.5], this model can be reformulated as a second-order cone program (see Appendix A.3 for definition of $\mathbb{B}_\rho^\infty$ and detailed derivation of reformulation).

(vi) the Optimal Transport based (distributionally robust) Conditional Mean-Variance model (OTCMV) where the portfolio allocation is the solution to problem (11) with the mean-variance loss function (5)

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathbb{R}} \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X = x_0) \geq \varepsilon} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X = x_0]. \tag{OTCMV}$$

The tuning parameters for this model includes the probability bound $\varepsilon \in \{0.1, 0.2, 0.5\}$ and the radius $\rho = a \times \rho_{\min}$, where $a \in \{1.1, 1.2, 1.5\}$ and $\rho_{\min}$ denotes the minimum distance between the training covariate $(\widehat{x}_i)_{i=1}^N$ and x_0. This model is equivalent to a second-order cone program thanks to Proposition 3.3.

Notice that the feasible set $\mathcal{B}$ is set to $\mathbb{R}$ in all methods. The ground metrics are chosen with $\mathbb{D}_{\mathcal{X}}(x, \widehat{x}) = \|x - \widehat{x}\|_2^2$ and $\mathbb{D}_{\mathcal{Y}}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$. All models can be solved using the MOSEK quadratic program solver [45].[7]

Experiments: We carry out 256 independent replications to ensure that the performance comparison is statistically significant. In each replication, we randomly sample $m = 20$ stocks from the data set to be used as a *stock pool*. The training procedure is carried out as follows. At each trading day t, we apply different portfolio optimization models to construct portfolio of the stock pool, using the historical observation of (X, Y) in the past 2-year window to form the nominal distribution $\widehat{\mathbb{P}}$ (more precisely, we use 252×2 observations precedent to t). Moreover, the covariate x_0 is chosen as the observation of X at time t. To obtain the

---

[6]More precisely, this is the empirical distribution supported on $\Delta_i = \mathbb{D}_{\mathcal{X}}(\widehat{x}_i, x_0)$ for $i = 1, \ldots, N$.

[7]The codes are available at `http://github.com/AndyZhang92/DR-Conditional-Port-Opt`.
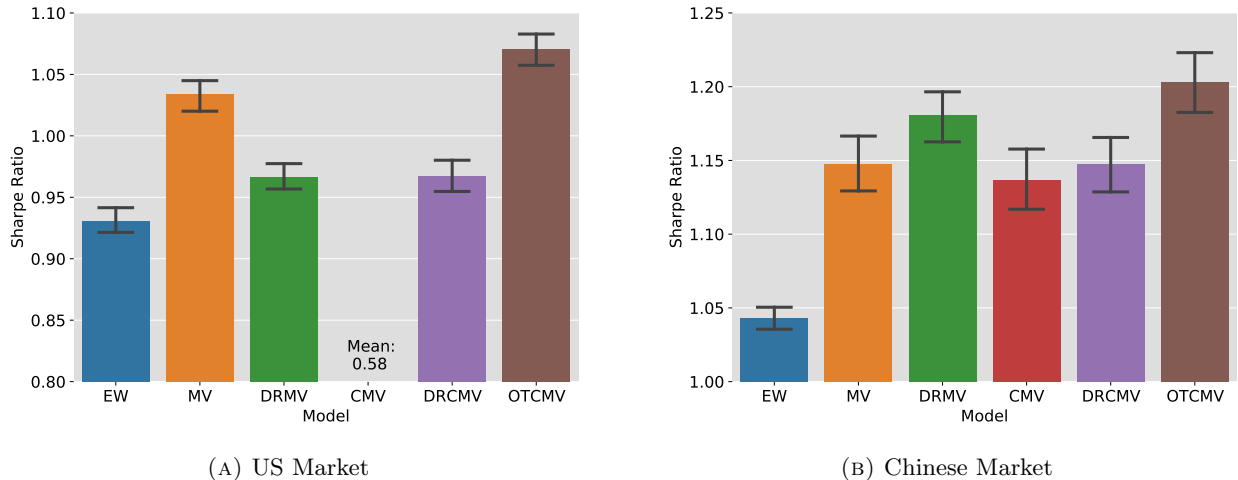
(A) US Market

(B) Chinese Market

FIGURE 2. Barplots of annualized Sharpe ratios in (A) US market and (B) Chinese market.

validation score, we compute the return for each time t in the period between January 1, 2012 and December 31, 2013, and then calculate the realized Sharpe ratio as the validation score. We first tune the parameter η of the loss function ℓ by tuning the MV model (EW) with $\eta \in [10^{-1}, 2 \cdot 10^1]$ with 7 equidistant points in the logscale to maximize the validation score of the MV model. This selected η is reused in all models: MV, DRMV, CMV, DRCMV and OTCMV. Subsequently, we tune the corresponding parameters for DRMV, CMV, DRCMV and OTCMV by selecting the hyper-parameters that maximize the validation score.

Test Results: To obtain the out-of-sample performance, we deploy the tuned model in a rolling horizon scheme on the testing data from January 1, 2014 to December 31, 2016. We report the experimental results of different models using Sharpe ratios of the realized returns, which are usually used to compare the performance of different trading strategies. For each sampled experiment and a given portfolio optimization model, we denote the daily percentage return of the model by r_i for $i \in [T]$, where T is the total number of trading days in the test period. The annualized Sharpe ratio is computed using the following formula:

$$\text{annualized Sharpe ratio} \triangleq \sqrt{252} \times \left(\widehat{\text{mean}}(\{r_i\}_{i \in [T]}) \Big/ \widehat{\text{std}}(\{r_i\}_{i \in [T]}) \right),$$

where $\widehat{\text{mean}}$ and $\widehat{\text{std}}$ denote the empirical mean and standard deviation estimator, respectively. Thus, for each model we compute 256 annualized Sharpe ratios obtained from 256 independent experiments.

In Figure 2, we present the bar charts of the annualized Sharpe ratios for both the US and the Chinese market. The average annualized Sharpe ratios across 256 experiments are represented by the height of the bar, with error bars depicting the standard errors. From Figure 2 we can remark that OTCMV achieves the best Sharpe ratio performance in both markets. To test whether the mean annualized Sharpe ratio of OTCMV is larger than another model, we conduct the two-sample one-sided t-test with the hypotheses:

$$\mathcal{H}_{\text{null}} : \text{Sharpe ratio of OTCMV is smaller}, \qquad \mathcal{H}_{\text{alt}} : \text{Sharpe ratio of OTCMV is bigger}.$$

The p-values for the tests are reported in Table 1. If we choose a significance level of 0.05, then the null hypothesis is rejected in the US market for all competing approaches. In the Chinese market, the only null hypothesis that the test fails to reject is the comparison between OTCMV and DRMV (middle column).

In order to demonstrate that our proposed method is capable of producing robust portfolio returns, the bottom 5%-quantile of the daily realized returns $\{r_i\}_{i \in [T]}$ is also reported in Figure 3. As we are viewing the bottom 5% return, the barplots fall on the negative side. In both the Chinese and US markets, the 5%-quantile of OTCMV's daily returns is remarkably higher than its counterparts, closely followed by MV and the other two distributionally robust methods DRMV and DRCMV.

Model	EW	MV	DRMV	CMV	DRCMV
US Market	< 0.001	0.037	< 0.001	< 0.001	< 0.001
Chinese Market	< 0.001	0.044	0.397	0.021	0.045

TABLE 1. p-values of two sample t-test of whether the mean Sharpe ratio of OTCMV is larger than the competing approaches.

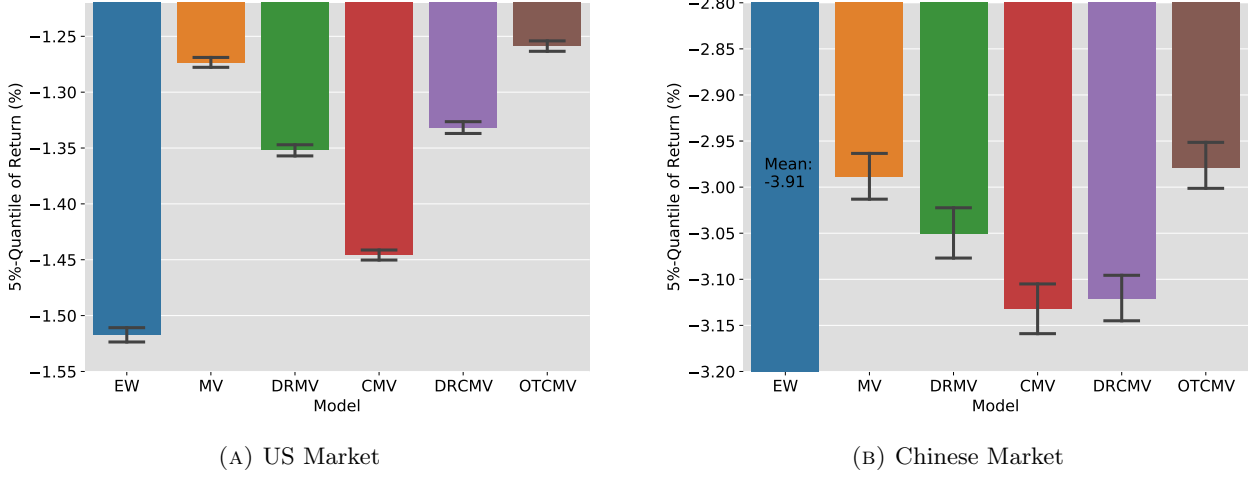(A) US Market

(B) Chinese Market

FIGURE 3. Barplots of 5%-quantile of daily return in (A) US market and (B) Chinese market. The 5%-quantile for model EW in the Chinese market is significantly lower than other models, thus we only report its value in (B).

APPENDIX A. SUPPLEMENTARY MATERIAL

A.1. **Minimax result**

The following results are needed to justify Lemmas 2.2 and 2.3.

Lemma A.1 (Interchange). *Suppose that $\mathcal{Y}$ and $\mathcal{B}$ are compact and that ℓ is continuous in Y and convex in β. Then we have*

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \inf_{\beta \in \mathcal{B}} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)] = \inf_{\beta \in \mathcal{B}} \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)].$$

The proof of Lemma A.1 follows from the convexity of the conditional ambiguity set. To this end, define the following ambiguity set

$$\mathcal{B}_{x_0, \gamma, \varepsilon}(\mathbb{B}_\rho) \triangleq \left\{ \mu_{x_0} \in \mathcal{M}(\mathcal{Y}) : \begin{array}{l} \exists \mathbb{Q} \in \mathbb{B}_\rho \text{ such that } \mathbb{Q}(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq \varepsilon \\ \mathbb{Q}(\mathcal{N}_\gamma(x_0) \times A) = \mu_{x_0}(A)\mathbb{Q}(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \, \forall A \subseteq \mathcal{Y} \text{ measurable} \end{array} \right\}.$$

Lemma A.2 (Convexity of $\mathcal{B}_{x_0, \varepsilon, \gamma}(\mathbb{B}_\rho)$). *If $\mathbb{B}_\rho$ is convex, then the conditional ambiguity set $\mathcal{B}_{x_0, \gamma, \varepsilon}(\mathbb{B}_\rho)$ is also convex.*

Proof of Lemma A.2. Let μ_0^0, $\mu_0^1 \in \mathcal{B}_{x_0,\gamma,\varepsilon}(\mathbb{B}_\rho)$ be two arbitrary probability measures supported on $\mathcal{Y}$. Associated with each μ_0^j, $j \in \{0, 1\}$, is a corresponding joint measure $\mathbb{Q}^j \in \mathcal{M}(\mathcal{X} \times \mathcal{Y})$ such that

$$\mathbb{Q}^j(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq \varepsilon \quad \text{and} \quad \frac{\mathbb{Q}^j(\mathcal{N}_\gamma(x_0) \times A)}{\mathbb{Q}^j(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})} = \mu_0^j(A) \quad \forall A \subseteq \mathcal{Y} \text{ measurable.}$$

Select any $\lambda \in (0, 1)$. We now show that $\mu_0^\lambda = \lambda \mu_0^1 + (1 - \lambda)\mu_0^0 \in \mathcal{B}_{x_0,\gamma,\varepsilon}(\mathbb{B}_\rho)$. Indeed, consider the joint measure

$$\mathbb{Q}^\lambda = \theta \mathbb{Q}^1 + (1 - \theta)\mathbb{Q}^0,$$

where θ is defined as

$$\theta = \frac{\lambda \mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})}{\lambda \mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) + (1 - \lambda)\mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})} \in [0, 1].$$

By definition, we have $\mathbb{Q}^\lambda(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) \geq \varepsilon$, and by the convexity of $\mathbb{B}_\rho$, we have $\mathbb{Q}^\lambda \in \mathbb{B}_\rho$. Moreover, for any set $A \subseteq \mathcal{Y}$ measurable, we find

$$\frac{\mathbb{Q}^\lambda(\mathcal{N}_\gamma(x_0) \times A)}{\mathbb{Q}^\lambda(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})} = \frac{\theta \mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times A) + (1 - \theta)\mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times A)}{\theta \mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times \mathcal{Y}) + (1 - \theta)\mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})}$$

$$= \frac{\lambda \mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})\mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times A) + (1 - \lambda)\mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})\mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times A)}{\mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})\mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})}$$

$$= \frac{\lambda \mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times A)}{\mathbb{Q}^1(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})} + \frac{(1 - \lambda)\mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times A)}{\mathbb{Q}^0(\mathcal{N}_\gamma(x_0) \times \mathcal{Y})}$$

$$= \lambda \mu_0^1(A) + (1 - \lambda)\mu_0^0(A),$$

where the second equality holds thanks to the definition of θ. This line of argument implies that $\mu_0^\lambda \in \mathcal{B}_{x_0,\gamma,\varepsilon}(\mathbb{B}_\rho)$, and further asserts the convexity of $\mathcal{B}_{x_0,\gamma,\varepsilon}(\mathbb{B}_\rho)$. □

We are now ready to prove Lemma A.1.

Proof of Lemma A.1. By rewriting the conditional expectation using the conditional measure μ_0, we have for any value of α

$$\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \inf_{\beta \in \mathcal{B}} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta)|X \in \mathcal{N}_\gamma(x_0)] = \sup_{\mu_0 \in \mathcal{B}_{x_0,\gamma,\varepsilon}(\mathbb{B}_\rho)} \inf_{\beta \in \mathcal{B}} \mathbb{E}_{\mu_0}[\ell(Y, \alpha, \beta)].$$

The set $\mathcal{B}_{x_0,\gamma,\varepsilon}(\mathbb{B}_\rho)$ is convex by Lemma A.2. Moreover, because ℓ is continuous in Y, the mapping $\mu_0 \mapsto \mathbb{E}_{\mu_0}[\ell(Y, \alpha, \beta)]$ is upper semicontinuous in the weak topology thanks to the compactness of $\mathcal{Y}$. Finally, $\mathcal{B}$ is compact and ℓ is convex in β. The interchangeability of the supremum and the infimum operators is now a consequence of the Sion's minimax theorem [60]. □

Proof of Lemma 2.2. It is well-known that

$$\text{Variance}_{\mathbb{Q}}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)] = \min_{\beta \in \mathbb{R}} \mathbb{E}_{\mathbb{Q}}[(Y^\top \alpha - \beta)^2 | X \in \mathcal{N}_\gamma(x_0)]$$

for any probability measure $\mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y})$, where the optimal β is the conditional mean of $Y^\top \alpha$ given $X \in \mathcal{N}_\gamma(x_0)$ under $\mathbb{Q}$. When $\mathcal{A}$ and $\mathcal{Y}$ are compact, the random variable $Y^\top \alpha$ has bounded mean for any $\mathbb{Q} \in \mathbb{B}_\rho$. More precisely, we have

$$\mathbb{E}_{\mathbb{Q}}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)] \in \mathcal{B},$$

where $\mathcal{B}$ is defined as in the statement of the lemma. Therefore, it is without any loss of optimality to restrict $\beta \in \mathcal{B}$. We thus find

$$\min_{\alpha \in \mathcal{A}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \text{Variance}_\mathbb{Q}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)] - \eta \cdot \mathbb{E}_\mathbb{Q}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)]$$

$$= \min_{\alpha \in \mathcal{A}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \min_{\beta \in \mathbb{R}} \quad \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]$$

$$= \min_{\alpha \in \mathcal{A}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \min_{\beta \in \mathcal{B}} \quad \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]$$

$$= \min_{\alpha \in \mathcal{A}} \min_{\beta \in \mathcal{B}} \quad \sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) \geq \varepsilon} \quad \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)],$$

where the last equality follows from Lemma A.1. This completes the proof. $\qquad \square$

Proof of Lemma 2.3. It is well-known that

$$\text{CVaR}_\mathbb{Q}^{1-\tau}[Y^\top \alpha | X \in \mathcal{N}_\gamma(x_0)] = \min_{\beta \in \mathbb{R}} \ \mathbb{E}_\mathbb{Q}[\beta + \frac{1}{\tau}(-Y^\top \alpha - \beta)^+ | X \in \mathcal{N}_\gamma(x_0)]$$

for any probability measure $\mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y})$, where the optimal β is the conditional $(1 - \tau)$-quantile of $Y^\top \alpha$ given $X \in \mathcal{N}_\gamma(x_0)$ under $\mathbb{Q}$. When $\mathcal{A}$ and $\mathcal{Y}$ are compact, the random variable $Y^\top \alpha$ has uniformly bounded support $\mathcal{B}$ for any $\mathbb{Q} \in \mathbb{B}_\rho$ and $\alpha \in \mathcal{A}$, where $\mathcal{B}$ is defined as in the statement of the lemma. Therefore, it is without any loss of optimality to restrict $\beta \in \mathcal{B}$. The rest of the proof follows the same argument in the proof of Lemma 2.2. $\qquad \square$

A.2. Auxiliary Results

The results in this section is used in the proofs of Section 3. For any $N \in \mathbb{N}$, given some $c \in \mathbb{R}^N$, $d \in \mathbb{R}_+^N$ and $\varepsilon \in (0, 1)$, consider the following two optimization problems

$$\begin{cases} \sup & \dfrac{\sum_i v_i c_i}{N \sum_i v_i} \\ \text{s.t.} & v \in [0, 1]^N \\ & \displaystyle\sum_i v_i \geq \varepsilon N, \ \sum_i v_i d_i \leq \rho \end{cases} \tag{23a}$$

and

$$\begin{cases} \sup & \dfrac{\sum_i v_i c_i}{N \sum_i v_i} \\ \text{s.t.} & v \in [0, 1]^N \\ & \displaystyle\sum_i v_i = \varepsilon N, \ \sum_i v_i d_i \leq \rho, \end{cases} \tag{23b}$$

where the summations are taken over $i \in [N]$. Notice that the summation constraint of v_i in (23a) is an *in*equality constraint, while in (23b) it is an equality constraint. The next result asserts that the inequality in (23a) can be strengthened to an equality as in (23b) without any loss of optimality.

Lemma A.3. *The optimal values of problems* (23a) *and* (23b) *are equal.*

Proof of Lemma A.3. Let $\bar{v} \in [0, 1]^N$ be such that $\sum_i \bar{v}_i > \varepsilon N$ and $\sum_i \bar{v}_i d_i \leq \rho$. One can construct $v' = (\varepsilon N / \sum_i \bar{v}_i)\bar{v}$ which satisfies $v' \in [0, 1]^N$, $\sum_i v'_i = \varepsilon N$, and $\sum_i v'_i d_i \leq \rho$ since $d \geq 0$. Furthermore, it reaches the same objective value

$$\frac{1}{N \sum_i v'_i} \sum_i v'_i c_i = \frac{1}{\varepsilon N^2} \sum_i \frac{\varepsilon N}{\sum_i \bar{v}_i} \bar{v}_i c_i = \frac{1}{N \sum_i \bar{v}_i} \sum_i \bar{v}_i c_i,$$

which finishes the proof. $\qquad \square$

Proposition A.4 (Equivalent representation). *Given $\varepsilon \in (0, 1]$ and $\rho > \rho_{\min}(x_0, 0, \varepsilon)$, then for any feasible solution (α, β), we have*

$$
\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X = x_0) \geq \varepsilon} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X = x_0] =
\begin{cases}
\sup & (N\varepsilon)^{-1} \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)] \\
\text{s.t.} & v \in [0, 1]^N, \ \mu_0^i \in \mathcal{M}(\mathcal{Y}) \quad \forall i \in [N] \\
& \sum_{i \in [N]} v_i = N\varepsilon \\
& \sum_{i \in [N]} v_i \big(\mathbb{D}_\mathcal{X}(x_0, \widehat{x}_i) + \mathbb{E}_{\mu_0^i}[\mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \big) \leq N\rho.
\end{cases}
$$

Proof of Proposition A.4. By exploiting the definition of the optimal transport cost and the fact that any joint probability measure $\pi \in \Pi(\mathbb{Q}, \widehat{\mathbb{P}})$ can be written as $\pi = N^{-1} \sum_{i \in [N]} \pi_i \otimes \delta_{(\widehat{x}_i, \widehat{y}_i)}$, where each π_i is a probability measure on $\mathcal{X} \times \mathcal{Y}$, we have

$$
\sup_{\mathbb{Q} \in \mathbb{B}_\rho, \mathbb{Q}(X = x_0) \geq \varepsilon} \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X = x_0] =
\begin{cases}
\sup & \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X = x_0] \\
\text{s.t.} & \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) \quad \forall i \in [N] \\
& \sum_{i \in [N]} \pi_i(\{x_0\} \times \mathcal{Y}) \geq N\varepsilon \\
& \mathbb{Q} = N^{-1} \sum_{i \in [N]} \pi_i, \ \sum_{i \in [N]} \mathbb{W}(\pi_i, \delta_{(\widehat{x}_i, \widehat{y}_i)}) \leq N\rho
\end{cases}
$$

$$
=
\begin{cases}
\sup_{v \in \mathcal{U}} \ \sup & \mathbb{E}_\mathbb{Q}[\ell(Y, \alpha, \beta) | X = x_0] \\
\text{s.t.} & \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) & \forall i \in [N] \\
& \pi_i(\{x_0\} \times \mathcal{Y}) = v_i & \forall i \in [N] \\
& \mathbb{Q} = N^{-1} \sum_{i \in [N]} \pi_i, \ \sum_{i \in [N]} \mathbb{W}(\pi_i, \delta_{(\widehat{x}_i, \widehat{y}_i)}) \leq N\rho,
\end{cases}
$$

where the set $\mathcal{U}$ is defined as

$$
\mathcal{U} \triangleq \left\{ v \in [0, 1]^N : \sum_{i \in [N]} v_i \geq N\varepsilon \right\}.
$$

Define the following two functions $g, h : \mathcal{A} \to \mathbb{R}$ as

$$
g(v) \triangleq
\begin{cases}
\sup & \big(\sum_{i \in [N]} v_i \big)^{-1} \sum_{i \in [N]} \int_\mathcal{Y} \ell(y, \alpha, \beta) \pi_i(\{x_0\} \times \mathrm{d}y) \\
\text{s.t.} & \pi_i \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) & \forall i \in [N] \\
& \pi_i(\{x_0\} \times \mathcal{Y}) = v_i & \forall i \in [N] \\
& \sum_{i \in [N]} \mathbb{W}(\pi_i, \delta_{(\widehat{x}_i, \widehat{y}_i)}) \leq N\rho
\end{cases}
\tag{24}
$$

and

$$
h(v) \triangleq
\begin{cases}
\sup & \big(\sum_{i \in [N]} v_i \big)^{-1} \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)] \\
\text{s.t.} & \mu_0^i \in \mathcal{M}(\mathcal{Y}) & \forall i \in [N] \\
& \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_\mathcal{X}(x_0, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq N\rho.
\end{cases}
\tag{25}
$$

We can show that $\sup_{v \in \mathcal{U}} g(v) = \sup_{v \in \mathcal{U}} h(v)$. First, to show $\sup_{v \in \mathcal{U}} g(v) \leq \sup_{v \in \mathcal{U}} h(v)$, we fix an arbitrary value $v \in \mathcal{U}$. For any $\{\pi_i\}_{i \in [N]}$ that is feasible for (24), define $\mu_0^i \in \mathcal{M}(\mathcal{Y})$ such that

$$
v_i \mu_0^i(S) = \pi_i(\{x_0\} \times S) \qquad \forall S \subseteq \mathcal{Y} \text{ measurable} \qquad \forall i \in [N].
$$

One can verify that $\{\mu_0^i\}_{i \in [N]}$ is a feasible solution to (25), notably because

$$
\sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_\mathcal{X}(x_0, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] \leq \sum_{i \in [N]} \mathbb{E}_{\pi_i}[\mathbb{D}_\mathcal{X}(X, \widehat{x}_i) + \mathbb{D}_\mathcal{Y}(Y, \widehat{y}_i)] = \sum_{i \in [N]} \mathbb{W}(\pi_i, \delta_{(\widehat{x}_i, \widehat{y}_i)}) \leq N\rho,
$$

where the first inequality follows from the non-negativity of $\mathbb{D}(\cdot, \widehat{x})$ and $\mathbb{D}(\cdot, \widehat{y})$, and the second inequality is from the feasibility of $\{\pi_i\}_{i \in [N]}$ in (24). Moreover, the optimal value of $\{\pi_i\}_{i \in [N]}$ in (24) and the optimal

value of $\{\mu_0^i\}_{i\in[N]}$ in (25) coincide because

$$\int_{\mathcal{Y}} \ell(y,\alpha,\beta)\pi_i(\{x_0\}\times \mathrm{d}y) = \upsilon_i\int_{\mathcal{Y}} \ell(y,\alpha,\beta)\mu_0^i(\mathrm{d}y) = \upsilon_i\mathbb{E}_{\mu_0^i}[\ell(Y,\alpha,\beta)] \qquad \forall i\in[N].$$

This implies that $g(\upsilon) \le h(\upsilon)$ for any $\upsilon\in\mathcal{U}$, and thus we have

$$\sup_{\upsilon\in\mathcal{U}} g(\upsilon) \le \sup_{\upsilon\in\mathcal{U}} h(\upsilon). \tag{26}$$

Next, we will establish the reverse direction of the inequality in (26). To this end, consider any $\{\mu_0^i\}_{i\in[N]}$ that is feasible for (25), we will construct explicitly a sequence of probability families $\{\pi_{i,k}\}_{i\in[N],k\in\mathbb{N}}$ that is feasible for (24) and attains the same objective value in the limit as $k\to\infty$.

In doing so, we start by supposing that $\sum_{i\in[N]} \upsilon_i\mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{X}}(x_0,\widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_i)] < N\rho$. Let us define the measures $\{\pi_i\}_{i\in[N]}$ as

$$\pi_i = \upsilon_i\delta_{x_0}\otimes\mu_0^i + (1-\upsilon_i)\delta_{(x_i,\widehat{y}_i)},$$

for some $x_i\in\mathcal{X}\setminus\{\widehat{x}_i\}$ that satisfy

$$\sum_{i\in[N]} (1-\upsilon_i)\mathbb{D}_{\mathcal{X}}(x_i,\widehat{x}_i) \le N\rho - \sum_{i\in[N]} \upsilon_i\big(\mathbb{D}_{\mathcal{X}}(x_0,\widehat{x}_i) + \mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_i)]\big).$$

Notice that under the condition of this case, the right-hand side is strictly positive, and the existence of such x_i's is guaranteed thanks to the continuity of $\mathbb{D}_{\mathcal{X}}$ and the fact that $\mathbb{D}_{\mathcal{X}}(\widehat{x},\widehat{x}) = 0$. It is now easy to verify that $\{\pi_i\}_{i\in[N]}$ is feasible for (24), and moreover, the objective value of $\{\pi_i\}_{i\in[N]}$ in (24) amounts to

$$\big(\sum_{i\in[N]} \upsilon_i\big)^{-1}\sum_{i\in[N]}\int_{\mathcal{Y}} \ell(y,\alpha,\beta)\pi_i(\{x_0\}\times\mathrm{d}y) = \big(\sum_{i\in[N]} \upsilon_i\big)^{-1}\sum_{i\in[N]}\Big[\upsilon_i\mathbb{E}_{\mu_0^i}[\ell(Y,\alpha,\beta)]\Big].$$

The case where $\sum_{i\in[N]} \upsilon_i\mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{X}}(x_0,\widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_i)] = N\rho$ is more complex. In particular, we start by focusing on the situation where there exists some $\widehat{i}$ for which $\upsilon_{\widehat{i}}\ne 0$ and $\mu_0^{\widehat{i}}\ne\delta_{\widehat{y}_{\widehat{i}}}$. Here, we can construct a sequence of measure $\{\pi_{i,k}\}_{i\in[N],k\in\mathbb{N}}$ as

$$\pi_{i,k} = \begin{cases} (\upsilon_i-\gamma_k)\delta_{x_0}\otimes\mu_0^i + \gamma_k\delta_{(x_0,\widehat{y}_i)} + (1-\upsilon_i)\delta_{(x_{i,k},\widehat{y}_i)} & \text{if } i=\widehat{i}, \\ \upsilon_i\delta_{x_0}\otimes\mu_0^i + (1-\upsilon_i)\delta_{(x_{i,k},\widehat{y}_i)} & \text{otherwise,} \end{cases}$$

for some $\gamma_k\in[0,\upsilon_{\widehat{i}}]$, $\lim_{k\to\infty}\gamma_k = 0$, and some $x_{i,k}\in\mathcal{X}\setminus\{\widehat{x}_i\}$

$$x_{i,k} \xrightarrow{k\to\infty} \widehat{x}_i \quad\forall k.$$

Furthermore, let the sequences γ_k and $x_{i,k}$ satisfy for any $k\in\mathbb{N}$

$$\sum_{i\in[N]} (1-\upsilon_i)\mathbb{D}_{\mathcal{X}}(x_{i,k},\widehat{x}_i) \le \gamma_k\mathbb{E}_{\mu_0^{\widehat{i}}}[\mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_{\widehat{i}})].$$

Notice that under the condition of this case, the right-hand side is strictly positive, and the existence of the sequence $(x_{i,k},\gamma_k)$ is again guaranteed thanks to the continuity of $\mathbb{D}_{\mathcal{X}}$ and the fact that $\mathbb{D}_{\mathcal{X}}(\widehat{x},\widehat{x}) = 0$. It is now easy to verify that for any k, $\{\pi_{i,k}\}_{i\in[N]}$ is feasible for (24):

$$\sum_{i\in[N]} \mathrm{W}(\pi_{i,k},\delta_{(\widehat{x}_i,\widehat{y}_i)}) = \sum_{i\in[N]} \upsilon_i(\mathbb{D}_{\mathcal{X}}(x_0,\widehat{x}_i)) + \mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_i)]) + (1-\upsilon_i)\mathbb{D}_{\mathcal{X}}(x_{i,k},\widehat{x}_i) - \gamma_k\mathbb{E}_{\mu_0^{\widehat{i}}}[\mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_{\widehat{i}})]$$

$$\le \sum_{i\in[N]} \upsilon_i(\mathbb{D}_{\mathcal{X}}(x_0,\widehat{x}_i) + \mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{Y}}(Y,\widehat{y}_i)])) \le N\rho.$$

Moreover, the objective value of $\{\pi_{i,k}\}_{i\in[N]}$ in (24) amounts to

$$\big(\sum_{i\in[N]} \upsilon_i\big)^{-1}\left(\sum_{i\in[N]} \upsilon_i\mathbb{E}_{\mu_0^i}[\ell(Y,\alpha,\beta)] + \gamma_k(\ell(\widehat{y}_i,\alpha,\beta) - \mathbb{E}_{\mu_0^{\widehat{i}}}[\ell(Y,\alpha,\beta)])\right) \xrightarrow{k\to\infty} \big(\sum_{i\in[N]} \upsilon_i\big)^{-1}\sum_{i\in[N]} \upsilon_i\mathbb{E}_{\mu_0^i}[\ell(Y,\alpha,\beta)],$$

where the limit holds because $\lim_{k\to\infty}\gamma_k = 0$.

In the final case, we have $\mu_0^i = \delta_{\widehat{y}_i}$ for all $i \in [N]$ and $\sum_{i \in [N]} v_i \kappa_i = N\rho$. Since we have assumed that $\rho > \rho_{\min}(x_0, 0, \varepsilon)$, it must be that there exists an $\widehat{i}$ for which $v_{\widehat{i}} > 0$ and $\widehat{x}_{\widehat{i}} \neq x_0$. Indeed, otherwise we would have that $N\rho = \sum_{i \in [N]} v_i \kappa_i = 0 \leq N\rho_{\min}(x_0, 0, \varepsilon)$ which is a contradiction. Now let us one final time construct a sequence of measures

$$\pi_{i,k} = \begin{cases} (v_i - \gamma_k)\delta_{(x_0, \widehat{y}_i)} + (1 - v_i + \gamma_k)\delta_{(\widehat{x}_i, \widehat{y}_i)} & \text{if } i = \widehat{i}, \\ v_i \delta_{(x_0, \widehat{y}_i)} + (1 - v_i)\delta_{(\widehat{x}_i, \widehat{y}_i)} & \text{if } \widehat{x}_i \neq x_0 \\ v_i \delta_{(x_0, \widehat{y}_i)} + (1 - v_i)\delta_{(x_{i,k}, \widehat{y}_i)} & \text{otherwise,} \end{cases}$$

for some $\gamma_k \in (0, v_{\widehat{i}})$ and some $x_{i,k} \in \mathcal{X} \setminus \{x_0\}$, such that $x_{i,k} \xrightarrow{k \to \infty} \widehat{x}_i$, and that

$$\gamma_k \mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_{\widehat{i}}) \geq \sum_{i : \widehat{x}_i = x_0} (1 - v_i)\mathbb{D}_{\mathcal{X}}(x_0, x_{i,k}).$$

Indeed, it is easy to verify that $\{\pi_{i,k}\}_{i \in [N]}$ is always feasible for (24), and moreover, the objective value of $\{\pi_{i,k}\}_{i \in [N]}$ in (24) amounts to

$$(\sum_{i \in [N]} v_i)^{-1} \left(\sum_{i \in [N]} v_i \ell(\widehat{y}_i, \alpha, \beta) - \gamma_k \ell(\widehat{y}_{\widehat{i}}, \alpha, \beta) \right)$$
$$\xrightarrow{k \to \infty} (\sum_{i \in [N]} v_i)^{-1} \sum_{i \in [N]} v_i \ell(\widehat{y}_i, \alpha, \beta) = (\sum_{i \in [N]} v_i)^{-1} \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)],$$

where the limit holds because $\lim_{k \to \infty} \gamma_k \to 0$.

Combining the three cases, we can establish that

$$\sup_{v \in \mathcal{U}} g(v) \geq \sup_{v \in \mathcal{U}} h(v),$$

and by considering the above inequality along with (26), we can claim that

$$\sup_{v \in \mathcal{U}} g(v) = \sup_{v \in \mathcal{U}} h(v).$$

One now can rewrite

$$\sup_{v \in \mathcal{U}} h(v) = \begin{cases} \sup_{\mu_0^i \in \mathcal{M}(\mathcal{Y}) \ \forall i} \ \sup \ (\sum_{i \in [N]} v_i)^{-1} \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)] \\ \quad \text{s.t.} \quad v \in [0,1]^N \\ \qquad\qquad \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(Y, \widehat{y}_i)] \leq N\rho \\ \qquad\qquad \sum_{i \in [N]} v_i \geq N\varepsilon. \end{cases} \tag{27a}$$

$$= \begin{cases} \sup_{\mu_0^i \in \mathcal{M}(\mathcal{Y}) \ \forall i} \ \sup \ (\sum_{i \in [N]} v_i)^{-1} \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)] \\ \quad \text{s.t.} \quad v \in [0,1]^N \\ \qquad\qquad \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(Y, \widehat{y}_i)] \leq N\rho \\ \qquad\qquad \sum_{i \in [N]} v_i = N\varepsilon. \end{cases} \tag{27b}$$

$$= \begin{cases} \sup_{v \in [0,1]^N, \sum_{i \in [N]} v_i = N\varepsilon} \ \sup \ (\sum_{i \in [N]} v_i)^{-1} \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\ell(Y, \alpha, \beta)] \\ \quad \text{s.t.} \quad \mu_0^i \in \mathcal{M}(\mathcal{Y}) \quad \forall i \in [N] \\ \qquad\qquad \sum_{i \in [N]} v_i \mathbb{E}_{\mu_0^i}[\mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) + \mathbb{D}_{\mathcal{Y}}(Y, \widehat{y}_i)] \leq N\rho, \end{cases} \tag{27c}$$

where equality (27a) and (27c) are by interchanging the order of the two supremum operators, equality (27b) is from Lemma A.3. This completes the proof. $\qquad\square$

A.3. **Portfolio Allocation with type-∞ Optimal Transport Cost Ambiguity Set**

We first revisit the definition of the type-∞ optimal transport distance.

Definition A.5 (Type-∞ optimal transport cost)**.** *Let $\mathbb{D}$ be a nonnegative and continuous function on $\Xi \times \Xi$. The type-∞ optimal transport cost between two distributions $\mathbb{Q}_1$ and $\mathbb{Q}_2$ supported on Ξ is defined as*

$$\mathrm{W}_\infty(\mathbb{Q}_1, \mathbb{Q}_2) \triangleq \inf \left\{ \operatorname*{ess\,sup}_\pi \left\{ \mathbb{D}(\xi_1, \xi_2) : (\xi_1, \xi_2) \in \Xi \times \Xi \right\} : \pi \in \Pi(\mathbb{Q}_1, \mathbb{Q}_2) \right\}.$$

The type-∞ ambiguity set can be formally defined as

$$\mathbb{B}_\rho^\infty = \left\{ \mathbb{Q} \in \mathcal{M}(\mathcal{X} \times \mathcal{Y}) : \mathrm{W}_\infty(\mathbb{Q}, \widehat{\mathbb{P}}) \leq \rho \right\}.$$

We now provide the reformulation for the mean-variance portfolio allocation problem. To this end, recall that the parameters κ are defined as in (7). In addition, define the following set

$$\mathcal{J} \triangleq \left\{ i \in [N] : \mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) \leq \rho + \gamma \right\},$$

and $\mathcal{J}$ is decomposed further into two disjoint subsets

$$\mathcal{J}_1 = \{ i \in \mathcal{J} : \mathbb{D}_{\mathcal{X}}(x_0, \widehat{x}_i) + \rho \leq \gamma \} \text{ and } \mathcal{J}_2 = \mathcal{J} \backslash \mathcal{J}_1.$$

Proposition A.6 (Mean-variance loss function)**.** *Suppose that ℓ is the mean-variance loss function of the form (5), $\gamma \in \mathbb{R}_+$. Suppose in addition that $\mathcal{X} = \mathbb{R}^n, \mathcal{Y}_0 = \mathbb{R}^m$, $\mathbb{D}_{\mathcal{X}}(x, \widehat{x}) = \|x - \widehat{x}\|^2$, $\mathbb{D}_{\mathcal{Y}}(y, \widehat{y}) = \|y - \widehat{y}\|_2^2$ and $\rho > \min_{i \in [N]} \kappa_i$. The distributionally robust portfolio allocation model with side information*

$$\min_{\alpha \in \mathcal{A}, \ \beta \in \mathbb{R}} \sup_{\mathbb{Q} \in \mathbb{B}_\rho^\infty, \mathbb{Q}(X \in \mathcal{N}_\gamma(x_0)) > 0} \mathbb{E}_{\mathbb{Q}}[\ell(Y, \alpha, \beta) | X \in \mathcal{N}_\gamma(x_0)]$$

is equivalent to the second-order cone program

$$
\begin{aligned}
\min \quad & \lambda \\
\mathrm{s.\,t.} \quad & \alpha \in \mathcal{A}, \ \beta \in \mathbb{R}, \ \lambda \in \mathbb{R}, \ u_i \in \mathbb{R} \ \forall i \in \mathcal{J}_1, \ u_i \in \mathbb{R}_+ \ \forall i \in \mathcal{J}_2, \ t \in \mathbb{R}^N, \ z \in \mathbb{R}_+^N \\
& \sum_{i \in \mathcal{J}} u_i \leq 0 \\
& \left.
\begin{array}{l}
\left\| \begin{bmatrix} 2 z_i \\ 1 - \lambda - u_i - \eta\beta - \frac{1}{4}\eta^2 \end{bmatrix} \right\|_2 \leq 1 + \lambda + u_i + \eta\beta + \frac{1}{4}\eta^2 \\
\widehat{y}_i^\top \alpha - \beta - 0.5\eta \leq t_i, \ -\widehat{y}_i^\top \alpha + \beta + \frac{1}{2}\eta \leq t_i \\
t_i + (\rho - \mathbb{D}_{\mathcal{X}}(\widehat{x}_i^p, \widehat{x}_i))^{1/2} \|\alpha\|_2 \leq z_i
\end{array}
\right\} \quad \forall i \in \mathcal{J}.
\end{aligned}
$$

Proof of Proposition A.6. Using [48, Theorem 2.3], we have

$$
\begin{aligned}
\min \quad & \lambda \\
\mathrm{s.\,t.} \quad & \alpha \in \mathcal{A}, \ \beta \in \mathbb{R}, \ \lambda \in \mathbb{R}, \ u_i \in \mathbb{R} \ \forall i \in \mathcal{J}_1, \ u_i \in \mathbb{R}_+ \ \forall i \in \mathcal{J}_2 \\
& \lambda + u_i \geq v_i^\star(\alpha, \beta) \quad \forall i \in \mathcal{J} \\
& \sum_{i \in \mathcal{J}} u_i \leq 0,
\end{aligned}
$$

where for each $i \in \mathcal{J}$, the value $v_i^\star(\alpha, \beta)$ is

$$
\begin{aligned}
v_i^\star(\alpha, \beta) &= \sup \left\{ \ell(y_i, \alpha, \beta) : y_i \in \mathcal{Y}, \ \|y_i - \widehat{y}_i\|_2^2 \leq \rho - \mathbb{D}_{\mathcal{X}}(\widehat{x}_i^p, \widehat{x}_i) \right\} \\
&= \left(|\widehat{y}_i^\top \alpha - \beta - \frac{1}{2}\eta| + \|\alpha\|_2 (\rho - \mathbb{D}_{\mathcal{X}}(\widehat{x}_i^p, \widehat{x}_i))^{1/2} \right)^2 - \eta\beta - \frac{1}{4}\eta^2.
\end{aligned}
$$

Formulating each constraint $\lambda + u_i \geq v_i^\star(\alpha, \beta)$ using an epigraphical formulation of the form

$$|\widehat{y}_i^\top \alpha - \beta - \frac{1}{2}\eta| \leq t_i, \ t_i + \|\alpha\|_2 (\rho - \mathbb{D}_{\mathcal{X}}(\widehat{x}_i^p, \widehat{x}_i))^{1/2} \leq z_i, \ v_i^\star(\alpha, \beta) \geq z_i^2 - \eta\beta - \frac{1}{4}\eta^2$$

and formulating them as linear or second-order cone constraints completes the proof. $\qquad \square$

REFERENCES

[1] C. D. ALIPRANTIS AND K. C. BORDER, *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer, 2006.

[2] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, *Coherent measures of risk*, Mathematical Finance, 9 (1999), pp. 203–228.

[3] S. Athey, J. Tibshirani, and S. Wager, *Generalized random forests*, Annals of Statistics, 47 (2019), pp. 1148–1178.

[4] M. Baker and J. Wurgler, *Investor sentiment and the cross-section of stock returns*, The Journal of Finance, 61 (2006), pp. 1645–1680.

[5] G.-Y. Ban and C. Rudin, *The big data newsvendor: Practical insights from machine learning*, Operations Research, 67 (2019), pp. 90–108.

[6] S. Benati and R. Rizzi, *A mixed integer linear programming formulation of the optimal mean/value-at-risk portfolio problem*, European Journal of Operational Research, 176 (2007), pp. 423–434.

[7] D. Bertsimas and N. Kallus, *From predictive to prescriptive analytics*, Management Science, 66 (2020), pp. 1025–1044.

[8] D. Bertsimas and C. McCord, *From predictions to prescriptions in multistage optimization problems*, arXiv preprint arXiv:1904.11637, (2019).

[9] F. Black, *Capital market equilibrium with restricted borrowing*, The Journal of Business, 45 (1972), pp. 444–455.

[10] J. Blanchet, L. Chen, and X. Y. Zhou, *Distributionally robust mean-variance portfolio selection with Wasserstein distances*, arXiv preprint arXiv:1802.04885, (2018).

[11] J. Blanchet, Y. Kang, and K. Murthy, *Robust Wasserstein profile inference and applications to machine learning*, Journal of Applied Probability, 56 (2019), pp. 830–857.

[12] J. Blanchet and K. Murthy, *Quantifying distributional model risk via optimal transport*, Mathematics of Operations Research, 44 (2019), pp. 565–600.

[13] J. Blanchet, K. Murthy, and F. Zhang, *Optimal transport based distributionally robust optimization: Structural properties and iterative schemes*, arXiv preprint arXiv:1810.02403, (2018).

[14] M. W. Brandt, P. Santa-Clara, and R. Valkanov, *Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns*, The Review of Financial Studies, 22 (2009), pp. 3411–3447.

[15] M. Britten-Jones, *The sampling error in estimates of mean-variance efficient portfolio weights*, The Journal of Finance, 54 (1999), pp. 655–671.

[16] Z. Chen, M. Sim, and P. Xiong, *Robust stochastic optimization made easy with RSOME*, Management Science, 66 (2020), pp. 3329–3339.

[17] V. K. Chopra and W. T. Ziemba, *The effect of errors in means, variances, and covariances on optimal portfolio choice*, Journal of Portfolio Management, 19 (1993), p. 6.

[18] G. B. Dantzig, *Discrete-variable extremum problems*, Operations Research, 5 (1957), pp. 266–277.

[19] E. Delage and Y. Ye, *Distributionally robust optimization under moment uncertainty with application to data-driven problems*, Operations Research, 58 (2010), pp. 595–612.

[20] V. DeMiguel, L. Garlappi, F. J. Nogales, and R. Uppal, *A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms*, Management Science, 55 (2009), pp. 798–812.

[21] V. DeMiguel, L. Garlappi, and R. Uppal, *Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?*, The Review of Financial Studies, 22 (2009), pp. 1915–1953.

[22] V. DeMiguel, Y. Plyakha, R. Uppal, and G. Vilkov, *Improving portfolio selection using option-implied volatility and skewness*, Journal of Financial and Quantitative Analysis, (2013), pp. 1813–1845.

[23] J. Duchi and H. Namkoong, *Learning models with uniform performance via distributionally robust optimization*, arXiv preprint arXiv:1810.08750, (2018).

[24] A. Esteban-Pérez and J. M. Morales, *Distributionally robust stochastic programs with side information based on trimmings*, arXiv preprint arXiv:2009.10592, (2020).

[25] E. F. Fama and K. R. French, *The cross-section of expected stock returns*, The Journal of Finance, 47 (1992), pp. 427–465.

[26] ———, *Multifactor explanations of asset pricing anomalies*, The Journal of Finance, 51 (1996), pp. 55–84.

[27] M. J. Flannery and A. A. Protopapadakis, *Macroeconomic factors do influence aggregate stock returns*, The Review of Financial Studies, 15 (2002), pp. 751–782.

[28] H. Föllmer and A. Schied, *Convex measures of risk and trading constraints*, Finance and Stochastics, 6 (2002), pp. 429–447.

[29] R. Gao and A. J. Kleywegt, *Distributionally robust stochastic optimization with Wasserstein distance*, arXiv preprint arXiv:1604.02199, (2016).

[30] L. E. Ghaoui, M. Oks, and F. Oustry, *Worst-case value-at-risk and robust portfolio optimization: A conic programming approach*, Operations Research, 51 (2003), pp. 543–556.

[31] J. Goh and M. Sim, *Robust optimization made easy with ROME*, Operations Research, 59 (2011), pp. 973–985.

[32] M. Grinblatt, S. Titman, and R. Wermers, *Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior*, The American Economic Review, (1995), pp. 1088–1105.

[33] R. Ji and M. A. Lejeune, *Data-driven optimization of reward-risk ratio measures*, INFORMS Journal on Computing, (2021).

[34] N. Kallus and X. Mao, *Stochastic optimization forests*, arXiv preprint arXiv:2008.07473, (2020).

[35] R. Kannan, G. Bayraksan, and J. R. Luedtke, *Data-driven sample average approximation with covariate information*, arXiv preprint arXiv:1904.11637, (2020).

[36] ———, *Residuals-based distributionally robust optimization with covariate information*, Available from Optimization Online, (2020).

[37] H. Konno and H. Yamazaki, *Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market*, Management Science, 37 (1991), pp. 519–531.

[38] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, Springer, 2007.

[39] D. Kuhn, P. Mohajerin Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, *Wasserstein distributionally robust optimization: Theory and applications in machine learning*, in Operations Research & Management Science in the Age of Analytics, 2019, pp. 130–166.

[40] M. S. Lobo and S. Boyd, *The worst-case risk of a portfolio*, Unpublished Manuscript, (2000).

[41] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, *Applications of second-order cone programming*, Linear Algebra and Its Applications, 284 (1998), pp. 193–228.

[42] H.-J. Lüthi and J. Doege, *Convex risk measures for portfolio optimization and concepts of flexibility*, Mathematical Programming, 104 (2005), pp. 541–559.

[43] H. Markowitz, *Portfolio selection*, The Journal of Finance, 7 (1952), pp. 77–91.

[44] 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.

[45] MOSEK ApS, *MOSEK Optimizer API for Python 9.2.10*, 2019.

[46] K. Natarajan, M. Sim, and J. Uichanco, *Tractable robust expected utility and risk models for portfolio optimization*, Mathematical Finance, 20 (2010), pp. 695–731.

[47] C. J. Neely, D. E. Rapach, J. Tu, and G. Zhou, *Forecasting the equity risk premium: The role of technical indicators*, Management Science, 60 (2014), pp. 1772–1791.

[48] V. A. Nguyen, F. Zhang, J. Blanchet, E. Delage, and Y. Ye, *Distributionally robust local non-parametric conditional estimation*, in Advances in Neural Information Processing Systems 33, 2020.

[49] L. Pástor and R. F. Stambaugh, *Liquidity risk and expected stock returns*, Journal of Political Economy, 111 (2003), pp. 642–685.

[50] G. Pflug, A. Pichler, and D. Wozabal, *The 1/N investment strategy is optimal under high model ambiguity*, Journal of Banking & Finance, 36 (2012), pp. 410–417.

[51] G. Pflug and D. Wozabal, *Ambiguity in portfolio selection*, Quantitative Finance, 7 (2007), pp. 435–442.

[52] I. Pólik and T. Terlaky, *A survey of the S-lemma*, SIAM Review, 49 (2007), pp. 371–418.

[53] H. Rahimian and S. Mehrotra, *Distributionally robust optimization: A review*, arXiv preprint arXiv:1908.05659, (2020).

[54] H. Rahimian and B. Pagnoncelli, *Contextual chance-constrained programming*, Available on Optimization Online, (2020).

[55] D. E. Rapach, M. C. Ringgenberg, and G. Zhou, *Short interest and aggregate stock returns*, Journal of Financial Economics, 121 (2016), pp. 46–65.

[56] R. T. Rockafellar and S. Uryasev, *Optimization of Conditional Value-at-Risk*, Journal of Risk, 2 (2000), pp. 21–41.

[57] S. Schaible, *Parameter-free convex equivalent and dual programs of fractional programming problems*, Zeitschrift für Operations Research, 18 (1974), pp. 187–196.

[58] A. Shapiro, *On duality theory of conic linear problems*, in Semi-Infinite Programming, Kluwer Academic Publishers, 2001, pp. 135–165.

[59] W. F. Sharpe, *Capital asset prices: A theory of market equilibrium under conditions of risk*, The Journal of Finance, 19 (1964), pp. 425–442.

[60] M. Sion, *On general minimax theorems*, Pacific Journal of Mathematics, 8 (1958), pp. 171–176.

[61] J. E. Smith, *Generalized Chebychev inequalities: Theory and applications in decision analysis*, Operations Research, 43 (1995), pp. 807–825.

[62] C. Villani, *Optimal Transport: Old and New*, Springer, 2008.

[63] C. Zhao and Y. Guan, *Data-driven risk-averse stochastic optimization with Wasserstein metric*, Operations Research Letters, 46 (2018), pp. 262 – 267.

[64] Y. Zhao, Y. Liu, J. Zhang, and X. Yang, *Distributionally robust reward-risk ratio programming with Wasserstein metric*, Available on Optimization Online, (2017).

[65] S. Zhu and M. Fukushima, *Worst-case conditional value-at-risk with application to robust portfolio management*, Operations Research, 57 (2009), pp. 1155–1168.

[66] S. Zymler, D. Kuhn, and B. Rustem, *Distributionally robust joint chance constraints with second-order moment information*, Mathematical Programming, 137 (2013), pp. 167–198.