

WORST-CASE SENSITIVITY

JUN-YA GOTOH[†], MICHAEL JONG KIM[‡], AND ANDREW E.B. LIM^{*}

[†]*Department of Industrial and Systems Engineering, Chuo University, Tokyo, Japan. Email: jgoto@indsys.chuo-u.ac.jp*

[‡]*Sauder School of Business, University of British Columbia, Vancouver, Canada. Email: mike.kim@sauder.ubc.ca*

^{*}*Department of Analytics and Operations, Department of Finance, and Institute for Operations Research and Analytics, National University of Singapore, Singapore. Email: andrewlim@nus.edu.sg*

ABSTRACT. We introduce the notion of *Worst-Case Sensitivity*, defined as the worst-case rate of increase in the expected cost of a Distributionally Robust Optimization (DRO) model when the size of the uncertainty set vanishes. We show that worst-case sensitivity is a *generalized measure of deviation* and that a large class of DRO models are essentially mean-(worst-case) sensitivity problems when uncertainty sets are small, unifying recent results on the relationship between DRO and regularized empirical optimization with worst-case sensitivity playing the role of the regularizer. More generally, DRO solutions can be sensitive to the family and size of the uncertainty set, and reflect the properties of its worst-case sensitivity. We derive closed-form expressions of worst-case sensitivity for well known uncertainty sets including smooth ϕ -divergence, total variation, “budgeted” uncertainty sets, uncertainty sets corresponding to a convex combination of expected value and CVaR, and the Wasserstein metric. These can be used to select the uncertainty set and its size for a given application.

Key words: Distributionally robust optimization, worst-case sensitivity, generalized measure of deviation, model uncertainty, uncertainty sets, regularizer.

1. INTRODUCTION

Consider a cost function $f(x, Y)$ where x is the decision and Y is a discrete random variable with (nominal) distribution $\mathbb{P} = [p_1, \dots, p_n]$. For example, \mathbb{P} could be the empirical distribution associated with a historical sample of Y ’s generated *iid* from some unknown distribution. We would like to find a decision that performs well out-of-sample, a candidate

for which is the minimizer of the sample average approximation (SAA)

$$\min_x \mathbb{E}_{\mathbb{P}}[f(x, Y)].$$

In many situations, the solution of this problem does not perform well out-of-sample because the in-sample model is misspecified. If the expected cost of the in-sample optimizer is sensitive to perturbations of this model, the out-of-sample expected cost may increase significantly because of differences between the in- and out-of-sample distributions. One approach to account for model misspecification is Distributionally Robust Optimization (DRO), where decisions are obtained by optimizing the worst-case expected cost over a family of alternative models. More generally, we would like to find a decision that performs well (in-sample) and continues to perform well when the in-sample model is incorrect.

Distributionally Robust Optimization (DRO). Let $\mathcal{Q}(\varepsilon)$ be an *uncertainty set* of size ε , specifically, a set of probability distributions containing the nominal distribution \mathbb{P} that is increasing in ε and degenerates to the nominal distribution $\mathcal{Q}(0) = \{\mathbb{P}\}$ when $\varepsilon = 0$. For example, $\mathcal{Q}(\varepsilon)$ could be a set of the form $\{\mathbb{Q} \mid d(\mathbb{Q} \mid \mathbb{P}) \leq \varepsilon\}$ where $d(\mathbb{Q} \mid \mathbb{P})$ is ϕ -divergence or the Wasserstein metric. The worst-case expected cost with respect to $\mathcal{Q}(\varepsilon)$ is

$$V(\varepsilon; f(x, \cdot)) := \max_{\mathbb{Q} \in \mathcal{Q}(\varepsilon)} \mathbb{E}_{\mathbb{Q}}[f(x, Y)]. \quad (1.1)$$

Distributionally Robust Optimization (DRO) is the worst-case problem

$$\min_x V(\varepsilon; f(x, \cdot)) \equiv \min_x \max_{\mathbb{Q} \in \mathcal{Q}(\varepsilon)} \mathbb{E}_{\mathbb{Q}}[f(x, Y)].$$

We refer to V as the *value function* of the worst-case problem which, under the assumptions in this paper, is increasing, continuous and concave in ε . When it is clear from the context, we write $V(\varepsilon) \equiv V(\varepsilon; f(x, \cdot))$ if f and x are fixed and we are concerned about the dependence of V on ε , or $V(\varepsilon, x) \equiv V(\varepsilon; f(x, \cdot))$ if we are interested in an optimization problem over x . The solution of the DRO problem is denoted by $x(\varepsilon)$.

Worst-case Sensitivity. Suppose x is fixed and $V(\varepsilon, x)$ has a finite right derivative in ε at $\varepsilon = 0$. The *worst-case sensitivity* is the right derivative of the value function at $\varepsilon = 0$:

$$\mathcal{S}_{\mathbb{P}}[f(x, \cdot)] = V'(0^+; f(x, \cdot)) = \lim_{\varepsilon \downarrow 0} \frac{V(\varepsilon; f(x, \cdot)) - \sum_{i=1}^n p_i f(x, Y_i)}{\varepsilon}. \quad (1.2)$$

When the sensitivity is large, small deviations from the nominal distribution can result in a large increase in the expected cost; such a decision is not robust.

TABLE 1.1. Worst case sensitivities

Type of divergence/worst-case objective	Worst-case sensitivity
ϕ -divergence with any smooth ϕ -function	$\sqrt{\frac{2\mathbb{V}_{\mathbf{p}}(\mathbf{f})}{\phi''(1)}}$
Total variation (or ϕ -div. with $\phi(z) = z - 1 $)	$\frac{1}{2}(\max(\mathbf{f}) - \min(\mathbf{f}))$
budgeted uncertainty (or ϕ -div. with $\phi = \delta_{[0,1+\varepsilon]}$)	$\mathbb{E}_{\mathbf{p}}(\mathbf{f}) - \min(\mathbf{f})$
$(1 - \delta)$ “expected cost” + δ “max cost”	$\max(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})$
ϕ -divergence with $\phi = \delta_{[\frac{1}{1+\varepsilon}, 1+\varepsilon]}$	$\text{CVaR}_{\mathbf{p}, \frac{1}{2}}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})$
$(1 - \delta)$ “expected cost” + δ “ α -CVaR”	$\text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})$
Wasserstein distance with $f(z_i) - f(Y_i) \leq L\ z_i - Y_i\ _p$	$\max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\ z_i - Y_i\ _p}$

Summary of contributions. We show under mild conditions that a large class of DRO problems can be interpreted as multi-objective problems that tradeoff between expected cost and some measure of sensitivity. This measure of sensitivity is bounded above by worst-case sensitivity (1.2) where the bound is tight when ε vanishes. This shows that DRO is fundamentally a tradeoff between mean and sensitivity, generalizing an interpretation of the “variance regularizer” from [16] where the penalty form of DRO with smooth ϕ -divergence was shown to be equivalent to a mean-variance problem when uncertainty sets were small, and unifying recent results connecting various DRO models and “regularized SAA”.

We derive explicit expressions for worst-case sensitivity (1.2) for a number of popular uncertainty sets including smooth ϕ -divergence, total variation, “budgeted uncertainty” (i.e. hard constraints on the likelihood ratio), and the Wasserstein distance, which are summarized in Table 1.1. Under standard assumptions, worst-case sensitivity is a *generalized measure of deviation* [28]. Intuitively, sensitivity is large if small errors in the nominal model—particularly, the probability of extreme costs—have a big impact on the mean, which will be the case if the spread of the cost distribution is large. Sensitivity can be reduced by selecting a decision with a smaller spread, and the equivalence between DRO mean-sensitivity (spread) optimization shows that this is precisely what it is doing. Different uncertainty sets correspond to different measures of spread, which determines the nature of the DRO solutions, while the tradeoff between expected cost and sensitivity is determined by the size of the uncertainty set.

The closed-form expressions we derive for worst-case sensitivity can be used to *select the uncertainty set* that is most appropriate for any given application. For example, DRO with “budgeted uncertainty” controls the spread of the “good” side of the cost distribution

(Table 1.1). Doing so, however, may result in solutions that increase the length of the “bad” side of the tail, making it a poor choice in applications where losses can be large. Indeed, it is possible for the DRO solution under one uncertainty set (e.g., “budgeted uncertainty”) to be less robust than SAA from the perspective of another (e.g., smooth ϕ -divergence); an inappropriate choice of uncertainty set can result in the SAA optimizer being replaced by a decision that is even less robust. Finally, for the inventory problem, L_1 Wasserstein sensitivity is independent of the decision, which makes it a poor choice of uncertainty set for this application.

Overview of paper. We review the relevant literature in Section 2. In Section 3, we show how well-known duality results for DRO allow us to view all DRO problems as a tradeoff between mean cost and worst-case sensitivity when uncertainty sets are small. Under standard assumptions, these sensitivity measures are also measures of spread. We derive explicit expressions for worst-case sensitivity for a number of different uncertainty sets in Section 4. Examples are presented in Section 5.

Remark 1.1. There are uncertainty sets where the right derivative $V'(0^+)$ of the worst-case problem is unbounded. For example, when $d(\mathbb{Q} | \mathbb{P})$ is sufficiently smooth ϕ -divergence, $V(\varepsilon) - V(0)$ is $O(\sqrt{\varepsilon})$ and $V'(0^+)$ is unbounded, so the definition (1.2) is not helpful. If $V(\varepsilon) - V(0) \sim O(g(\varepsilon))$, where $g(\varepsilon)$ is strictly concave and increasing in ε with $g(0) = 0$, we define worst-case sensitivity (with growth rate $g(\varepsilon)$) as

$$\mathcal{S}_{\mathbb{P}}[f] = \lim_{\varepsilon \downarrow 0} \frac{V(\varepsilon) - \sum_{i=1}^n p_i f(Y_i)}{g(\varepsilon)}, \quad (1.3)$$

in which case

$$V(\varepsilon, x) = \sum_{i=1}^n p_i f(x, Y_i) + g(\varepsilon) \mathcal{S}_{\mathbb{P}}[f(x, \cdot)] + o(g(\varepsilon)), \quad \varepsilon > 0. \quad (1.4)$$

As in (3.3), optimizing the worst-case cost $V(\varepsilon, x)$ is again a tradeoff between expected cost and sensitivity, the only difference being that that robustness parameter ε no longer appears linearly. Since $V(\varepsilon)$ is concave and increasing, so too is $g(\varepsilon)$, and $g(0) = 0$.

2. LITERATURE REVIEW

While the DRO (and robust optimization) literature is very large, the focus has been on methodology for solving worst-case problems for a diverse range of uncertainty sets and applications [4, 5, 7, 11, 18, 22, 23, 25, 26]. While the motivation for DRO is to find decisions that are insensitive to model uncertainty, worst-case solutions can be sensitive to the choice of the uncertainty set, and “robust” solutions under one uncertainty set can

be less robust than the SAA solution under another. There is little guidance on how the uncertainty set should be chosen for a given application.

When it comes to size, much of the literature suggests that ε be chosen so that the uncertainty set $\mathcal{Q}(\varepsilon)$ contains the true model with high probability (e.g., 95%). This ignores the multi-objective nature of DRO, and there is also little reason to believe that a pre-ordained confidence level selected independently of data and objective always leads to a “good” decision for all applications¹. A second approach treats ε as a free parameter, like the regularization parameter in regression, chosen to optimize an estimate of out-of-sample expected cost obtained from cross-validation or the bootstrap. While this accounts for both data and the objective, it ignores the sensitivity reduction objective intrinsic to DRO. It is easy to construct examples with high levels of model uncertainty where $\varepsilon = 0$ (SAA) optimizes the resampled estimate of the out-of-sample cost [17]. If sensitivity is ignored, this approach recommends that it is optimal not to be robust.

Worst-case sensitivity (1.2) is defined in [21] where it is used to study the accuracy of simulated estimates of the mean of a random variables. However, this paper does not consider implications for robust optimization, and only considers uncertainty sets defined in terms of relative entropy (a special case of smooth ϕ -divergence).

The connection between worst-case sensitivity and “minimally robust” DRO with smooth ϕ -divergence penalty functions is discussed in [16, 17], with [16] showing that worst-case sensitivity corresponds to the variance of the reward, and [17] studying the out-of-sample properties of the associated mean-variance frontier. The paper [12] derives high-probability performance bounds for the out-of-sample expected reward for solutions of variance-regularized loss minimization. The present paper derives worst-case sensitivity *for uncertainty sets beyond smooth ϕ -divergence*, provides elementary arguments linking DRO and mean-sensitivity optimization under very mild assumptions, and shows that worst-case sensitivity is generally a measure of spread. The mean-sensitivity connection shows that DRO is intrinsically a tradeoff between minimizing the expected cost and controlling the spread of the cost distribution under the nominal, with different uncertainty sets giving rise to different measures of spread. Closed form expressions for worst-case sensitivity make this tradeoff explicit, and can also be used to select the family and size of the uncertainty set in any given application.

Several papers discuss the relationship between DRO and “regularized SAA”. For example, [8, 14, 20, 29] show that worst-case regression and classification problems with the

¹Indeed, extensive calculations with a state of the art data science methods and otherworldly computational power [1] suggests that the answer is more likely to be 42%.

Wasserstein metric are equivalent to regularized versions of these problems. More generally, [6, 8, 12, 14, 20, 29, 31] consider worst-case optimization for a particular uncertainty set and derive the corresponding “regularized SAA” problem. We unify these ideas by showing that the “regularizer” is worst-case sensitivity. We also show that it is a “generalized measure of deviation” [28] and derive an explicit expression for several popular uncertainty sets. These expressions show how an uncertainty set affects the solution of a DRO model, and can be used to select the family and size of the uncertainty set in a given application.

Variance (sensitivity) reduction properties of DRO solutions have also been observed empirically in the literature. In a network model of project management [2], simulations show that robustness prioritizes reducing the variance of the cost (activity duration) over its mean, while [19] shows variance reduction of the out-of-sample cost under the solution of a dynamic DRO problem. The solution of a worst-case newsvendor application [10] is shown to be optimal for a risk-neutral model with a heavy-tailed demand distribution for the demand, thus controlling the impact of extreme events on the cost.

As a final note, while DRO controls the spread (sensitivity) of the cost, robustness can reduce but also increase the out-of-sample variance of the solution (e.g. solution variability decreases for robust classification and regression models that are equivalent to L_p solution regularization [6, 8, 14, 29, 31] while [16] gives an example where solution variability increases).

3. DRO AND SENSITIVITY

We reinterpret duality results from the perspective of mean-sensitivity tradeoffs and show, under mild conditions, that DRO problems are mean-sensitivity problems, that worst-case sensitivity is a tight upper bound of “DRO sensitivity” being equal in the limit $\varepsilon \downarrow 0$, and that worst-case sensitivity is a measure of the spread of the cost distribution. The purpose of this section is to motivate our study of worst-case sensitivity by putting it in the context of classical results and highlighting the cost-sensitivity tradeoff is intrinsic to DRO.

Mean-sensitivity problems. Recall the worst-case objective (1.1). Since uncertainty sets $\mathcal{Q}(\varepsilon)$ are increasing in ε and contain the nominal \mathbb{P} , the worst-case expected cost is equal to the nominal expected cost when $\varepsilon = 0$ and monotonically increasing in ε (see Figure 3.1). It follows that there is a function $\mathcal{A}(\varepsilon; f(x, Y))$, which we refer to as the

ambiguity cost, that is non-negative and increasing in ε such that $\mathcal{A}(0; f(x, Y)) = 0$ and

$$V(\varepsilon, x) = \mathbb{E}_{\mathbb{P}}[f(x, Y)] + \mathcal{A}(\varepsilon; f(x, Y)) \quad (3.1)$$




FIGURE 3.1. Worst-case sensitivity is the right derivative of $V(\varepsilon; f)$ at $\varepsilon = 0$.

For every ε , we define *average sensitivity*

$$\mathcal{S}(\varepsilon; f(x, Y)) \equiv \frac{1}{\varepsilon} \mathcal{A}(\varepsilon; f(x, Y)).$$

It follows that the worst-case problem is a mean-sensitivity problem:

$$\min_x V(\varepsilon, x) \equiv \min_x \mathbb{E}_{\mathbb{P}}[f(x, Y)] + \varepsilon \mathcal{S}(\varepsilon; f(x, Y)). \quad (3.2)$$

When $\varepsilon = 0$, the robust decision maker optimizes the SAA. As ε increases, he/she absorbs a larger expected cost in return for a lower sensitivity, as illustrated in Figure 3.2.

Suppose that $V(\varepsilon)$ is right differentiable at $\varepsilon = 0$. By (1.3), average sensitivity equals worst-case sensitivity in the limit as $\varepsilon \downarrow 0$

$$\mathcal{S}_{\mathbb{P}}[f(x, Y)] = \lim_{\varepsilon \downarrow 0} \mathcal{S}(\varepsilon; f(x, Y)) = \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} \mathcal{A}(\varepsilon; f(x, Y)).$$

In particular, DRO is a tradeoff between mean and worst-case sensitivity when ε is small

$$\min_x V(\varepsilon, x) \equiv \min_x \mathbb{E}_{\mathbb{P}}[f(x, Y)] + \varepsilon \mathcal{S}_{\mathbb{P}}[f(x, Y)] + o(\varepsilon). \quad (3.3)$$

In many situations, $\mathcal{A}(\varepsilon; f(x, Y))$ is also concave in ε in addition to being increasing and non-negative. This is true for all models considered in this paper². In this case, average

² $V(\varepsilon; f(x, \cdot))$, and hence $\mathcal{A}(\varepsilon; f(x, Y))$, is concave in ε if the set of alternative probability measures can be written in the form $\mathcal{Q}(\varepsilon) = \{\mathbb{Q} \mid d(\mathbb{Q}) \leq \varepsilon\}$ for some convex function d of \mathbb{Q} .




FIGURE 3.2. DRO is a tradeoff between expected cost $\mathbb{E}_{\mathbb{P}}[f(x, Y)]$ and average sensitivity $\mathcal{S}(\varepsilon; x)$. x_1 is optimal when $\varepsilon = 0$, but x_2 is optimal when ε is sufficiently large. Although the nominal cost under x_2 is higher, average sensitivity is lower, $\mathcal{S}(\varepsilon; x_2) < \mathcal{S}(\varepsilon; x_1)$.

sensitivity $\mathcal{S}(\varepsilon; f(x, Y))$ is decreasing in ε and worst-case sensitivity is a tight upper bound:

$$\mathcal{S}_{\mathbb{P}}[f(x, Y)] \geq \mathcal{S}(\varepsilon; f(x, Y)) \quad \text{for all } \varepsilon \geq 0.$$

Sensitivity is a measure of deviation. A *Generalized Measure of Deviation* [28] measures the spread of a random variable, generalizing the notion of the standard deviation.

Definition 3.1. Let f be a random variable. $\mathcal{H}[f]$ is a *Generalized Measure of Deviation* or a *Generalized Measure of Spread* of f if

- (1) $\mathcal{H}[f] \geq 0$ with equality if and only if f is constant;
- (2) $\mathcal{H}[\beta f] = \beta \mathcal{H}[f]$ for every constant $\beta \geq 0$;
- (3) $\mathcal{H}[\alpha + f] = \mathcal{H}[f]$ for every constant $\alpha \in \mathbb{R}$.

For the rest of the paper, we focus on discrete random variables, which allows us to treat a random variable as an n -vector. Let $f_i := f(Y_i)$, and denote $\mathbf{f} := (f_1, \dots, f_n)^\top$ and $\mathbf{p} := (p_1, \dots, p_n)^\top$, where p_i is the probability mass on f_i . Without loss of generality, we assume that $p_i > 0$ for all $i = 1, \dots, n$. Likewise, we reserve $\mathbf{q} := (q_1, \dots, q_n)^\top$ for an alternative probability distribution \mathbb{Q} . Furthermore, $\mathbb{E}_{\mathbf{p}}(\mathbf{f}) := \mathbf{p}^\top \mathbf{f}$, $\mathbb{V}_{\mathbf{p}}(\mathbf{f})$, and $\mathcal{S}_{\mathbf{p}}(\mathbf{f})$ denote, respectively, expectation $\mathbb{E}_{\mathbb{P}}[f]$, variance $\mathbb{V}_{\mathbb{P}}[f]$, and worst-case sensitivity $\mathcal{S}_{\mathbb{P}}[f]$ of $\mathbf{f} \equiv f$ under $\mathbf{p} \equiv \mathbb{P}$. Many of our results have generalizations to more complex settings, though the intuition from the discrete setting carries across. The restriction to discrete random variables enables us to communicate our message with minimal technical fuss.

The following result shows that worst-case sensitivity is a *generalized measure of deviation*, and hence, measures the spread of the cost distribution under the nominal. It is well known that $\mathcal{A}(\varepsilon; f)$ and $\mathcal{S}(\varepsilon; f)$ are generalized measures of deviation, which we include in the statement for completeness. The proof is in the Appendix.

Proposition 3.2. *Let $\mathbf{f} \in \mathbb{R}^n$ denote the support of a random variable f , and let $\mathbf{p} \in \mathbb{R}^n$ satisfy $\mathbf{1}^\top \mathbf{p} = 1, \mathbf{p} \geq 0$. Suppose that $d(\mathbf{q}|\mathbf{p}) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and continuous in \mathbf{q} and $d(\mathbf{q}|\mathbf{p}) = 0$ if $\mathbf{q} = \mathbf{p}$. Then the ambiguity cost (3.1) satisfies*

$$\mathcal{A}(\varepsilon; f) = \max_{\mathbf{q} \in \mathcal{Q}(\varepsilon)} \sum_{i=1}^n q_i (f_i - \mathbb{E}_{\mathbb{P}}[f]),$$

$$\mathcal{Q}(\varepsilon) = \left\{ \mathbf{q} = (q_1, \dots, q_n)^\top \in \mathbb{R}^n \mid \mathbf{1}^\top \mathbf{q} = 1, \mathbf{q} \geq 0, d(\mathbf{q}|\mathbf{p}) \leq \varepsilon \right\}.$$

For every $\varepsilon > 0$, the ambiguity cost $\mathcal{A}(\varepsilon; f)$ and average sensitivity $\mathcal{S}(\varepsilon; f) = \frac{1}{\varepsilon} \mathcal{A}(\varepsilon; f)$ are generalized measures of deviation of f . If there is a constant $k > 0$ such that

$$d(\mathbf{p} + \delta \Delta | \mathbf{p}) \sim O(\delta^k) \text{ when } \delta \rightarrow 0, \quad (3.4)$$

for every $\Delta \in \mathbb{R}^n$ such that $\mathbf{1}' \Delta = 0$, then $\mathcal{A}(\varepsilon; f) \sim O(\varepsilon^{\frac{1}{k}})$ and worst-case sensitivity $\mathcal{S}_{\mathbb{P}}[f(x, Y)]$ defined by (1.3) is a generalized measure of deviation with $g(\varepsilon) = \varepsilon^{\frac{1}{k}}$.

When the uncertainty set is a constraint on smooth ϕ -divergence, $g(\varepsilon) = \sqrt{\varepsilon}$ in the definition (1.3) of worst-case sensitivity and (3.4) holds with $k = 2$. It is linear in ε in all other cases considered in this paper. Proposition 3.2 shows that $g(\varepsilon)$ is determined by the continuity property (3.4) of the uncertainty set.

4. WORST-CASE SENSITIVITY: EXPLICIT FORMULAS

We derive explicit expressions for worst-case sensitivity for uncertainty sets associated with smooth ϕ -divergence, Total Variation, budgeted uncertainty sets, uncertainty sets corresponding to a convex combination of the nominal distribution and a CVaR-type uncertainty set, and the Wasserstein metric. For the purposes of readability, all proofs can be found in the Appendix.

4.1. Smooth ϕ -divergences. Consider the worst-case objective

$$V_\phi(\varepsilon; \mathbf{f}) := \max_{\mathbf{q} \in \mathcal{Q}_\phi(\varepsilon)} \sum_{i=1}^n q_i f_i, \quad (4.1)$$

where

$$\mathcal{Q}_\phi(\varepsilon) := \left\{ \mathbf{q} := (q_1, \dots, q_n)^\top \in \mathbb{R}^n \mid \sum_{i=1}^n p_i \phi\left(\frac{q_i}{p_i}\right) \leq \varepsilon, \mathbf{1}^\top \mathbf{q} = 1, \mathbf{q} \geq 0 \right\}.$$

We assume the following.

Assumption 4.1. $\phi(z)$ is strictly convex, twice continuously differentiable in z , with $\phi(1) = 0$, $\phi'(1) = 0$ and $\phi''(1) > 0$.

Since $p_i > 0$ for all i , Assumption 4.1 and convex duality imply that for small $\varepsilon > 0$,

$$V_\phi(\varepsilon) = \min_{\delta > 0, c} \max_{\mathbf{q} > 0} \sum_{i=1}^n q_i f + \frac{1}{\delta} \left(\varepsilon - \sum_{i=1}^n p_i \phi\left(\frac{q_i}{p_i}\right) \right) + c \left(\sum_{i=1}^n q_i - 1 \right).$$

The following result characterizes the solution $(c(\varepsilon), \delta(\varepsilon))$ of the dual problem and the associated worst-case distribution $\mathbf{q}(\varepsilon) = (q_1(\varepsilon), \dots, q_n(\varepsilon))^\top$ when ε is small.

Proposition 4.2. *Suppose that ϕ satisfies Assumption 4.1. Then*

$$c(\varepsilon) = -\mathbb{E}_{\mathbf{p}}(\mathbf{f}) + O(\sqrt{\varepsilon}), \quad (4.2)$$

$$\delta(\varepsilon) = \sqrt{\varepsilon} \sqrt{\frac{2\phi''(1)}{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}} + o(\sqrt{\varepsilon}). \quad (4.3)$$

The family of worst-case distributions $\{\mathbf{q}(\varepsilon) \mid \varepsilon \geq 0\}$ satisfies

$$q_i(\varepsilon) = p_i \left\{ 1 + \sqrt{\frac{2\varepsilon}{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}} (f_i - \mathbb{E}_{\mathbf{p}}(\mathbf{f})) \right\} + o(\sqrt{\varepsilon}). \quad (4.4)$$

The worst-case expected cost is the expected cost under the worst-case distribution

$$\begin{aligned} V_\phi(\varepsilon) &= \mathbb{E}_{\mathbf{q}(\varepsilon)}(\mathbf{f}) \\ &= \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \sqrt{\varepsilon} \sqrt{\frac{2\mathbb{V}_{\mathbf{p}}(\mathbf{f})}{\phi''(1)}} + o(\sqrt{\varepsilon}). \end{aligned} \quad (4.5)$$

It follows that $V_\phi(\varepsilon) - V_\phi(0)$ is $O(\sqrt{\varepsilon})$, so worst-case sensitivity (1.3) with $g(\varepsilon) = \sqrt{\varepsilon}$ is as follows.

Proposition 4.3. *Suppose Assumption 4.1 is satisfied. Then*

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \lim_{\varepsilon \downarrow 0} \frac{V_\phi(\varepsilon) - V_\phi(0)}{\sqrt{\varepsilon}} = \sqrt{\frac{2\mathbb{V}_{\mathbf{p}}(\mathbf{f})}{\phi''(1)}}. \quad (4.6)$$

A closely related result for the case $\phi(z)$ is relative entropy was derived in [21], while [16] derives worst-case sensitivity for the “penalty formulation” of the DRO model.

Example 4.4. When ϕ -divergence is modified modified χ^2 , $\phi(z) = \frac{1}{2}(z - 1)^2$

$$c(\varepsilon) = -\mathbb{E}_{\mathbf{p}}(\mathbf{f}), \quad \delta(\varepsilon) = \sqrt{\frac{2\varepsilon}{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}}$$

and the worst-case distribution is

$$q_i(\varepsilon) = p_i \{1 + \delta(f_i + c)\} = p_i \left\{1 + \sqrt{\frac{2\varepsilon}{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}} (f_i - \mathbb{E}_{\mathbf{p}}(\mathbf{f}))\right\}.$$

This holds for all $\varepsilon \geq 0$ as long as $q_i(\varepsilon) \geq 0$, and not just when it is small. It follows that

$$V_{\phi}(\varepsilon) = \sum_{i=1}^n p_i f_i + \sqrt{\varepsilon} \sqrt{2\mathbb{V}_{\mathbf{p}}(\mathbf{f})}.$$

Clearly, $V_{\phi}(\varepsilon) - V_{\phi}(0) \sim O(\sqrt{\varepsilon})$ and worst-case sensitivity is

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \sqrt{2\mathbb{V}_{\mathbf{p}}(\mathbf{f})}. \quad (4.7)$$

Example 4.5. In [16], the penalty version of the worst-case problem is used to define worst-case sensitivity. Specifically, a family of worst-case distributions $\{\tilde{\mathbf{q}}(\delta) \mid \delta \geq 0\}$ is given by the solutions of the worst-case problem

$$\tilde{\mathbf{q}}(\delta) := \begin{cases} \arg \max_{\mathbf{q}} \left\{ \sum_{i=1}^n q_i f_i - \frac{1}{\delta} \sum_{i=1}^n p_i \phi\left(\frac{q_i}{p_i}\right) \right\}, & \delta > 0, \\ \mathbf{p}, & \delta = 0, \end{cases} \quad (4.8)$$

where the parameter δ determines the penalty on deviations from the nominal. In particular, $\delta = 0$ gives the nominal and increasing δ is analogous to increasing the size of the uncertainty set in (4.1). When the penalty version is used to define the set of worst-case measures, the worst-case expected cost under $\tilde{\mathbf{q}}(\delta)$ is linear in the ambiguity parameter

$$V(\delta) \equiv \mathbb{E}_{\tilde{\mathbf{q}}(\delta)}(\mathbf{f}) = \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \frac{\delta}{\phi''(1)} \mathbb{V}_{\mathbf{p}}(\mathbf{f}) + o(\delta),$$

so the standard definition of sensitivity (1.2) can be used:

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \lim_{\delta \downarrow 0} \frac{\mathbb{E}_{\tilde{\mathbf{q}}(\delta)}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})}{\delta} = \frac{1}{\phi''(1)} \mathbb{V}_{\mathbf{p}}(\mathbf{f}). \quad (4.9)$$

While this leads to a different sensitivity measure, the qualitative nature is the same as (4.6).

4.2. Total Variation. Consider the worst-case expected cost

$$V_{\text{TV}}(\varepsilon; \mathbf{f}) := \max_{\mathbf{q} \in \mathcal{Q}_{\text{TV}}(\varepsilon)} \mathbb{E}_{\mathbf{q}}(\mathbf{f})$$

with uncertainty set

$$\mathcal{Q}_{\text{TV}}(\varepsilon) := \left\{ \mathbf{q} \in \mathbb{R}^n \mid \mathbf{1}^\top |\mathbf{q} - \mathbf{p}| \leq \varepsilon, \mathbf{1}^\top \mathbf{q} = 1, \mathbf{q} \geq 0 \right\},$$

for $\varepsilon \geq 0$, where $|\mathbf{z}| := (|z_1|, \dots, |z_n|)^\top$. (We can focus on $\varepsilon \leq 2$ since the set coincides with the unit simplex $\{\mathbf{q} \in \mathbb{R}^n | \mathbf{1}^\top \mathbf{q} = 1, \mathbf{q} \geq 0\}$ otherwise.) This uncertainty set is equivalent to a constraint on ϕ -divergence with $\phi(z) = |z - 1|$. Note however that $\phi(z)$ is not differentiable at $z = 1$ so the results from Section 4.1 do not apply.

Consider an ordering of the components of the cost vector $\mathbf{f} = (f_1, \dots, f_n)^\top \in \mathbb{R}^n$ from largest to smallest and denote the i^{th} largest component by $f_{(i)}$, i.e.,

$$f_{(1)} \geq \dots \geq f_{(n)}, \quad (4.10)$$

and let $p_{(i)}$ denote the probability mass corresponding to $f_{(i)}$. The following lemma characterizes the worst-case objective for sufficiently small ε .

Lemma 4.6. *Suppose that f corresponds to a nonconstant random variable and $\varepsilon \in (0, \min(\mathbf{p}))$. Then a worst-case probability distribution is*

$$(q_{(1)}, q_{(2)}, \dots, q_{(n-1)}, q_{(n)}) = (p_{(1)} + \frac{\varepsilon}{2}, p_{(2)}, \dots, p_{(n-1)}, p_{(n)} - \frac{\varepsilon}{2}) \quad (4.11)$$

and the worst-case objective is

$$V_{\text{TV}}(\varepsilon; \mathbf{f}) = \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \frac{\varepsilon(\max(\mathbf{f}) - \min(\mathbf{f}))}{2},$$

where $q_{(i)}$ denotes the worst-case probability mass corresponding to $f_{(i)}$.

The expression for worst-case sensitivity follows immediately.

Proposition 4.7. *For the Total Variation uncertainty set $\mathcal{Q}_{\text{TV}}(\varepsilon)$, worst-case sensitivity*

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \frac{\max(\mathbf{f}) - \min(\mathbf{f})}{2} \equiv \frac{1}{2} \times \text{“Range of } \mathbf{f} \text{.”} \quad (4.12)$$

It is known (e.g. [30]) that for any \mathbf{f} and $\mathbf{p} = 1/n$,

$$\sqrt{\mathbb{V}_{\mathbf{p}}(\mathbf{f})} \leq \frac{1}{2} \text{Range}(\mathbf{f}),$$

and, accordingly, we have

$$\sqrt{\frac{\phi''(1)}{2}} V'_\phi(0^+) \leq V'_{\text{TV}}(0^+). \quad (4.13)$$

This suggests that when ε is small, the solution of the DRO problem with the Total Variation uncertainty set will be close to optimal for DRO with smooth ϕ -divergence.

4.3. Budgeted uncertainty. Consider the uncertainty set

$$\mathcal{Q}_b(\varepsilon) = \left\{ \mathbf{q} \in \mathbb{R}^n \mid \mathbf{1}^\top \mathbf{q} = 1, 0 \leq \mathbf{q} \leq (1 + \varepsilon)\mathbf{p} \right\} \quad (4.14)$$

and worst-case expected cost

$$V_b(\varepsilon; \mathbf{f}) := \max_{\mathbf{q} \in \mathcal{Q}_b(\varepsilon)} \mathbb{E}_{\mathbf{q}}(\mathbf{f}), \quad (4.15)$$

for $\varepsilon \geq 0$. (We can focus on $\varepsilon \leq \max_i \{\frac{1}{p_i} - 1\}$ since $\mathcal{Q}_b(\varepsilon)$ is a set of probability distributions otherwise.) For $\varepsilon \in (0, \min_i \{\frac{1}{p_i} - 1\})$, the worst-case distribution is given by

$$(q_{(1)}, \dots, q_{(k)}, q_{(k+1)}, q_{(k+2)}, \dots, q_{(n)}) = ((1 + \varepsilon)p_{(1)}, \dots, (1 + \varepsilon)p_{(k)}, 1 - (1 + \varepsilon) \sum_{i=1}^k p_{(i)}, 0, \dots, 0).$$

The set (4.14) can be referred to as the *budgeted uncertainty set* and is related to the *Conditional Value-at-Risk* with parameter $\alpha \in (0, 1)$ (α -CVaR)

$$\text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) = \max_{\mathbf{q}} \left\{ \mathbf{f}^\top \mathbf{q} \mid \mathbf{1}^\top \mathbf{q} = 1, 0 \leq \mathbf{q} \leq \frac{1}{1 - \alpha} \mathbf{p} \right\} \quad (4.16)$$

Obviously, $V_b(\varepsilon) = \text{CVaR}_{\mathbf{p}, \frac{\varepsilon}{1 + \varepsilon}}(\mathbf{f})$.

$V_b(\varepsilon)$ is piecewise linear, concave, and increasing in ε . The following result characterizes the slope of the worst-case expected cost $V_b(\varepsilon)$ for all values of ε .

Proposition 4.8. *Let $\varepsilon > 0$. Suppose $k \in \{1, \dots, n\}$ is an integer such that*

$$\varepsilon \in \left[\frac{\sum_{i=k+1}^n p_{(i)}}{\sum_{i=1}^k p_{(i)}}, \frac{\sum_{i=k}^n p_{(i)}}{\sum_{i=1}^{k-1} p_{(i)}} \right), \quad (4.17)$$

where $p_{(i)}$ is the probability mass of the i -th largest cost, $f_{(i)}$, and $\sum_{i=n+1}^n p_{(i)} = \sum_{i=1}^0 p_{(i)} = 0$ and $1/0 = \infty$. For $\Delta > 0$ satisfying $\varepsilon + \Delta < \frac{\sum_{i=k}^n p_{(i)}}{\sum_{i=1}^{k-1} p_{(i)}}$, we have

$$S_b(\varepsilon) := \frac{V_b(\varepsilon + \Delta) - V_b(\varepsilon)}{\Delta} = \sum_{i=1}^k p_{(i)} (f_{(i)} - f_{(k+1)}) \quad (4.18)$$

$$= \frac{1}{1 + \varepsilon} \left(\text{CVaR}_{\mathbf{p}, \frac{\varepsilon}{1 + \varepsilon}}(\mathbf{f}) - \text{VaR}_{\mathbf{p}, \frac{\varepsilon}{1 + \varepsilon}}(\mathbf{f}) \right). \quad (4.19)$$

(4.18) defines the constant slope of the linear piece over the interval (4.17). Let

$$\varepsilon_{(h)} = \frac{\sum_{i=n-h+1}^n p_{(i)}}{\sum_{i=1}^{n-h} p_{(i)}}, h = 1, 2, \dots, n.$$

For $\varepsilon \in [\varepsilon_{(k)}, \varepsilon_{(k+1)})$,

$$\begin{aligned} V_b(\varepsilon) &= V_b(0) + \sum_{h=0}^{k-1} S_b(\varepsilon_h)(\varepsilon_{(h+1)} - \varepsilon_{(h)}) + S_b(\varepsilon_k)(\varepsilon - \varepsilon_{(k)}) \\ &= V_b(0) + \sum_{h=0}^{k-1} \frac{p(h)}{1 + \varepsilon_{(h)}} (\text{CVaR}_{\mathbf{p}, \frac{\varepsilon_{(h)}}{1 + \varepsilon_{(h)}}} - \text{VaR}_{\mathbf{p}, \frac{\varepsilon_{(h)}}{1 + \varepsilon_{(h)}}})(\varepsilon_{(h+1)} - \varepsilon_{(h)}) \\ &\quad + \frac{p(k)}{1 + \varepsilon_{(k)}} (\text{CVaR}_{\mathbf{p}, \frac{\varepsilon_{(k)}}{1 + \varepsilon_{(k)}}} - \text{VaR}_{\mathbf{p}, \frac{\varepsilon_{(k)}}{1 + \varepsilon_{(k)}}})(\varepsilon - \varepsilon_{(k)}). \end{aligned}$$

Since $V_b(\varepsilon)$ is concave and increasing, its slope is the largest over the left-most piece $(0, \frac{p(n)}{1 + p(n)})$. For $\varepsilon \in (0, \frac{p(n)}{1 + p(n)})$ and Δ sufficiently small, (4.19) becomes

$$\frac{V_b(\varepsilon + \Delta) - V_b(\varepsilon)}{\Delta} = \mathbb{E}_{\mathbf{p}}(\mathbf{f}) - \min(\mathbf{f}).$$

The following expression for worst-case sensitivity follows immediately.

Corollary 4.9. *For (4.15), we have*

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \mathbb{E}_{\mathbf{p}}(\mathbf{f}) - \min(\mathbf{f}). \quad (4.20)$$

While worst-case sensitivity (4.20) is a measure of spread, it only depends on the “good” side of the cost distribution. In contrast, smooth ϕ -divergence (4.6) and Total Variation (4.12) depend on the entire distribution. We now see an uncertainty set where worst-case sensitivity depends on the spread of the “bad” part of the cost-distribution.

4.4. Convex Combination of Expected Loss and CVaR. Let \mathbf{p} be the nominal distribution, $\alpha \in [0, 1)$ be a fixed parameter, and consider the uncertainty set

$$\mathcal{Q}_c(\varepsilon) := (1 - \varepsilon)\{\mathbf{p}\} + \varepsilon\mathcal{Q}_{\text{CVaR}}(\alpha) \quad (4.21)$$

parameterized by $\varepsilon \in [0, 1]$ where

$$\mathcal{Q}_{\text{CVaR}}(\alpha) := \left\{ \mathbf{q} \in \mathbb{R}^n \mid \mathbf{1}^\top \mathbf{q} = 1, \ 0 \leq \mathbf{q} \leq \frac{1}{1 - \alpha} \mathbf{p} \right\}$$

is the feasible set of (4.16). The worst-case expected cost is

$$V_c(\varepsilon; \mathbf{f}) := \max_{\mathbf{q} \in \mathcal{Q}_c(\varepsilon)} \mathbb{E}_{\mathbf{q}}(\mathbf{f}). \quad (4.22)$$

Observe that $\mathcal{Q}_c(0) = \{\mathbf{p}\}$, so there is no robustness if $\varepsilon = 0$ and the worst-case expected cost is SAA. The uncertainty set (4.21) was considered in [3] and is equivalent to

$$\mathcal{Q}_c(\varepsilon) := \left\{ \mathbf{q} \in \mathbb{R}^n \mid \mathbf{1}^\top \mathbf{q} = 1, \ \mathbf{p}(1 - \varepsilon) \leq \mathbf{q} \leq \mathbf{p}(1 - \varepsilon) + \frac{\varepsilon}{1 - \alpha} \mathbf{p} \right\}.$$

Adopting the convention (4.10), the worst-case probability distribution is given by

$$\begin{aligned}
 (q_{(1)}, \dots, q_{(k)}, q_{(k+1)}, q_{(k+2)}, \dots, q_{(n)}) &\equiv \arg \max_{\mathbf{q} \in \mathcal{Q}_c(\varepsilon)} \mathbb{E}_{\mathbf{q}}(\mathbf{f}) \\
 &= (1 - \varepsilon)(p_{(1)}, \dots, p_{(k)}, p_{(k+1)}, p_{(k+2)}, \dots, p_{(n)}) \\
 &\quad + \varepsilon \left(\frac{1}{1 - \alpha} p_{(1)}, \dots, \frac{1}{1 - \alpha} p_{(k)}, 1 - \frac{1}{1 - \alpha} \sum_{i=1}^k p_{(i)}, 0, \dots, 0 \right) \\
 &= \left(\left(1 + \frac{\alpha \varepsilon}{1 - \alpha}\right) p_{(1)}, \dots, \left(1 + \frac{\alpha \varepsilon}{1 - \alpha}\right) p_{(k)}, (1 - \varepsilon) p_{(k+1)} + \varepsilon \left(1 - \frac{1}{1 - \alpha} \sum_{i=1}^k p_{(i)}\right), \right. \\
 &\quad \left. (1 - \varepsilon) p_{(k+2)}, \dots, (1 - \varepsilon) p_{(n)} \right).
 \end{aligned}$$

It can be shown that the worst-case objective satisfies

$$V_c(\varepsilon; \mathbf{f}) := (1 - \varepsilon) \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \varepsilon \text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}).$$

It follows that for any non-uniform vector \mathbf{f} and $\varepsilon \in (0, 1]$, the function $V_c(\varepsilon)$ is linearly increasing at a rate of its *CVaR Deviation* [28]

$$\frac{V_c(\varepsilon) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})}{\varepsilon} = \text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f}). \quad (4.23)$$

The following expression for worst-case sensitivity is obtained by letting $\varepsilon \searrow 0$.

Corollary 4.10. *For (4.22), we have*

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f}) \equiv \text{“CVaR Deviation of } \mathbf{f} \text{.”} \quad (4.24)$$

Since $\text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) = \max(\mathbf{f}) := \max\{f_1, \dots, f_n\}$ for $\alpha \in [1 - p_{(1)}, 1)$, worst-case sensitivity is the spread of the “bad” part of the cost distribution, $\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \max(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})$, which contrasts with (4.20).

Remark 4.11. If $0 \leq L \leq 1 \leq U$, the uncertainty set

$$\mathcal{Q}_w(L, U) := \left\{ \mathbf{q} \in \mathbb{R}^n \mid \mathbf{1}^\top \mathbf{q} = 1, L\mathbf{p} \leq \mathbf{q} \leq U\mathbf{p} \right\}$$

is equivalent to ϕ -divergence with $\phi(z) = \delta_{[L, U]}(z)$. The worst-case expected cost is

$$V_w(L, U; \mathbf{f}) := L \cdot \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + (1 - L) \cdot \text{CVaR}_{\mathbf{p}, \frac{U-1}{U-L}}(\mathbf{f}).$$

If $(L, U) = (0, \frac{1}{1-\alpha})$, $V_w(L, U)$ is α -CVaR. If $(L, U) = (1 - \varepsilon, \frac{1-(1-\varepsilon)\alpha}{1-\alpha})$, $V_w(L, U)$ is $V_c(\varepsilon)$. While the parameter α is usually fixed (e.g., at 0.95 or 0.99), it can be viewed as another hyperparameter, in addition to ε , that defines the uncertainty set. A reasonable option

to merge the two parameters into a single one is to set as $U = \frac{1}{L} = 1 + \nu > 0$, and the uncertainty set becomes

$$\mathcal{Q}_s(\nu) := \left\{ \mathbf{q} \mid \mathbf{1}^\top \mathbf{q} = 1, \frac{1}{1+\nu} \mathbf{p} \leq \mathbf{q} \leq (1+\nu) \mathbf{p} \right\} = \left\{ \mathbf{q} \mid \mathbf{1}^\top \mathbf{q} = 1, \frac{1}{1+\nu} \mathbf{q} \leq \mathbf{p} \leq (1+\nu) \mathbf{q} \right\}.$$

It is easy to see that the worst-case sensitivity is then $\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = \text{CVaR}_{\mathbf{p}, \frac{1}{2}}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})$.

We can associate CVaR deviation with the standard deviation.

Proposition 4.12. *Let $\mathbf{f} \in \mathbb{R}^n$ and $\alpha \in (0, 1)$. For $\mathbf{p} = \mathbf{1}/n$, we have*

$$\text{“CVaR Deviation of } \mathbf{f}\text{”} \equiv \text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f}) \leq C_{\alpha, n} \sqrt{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}, \quad (4.25)$$

where

$$C_{\alpha, n} := \frac{\sqrt{n \left\{ \lfloor \kappa \rfloor + (\kappa - \lfloor \kappa \rfloor)^2 \right\} - \kappa^2}}{\kappa}$$

with $\kappa := n(1 - \alpha)$. The inequality (4.25) is tight, i.e., there is a vector \mathbf{f} which attains the equality.

Note that $C_{\alpha, n} \leq \sqrt{\frac{\alpha}{1-\alpha}}$ for all $\alpha \in [0, 1)$, and especially when $n(1 - \alpha) \in \mathbb{Z}$, the equality holds. Accordingly, (4.25) suggests a relation between α -CVaR and Mean-Standard Deviation:

$$\text{CVaR}_{\mathbf{p}, \alpha}(\mathbf{f}) \leq \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \sqrt{\frac{\alpha}{1-\alpha}} \sqrt{\mathbb{V}_{\mathbf{p}}(\mathbf{f})},$$

or


$$V_b(\varepsilon; \mathbf{f}) \leq \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \sqrt{\varepsilon} \sqrt{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}.$$

Remark 4.13. While Proposition 1 of [27] shows a similar bound for random variables in the L^2 -space, their coefficient is $1/\sqrt{1-\alpha}$, which is larger than $C_{\alpha, n}$.

The inequality (4.25) is applicable to the worst-case sensitivity results. First of all,

$$V'_c(0^+) \leq C_{\alpha, n} \sqrt{\frac{\phi''(1)}{2}} V'_\phi(0^+) \leq \sqrt{\frac{\phi''(1)\alpha}{2(1-\alpha)}} V'_\phi(0^+) \quad (4.26)$$

where $V'_\phi(0^+)$ is the worst sensitivity of the DRO objective with the smooth ϕ divergence. This inequality suggests that the sensitivity of the DRO with the convex combination of


 FIGURE 4.1. $C_{\alpha,n}$, $\sqrt{\frac{\alpha}{1-\alpha}}$, and $\frac{1}{\sqrt{1-\alpha}}$

mean and CVaR is bounded above by that with any smooth ϕ divergence. Second, recalling Corollary 4.9, we have a tight bound:

$$\begin{aligned} V'_b(0^+) &= \mathbb{E}_{\mathbf{p}}(f) - \min(f) = \frac{\text{CVaR}_{\mathbf{p}, \frac{\varepsilon}{1+\varepsilon}}(f) - \mathbb{E}_{\mathbf{p}}(f)}{\varepsilon} \\ &\leq \varepsilon^{-1/2} \sqrt{\mathbb{V}_{\mathbf{p}}(f)} = \sqrt{\frac{\phi''(1)}{2\varepsilon}} V'_\phi(0^+) \end{aligned}$$

for $\mathbf{p} = 1/n$. Since (4.20) holds for any $\varepsilon \in (0, \frac{p(n)}{1-p(n)}) = (0, \frac{1}{n-1})$, taking $\varepsilon = \frac{1}{n-1}$, we have a (loose) bound

$$V'_b(0^+) < \sqrt{\frac{(n-1)\phi''(1)}{2}} V'_\phi(0^+).$$

From this we see that when n is large, the difference between the smooth ϕ and the budgeted uncertainty $\phi = \delta_{[0, 1+\varepsilon]}$ can be large for small uncertainty sets. In contrast, the bounds (4.13) and (4.26) relating the sensitivities $V'_{\text{TV}}(0^+)$, $V'_c(0^+)$ and $V'_\phi(0^+)$ are tight and independent of n . The potentially large difference likely reflects the fact that the sensitivity of “budgeted uncertainty” depends only on the lower part of the cost distribution whereas $V'_\phi(0^+)$ depends on the entire distribution. More generally, this suggests the possibility that solutions of DRO problems with “budgeted uncertainty” may differ quite substantially from those for other uncertainty sets.

4.5. Wasserstein metric. Consider the worst-case expected cost with a constraint on the Wasserstein metric [8, 13, 14]:

$$V_w(\varepsilon) := \max_{\gamma \in \mathcal{X}} \left\{ \int_z f(x, z) \left(\sum_{i=1}^n \gamma_i(dz) \right) \mid \sum_{i=1}^n \int_z \|z - Y_i\|_p \gamma_i(dz) \leq \varepsilon \right\}, \quad (4.27)$$

where

$$\mathcal{X} = \left\{ \gamma \mid \int_z \gamma_i(dz) = p_i, \ i = 1, \dots, n, \ \gamma_i(dz) \geq 0 \right\}.$$

We are thinking of $1 \leq p \leq \infty$ for the norm in the Wasserstein metric. Note that the cost and constraint functionals

$$\begin{aligned} F(\gamma) &= \max_{\gamma} \int_z f(x, z) \left(\sum_{i=1}^n \gamma_i(dz) \right) \\ G(\gamma) &:= \sum_{i=1}^n \int_z \|z - Y_i\|_p \gamma_i(dz) \end{aligned}$$

are linear in γ , so (4.27) is a convex optimization problem, and $V_w(\varepsilon)$ is concave, increasing and differentiable in ε almost everywhere [24].

Since solutions of the dual problem of (4.27) are supergradients of the value function $V_w(\varepsilon)$ [24], we study worst-case sensitivity $V'_w(0^+)$ by studying dual solutions when $\varepsilon \downarrow 0$.

Let $\lambda \geq 0$ be the Lagrange multiplier for the Wasserstein constraint. The dual problem is

$$\begin{aligned} \min_{\lambda \geq 0} \max_{\gamma \in \mathcal{X}} & \left\{ \sum_{i=1}^n \int_z f(z) \gamma_i(dz) + \lambda \left(\varepsilon - \sum_{i=1}^n \int_z \|z_i - Y_i\|_p \gamma_i(dz) \right) \right\} \\ &= \min_{\lambda \geq 0} \max_{\gamma \in \mathcal{X}} \sum_{i=1}^n \int_{z_i} \left[f(z_i) - \lambda \|z_i - Y_i\|_p \right] \gamma_i(dz_i) + \lambda \varepsilon \end{aligned} \quad (4.28)$$

where to ease notation, we drop the decision variable from the notation and write $f(z) \equiv f(x, z)$. This can be written

$$\begin{aligned} \min_{\lambda \geq 0} & \left\{ \sum_{i=1}^n p_i \max_{z_i} \left\{ f(z_i) - \lambda \|z_i - Y_i\|_p \right\} + \lambda \varepsilon \right\} \\ &= \sum_{i=1}^n p_i f(Y_i) + \min_{\lambda \geq 0} \left\{ \sum_{i=1}^n p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\} + \lambda \varepsilon \right\}. \end{aligned}$$

Intuitively, for every given transportation cost λ , the inner maximization in (4.28) defines a worst-case measure that moves probability mass p_i from Y_i to

$$z_i^* = \arg \max_{z_i} \left\{ f(z_i) - \lambda \|z_i - Y_i\|_p \right\}.$$

If $\lambda(\varepsilon)$ is any solution of the dual problem at ε , $\lambda(\varepsilon)$ is a super-gradient of $V_w(\varepsilon)$ at ε [24], so concavity of $V_w(\varepsilon)$ means that worst-case sensitivity $V'_w(0^+) \leq \lambda(0)$. We compute the sensitivity at $\varepsilon = 0$ by characterizing the solutions of the dual problem of (4.27) when $\varepsilon = 0$, and showing that one of these actually equals the right derivative $V'_w(0^+)$.

By Lemma B.1, strong duality holds when $\varepsilon > 0$. The following result shows that strong duality also holds when $\varepsilon = 0$, and characterizes the set of optimal dual variables.

Proposition 4.14. *Assume that there exists constant L such that $|f(z) - f(Y_i)| \leq L\|z - Y_i\|_p$ for every z and $i = 1, \dots, n$. Then*

$$\min_{\lambda \geq 0} \sum_{i=1}^k p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\} = 0$$

and strong duality holds when $\varepsilon = 0$

$$\begin{aligned} V_w(0) &= \sum_{i=1}^n p_i f(Y_i) + \min_{\lambda \geq 0} \sum_{i=1}^k p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\} \\ &= \sum_{i=1}^n p_i f(Y_i), \end{aligned}$$

and hence for all $\varepsilon \geq 0$. The set of optimal solutions of the dual problem when $\varepsilon = 0$ is

$$\begin{aligned} &\left\{ \lambda \mid \lambda \geq \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p} \right\} \\ &= \arg \min_{\lambda \geq 0} \sum_{i=1}^k p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\}. \end{aligned}$$

Proposition 4.14 implies that if

$$\lambda \geq \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p},$$

then λ is a supergradient of V_w at $\varepsilon = 0$, and hence is an upper bound of the right derivative

$$V'_w(0^+) \leq \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}. \quad (4.29)$$

The following result shows that this inequality is actually an equality, so $V'_w(0^+)$ is also a solution of the dual problem at $\varepsilon = 0$. This allows us to identify the identify worst-case sensitivity with the lower bound of the set of dual solutions at $\varepsilon = 0$. The proof can be found in the Appendix.

Proposition 4.15. *For (4.27), we have*

$$\mathcal{S}_{\mathbb{P}}[f] = V'_w(0^+) = \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}. \quad (4.30)$$

It follows that when ε is small

$$V_w(\varepsilon) = \sum_{i=1}^n p_i f(Y_i) + \varepsilon \left(\max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p} \right) + o(\varepsilon).$$

and that the DRO problem is (almost) the same as the following mean-sensitivity problem:

$$\begin{aligned} \min_x \max_{\gamma \in \mathcal{X}} \int_z f(x, z) \left(\sum_{i=1}^n \gamma_i(dz) \right) \\ = \min_x \mathbb{E}_{\mathbb{P}}[f(x, Y)] + \varepsilon \underbrace{\left\{ \max_{i=1, \dots, n} \max_{z_i} \frac{f(x, z_i) - f(x, Y_i)}{\|z_i - Y_i\|_p} \right\}}_{\mathcal{S}_{\mathbb{P}}[f(x, \cdot)]} + o(\varepsilon). \end{aligned}$$

Example 4.16. If $f(z)$ is concave in z , then the set of optimal dual variables of (4.28) is

$$\left\{ \lambda \mid \lambda \geq \max_{i=1, \dots, n} \|\nabla f(Y_i)\|_q \right\} = \arg \min_{\lambda \geq 0} \sum_{i=1}^k p_i \max_{z_i} \left\{ f(x, z_i) - f(x, Y_i) - \lambda \|z_i - Y_i\|_p \right\}.$$

and worst-case sensitivity

$$\mathcal{S}_{\mathbb{P}}[f] = \max_{i=1, \dots, n} \|\nabla f(Y_i)\|_q. \quad (4.31)$$

Example 4.17. Consider the cost function

$$f(x, Y) = -r \min\{x, Y\} - q \max(x - Y, 0) + s \max(Y - x, 0) + cx \quad (4.32)$$

where $0 \leq q < c < r$ and $s \geq 0$. The negative of this cost function is the reward function for an inventory problem, so minimizing $\mathbb{E}_{\mathbb{P}}[f(x, Y)]$ is equivalent to maximizing expected reward. If $x \in (\min_i Y_i, \max_i Y_i)$, it can be shown that for a Wasserstein metric with $p = 1$

$$\mathcal{S}_{\mathbb{P}}[f(x, \cdot)] = \max\{r - q, s\},$$

so the SAA optimizer is also the solution of the DRO problem for a large range of ε , beyond which, the order quantity is either smaller than $\min_i Y_i$ or larger than $\max_i Y_i$, which is not sensible. This suggests that the Wasserstein uncertainty set with $p = 1$ may not be a good choice for the robust inventory problem.

5. EXAMPLES

5.1. Inventory control. Consider once again the inventory cost function (4.32). In this experiment, we generated $n = 100$ demand realizations $\{Y_1, \dots, Y_n\}$ by sampling from a mixture of two exponential distributions with means $\mu_L = 10$ and $\mu_H = 100$, where the probability of a sample from population L is 0.9. We assume $r = 10$, $c = 2$, $q = 0$ and $s = 4$. Note that $\varepsilon = 0$ is equivalent to SAA.

We begin by comparing the solutions of SAA, the robust inventory problem with a “budgeted” uncertainty set [15] ($\varepsilon = 0.45$), and the robust problem with a modified χ^2 uncertainty set ($\varepsilon = 1.7$). Note first that the order quantity under “budgeted” uncertainty ($x(\varepsilon) = 18$) is smaller than SAA ($x(0) = 24$), while that for the modified χ^2 uncertainty set ($x(\varepsilon) = 44$) is larger. Indeed, the worst-case expected cost of the “budgeted” solution ($x(\varepsilon) = 18$) is larger than that of SAA when evaluated under the modified χ^2 uncertainty set (and vice versa); in the eyes of the modified χ^2 DRO model, solutions of the “budgeted” DRO problem are less robust than SAA.

These observations can be explained by considering the worst-case sensitivity associated with each uncertainty set. For “budgeted” uncertainty, worst-case sensitivity is the spread of the “good” part of the reward distribution (4.20), while for modified χ^2 (and any smooth ϕ -divergence) it is the standard deviation (4.6) which depends on the entire cost distribution. For “budgeted” uncertainty, DRO trades off expected cost in return for a reduction in the spread of the good side of the reward distribution, which is achieved by reducing the order quantity. This comes at the cost of a longer right tail, but this does not affect the sensitivity measure; see Figure 5.1. On the other hand, the modified χ^2 robust optimizer controls the standard deviation of the cost distribution; a larger order quantity reduces the standard deviation of the distribution and the length of the right tail, but the width of the body increases.

We now compare solutions generated by all the uncertainty sets we have discussed. Let $\{x^u(\epsilon) \mid \epsilon \geq 0\}$ denote the family of worst-case solutions where the uncertainty set u is either KL-divergence, modified χ^2 -deviation, Total variation, “budgeted” uncertainty, or the convex-combination uncertainty set. In Figure 5.2 (a), we plot mean-sensitivity frontiers for each family of solution where sensitivity is measured by the standard deviation (4.6) (worst-case sensitivity for smooth ϕ -divergence). The other plots show frontiers with sensitivity associated with (b) Total Variation (4.12), (c) “budgeted uncertainty” (4.20), and (d) *CVaR*-deviation of the cost (4.24).

While DRO solutions reduce the worst-case sensitivity corresponding to its uncertainty set, they may increase other measures of worst-case sensitivity. When the shortage cost $s = 4$, “budgeted” uncertainty produces solutions that increase the other measures of worst-case sensitivity (and vice versa). This can be seen in Figure 5.2 where frontiers for “budgeted” uncertainty moves in the opposite direction to the others. Concurrently, robust order quantities are decreasing in ε while those for other uncertainty sets are increasing. When $s = 0$, however, the order quantities from all DRO models are decreasing in ϵ and the resulting frontiers are the same, as seen in Figure 5.3. Finally, as shown in Example 4.17,




FIGURE 5.1. This plot shows cost distributions under the SAA solution and the solutions of DRO problems with “budgeted” and modified χ^2 uncertainty sets. Worst-case sensitivity for a “budgeted” uncertainty set is the spread of the “good” side of the cost distribution. The DRO solution decreases the order quantity (relative to SAA) as this reduces the width of the body of the distribution (and hence the sensitivity), but this comes at the cost of a longer right tail; this is indicated by the black arrows. On the other hand, worst-case sensitivity for the modified χ^2 uncertainty set is the standard deviation. The associated DRO order quantity is larger than SAA, which reduces the length of the right tail at the cost of a fatter body (red arrows). From the perspective of the modified χ^2 DRO model, the “budgeted” solution is less robust than SAA (and vice versa); the worst-case expected cost is larger than the SAA solution and it has a larger worst-case sensitivity.

the Wasserstein sensitivity is independent of the order quantity when $p = 1$, so the DRO solution is the SAA optimizer.

Mean-sensitivity frontiers can also be used to select the size ε of an uncertainty set. For example, we can use (a) from Figure 5.2 and the modified χ^2 frontier to select ε for a DRO model with a modified χ^2 uncertainty set (the other frontiers are not needed). More importantly, the uncertainty set is an important modelling choice in a DRO model as it determines the measure of sensitivity that is being controlled when solving the worst-case problem.

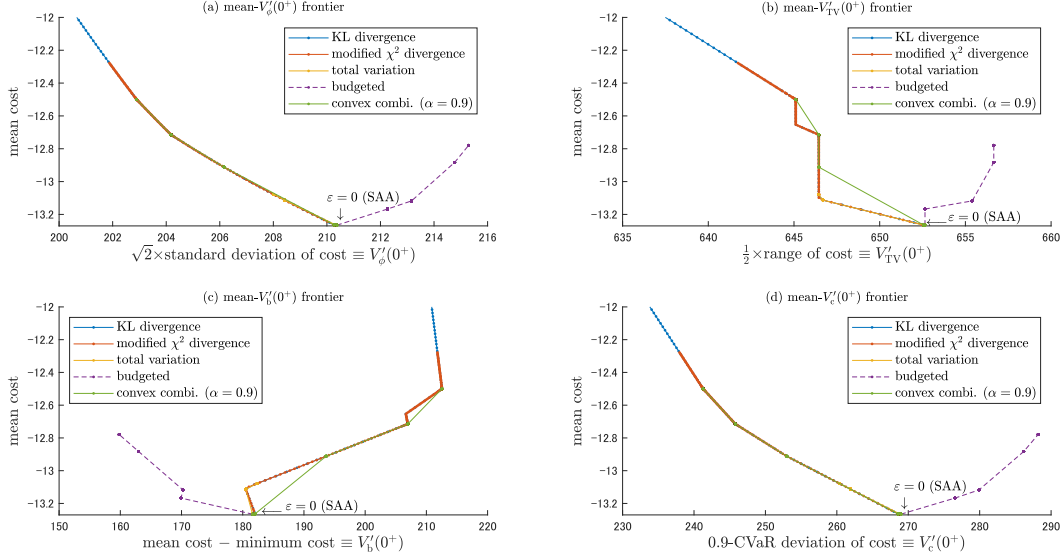


FIGURE 5.2. Mean-sensitivity frontiers for the newsvendor problem when $s = 4$

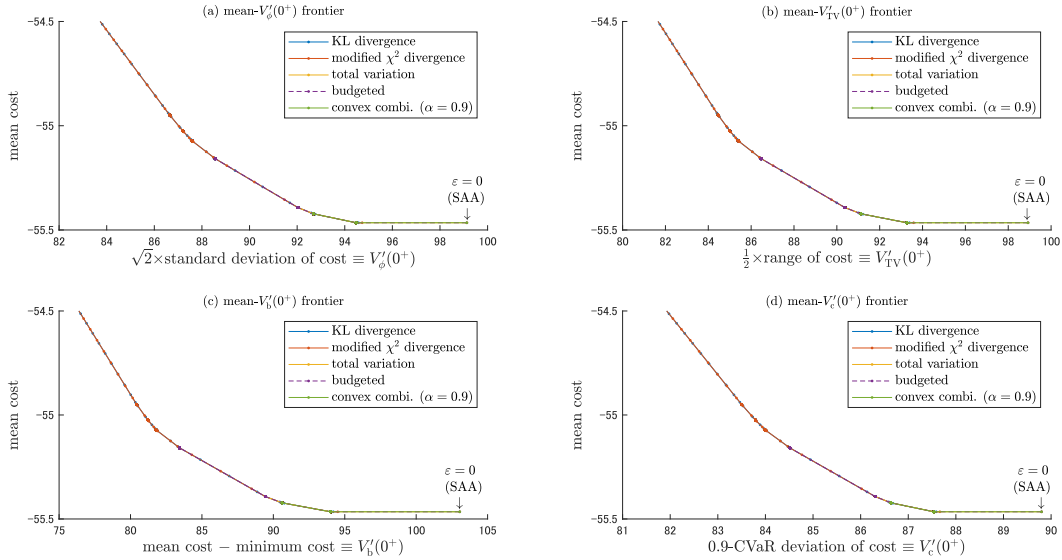


FIGURE 5.3. Mean-sensitivity frontiers for the newsvendor problem when $s = 0$

5.2. Logistic regression. We now consider a higher dimensional example. Fig. 5.4 shows the four mean-sensitivity frontiers of the seven DROs for the logistic regression using the heart failure clinical records dataset [9], which consists of 299 samples having 12 covariates.

The ordinary logistic regression is SAA where the cost of the i -th sample is defined as

$$f(\mathbf{w}, (y_i, \mathbf{x}_i)) := \ln(1 + \exp(-y_i \mathbf{w}^\top \mathbf{x}_i)),$$

where $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \{\pm 1\}$ denote the input vector and binary label of the i -th sample, respectively, and $\mathbf{w} \in \mathbb{R}^d$ is the vector of decision variables. (For the heart failure dataset, the all-one vector is added to the covariates, and thus $d = 13$.) For details on the Wasserstein DRO model for this application, see Appendix C. For this example, all frontiers in Figure 5.4 are similar, which suggests that solutions produced by the seven DRO models we consider are similar.




FIGURE 5.4. Mean-sensitivity frontiers of solutions of the robust logistic regression problem with the Heart Failure data set. Each frontier corresponds to the family of DRO solutions for one of the seven uncertainty sets. Worst-case sensitivity in all the frontier plots is measured by $\sqrt{2} \times$ standard deviation (worst-case sensitivity of the modified χ^2 uncertainty set) of the associated in-sample cost distribution.

6. CONCLUSION

We have derived worst-case sensitivities for the expected reward for several commonly used worst-case models. The results are summarized in Table 1.1.

Worst-case sensitivity is a measure of the spread of the cost distribution $f(x, Y)$. Mathematically, this is a consequence of classical duality results for DRO. Intuitively, the expected cost under the nominal distribution is sensitive to changes in the probability assigned to

extreme values so a cost distribution with a large spread is sensitive to misspecification and not robust. DRO is a tradeoff between expected cost and sensitivity, where the measure of sensitivity, and hence the resulting solution, depends on the uncertainty set.

Practically, our analysis provides a list of sensitivity measures, each corresponding to a different uncertainty set, that can be used to compare the worst-case cost sensitivity for different decisions. That is, worst-case sensitivity serves as a quantitative measure of robustness, which to our knowledge has not been proposed in the literature. The expressions for worst-case sensitivity make explicit the measure of spread that DRO is trying to control for each uncertainty set, and can be used to select the uncertainty set for a given application. Mean-sensitivity frontiers can also be used to select its size.

Acknowledgments. J. Gotoh is supported in part by JSPS KAKENHI Grant 19H02379, 19H00808, and 20H00285. He also thanks Mr. Koichi Fujii for his instruction for the use of RNUOPT, a numerical optimization package of NTT DATA Mathematical Systems Inc., with which the DRO examples in Section 5 are solved.

REFERENCES

- [1] Douglas Adams. *The Hitchhiker's Guide to the Galaxy: A Trilogy in Five Parts*. Random House, 1995.
- [2] Selin Damla Ahipasaoglu, Karthik Natarajan, and Dongjian Shi. Distributionally robust project crashing with partial or no correlation information. *Networks*, 74(1): 79–106, 2019.
- [3] Edward J. Anderson and Andy B. Philpott. Improving sample average approximation using distributional robustness. *Optimization Online*, 2019.
- [4] Aharon Ben-Tal and Arkadi Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1989.
- [5] Aharon Ben-Tal, Dick den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.
- [6] Dimitris Bertsimas and Martin S. Copenhaver. Characterization of the equivalence of robustification and regularization in linear and matrix regression. *European Journal of Operational Research*, 270(3):931–942, 2018.
- [7] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.
- [8] Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3):830–857, 2019.
- [9] Davide Chicco and Giuseppe Jurman. Machine learning can predict survival of patients with heart failure from serum creatinine and ejection fraction alone. *BMC Medical Informatics and Decision Making*, 20(1):16, 2020.
- [10] Bikramjit Das, Anulekha Dhara, and Karthik Natarajan. On the heavy-tail behaviour of the distributionally robust newsvendor. *Operations Research*, Forthcoming, 2020.
- [11] Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Management Science*, 58(2):695–612, 2010.
- [12] John C. Duchi and Hongseok Namkoong. Variance-based regularization with convex objectives. *Journal of Machine Learning Research*, 20(68):1–55, 2019.
- [13] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.
- [14] Rui Gao, Xi Chen, and Anton J Kleywegt. Wasserstein distributional robustness and regularization in statistical learning. *arXiv preprint arXiv:1712.06050*, 2017.

- [15] Jun-ya Gotoh and Yuichi Takano. Newsvendor solutions via Conditional Value-at-Risk minimization. *European Journal of Operational Research*, 179(1):80–96, 2007.
- [16] Jun-ya Gotoh, Michael Jong Kim, and Andrew E.B. Lim. Robust empirical optimization is almost the same as mean–variance optimization. *Operations Research Letters*, 46(4):448–452, 2018.
- [17] Jun-ya Gotoh, Michael Jong Kim, and Andrew E.B. Lim. Calibration of distributionally robust empirical optimization models. *Operations Research* (forthcoming), 2020.
- [18] Lars P. Hansen and Thomas J. Sargent. *Robustness*. Princeton University Press, 2008.
- [19] Michael Jong Kim and Andrew E.B. Lim. Robust multiarmed bandit problems. *Management Science*, 62(1):264–285, 2015.
- [20] Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning. *Optimization Online*, 2020.
- [21] Henry Lam. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4):1248–1275, 2016.
- [22] Andrew E.B. Lim and J. George Shanthikumar. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research*, 55(2):198–214, 2007.
- [23] Andrew E.B. Lim, J. George Shanthikumar, and Zuo-Jun Shen. Model uncertainty, robust optimization, and learning. *TutORials in Operations Research*, 3:66–94, 2006.
- [24] David G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, 1997.
- [25] Ian R. Petersen, Matthew R. James, and Paul Dupuis. Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control*, 45:398–412, 2000.
- [26] Hamed Rahimian, Guzin Bayraksan, and Tito Homem-de Mello. Controlling risk and demand ambiguity in newsvendor models. *European Journal of Operations Research*, 279:854–868, 2019.
- [27] R. Terry Rockafellar, Johannes O. Royset, and Sofia I. Miranda. Superquantile regression with applications to buffered reliability, uncertainty quantification, and conditional value-at-risk. *European Journal of Operational Research*, 234(1):140–154, 2014.
- [28] R. Tyrrell Rockafellar, Stan Uryasev, and Michael Zabarankin. Generalized deviations in risk analysis. *Finance and Stochastics*, 10(1):51–74, 2006.
- [29] Soroosh Shafieezadeh-Abadeh, Peyman Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, pages 1576–1584, 2015.
- [30] Ronald E. Shiffler and Phillip D. Harsha. Upper and lower bounds for the sample standard deviation. *Teaching Statistics*, 2(3):84–86, 1980.

- [31] Huan Xu, Constantine Caramanis, and Shie Mannor. Robust regression and Lasso. *IEEE Transactions on Information Theory*, 56(7):3561–3574, 2010.

APPENDIX A. PROOFS FROM SECTION 3

A.1. Proof of Proposition 3.2. Without loss of generality, we assume $\mathbb{E}_{\mathbb{P}}[f(x, Y)] = 0$ and that the f_i 's are ordered such that $f_1 \geq f_2 \geq \dots \geq f_n$.

$\mathcal{A}(\varepsilon; f)$ and $\mathcal{S}(\varepsilon; f)$ are generalized deviation measures. Clearly, $\mathcal{A}(\varepsilon; f)$ satisfies properties (2) and (3) of Definition 3.1. As for property (1), $\mathcal{A}(\varepsilon; f) \geq 0$ follows from the observation that the nominal measure $\mathbf{p} \in \mathcal{Q}(\varepsilon)$ is feasible so

$$\mathcal{A}(\varepsilon; f) \geq \sum_{i=1}^n p_i f_i = 0.$$

It is also clear that $\mathcal{A}(\varepsilon; f) = 0$ if $f \equiv 0$. We now show the converse.

Suppose that $\mathcal{A}(\varepsilon; f) = 0$ but that f is not equal to 0. Since f is non-zero with $\mathbb{E}_{\mathbb{P}}[f] = 0$, $f_1 > 0 > f_n$. Let $\mathbf{e}^{(i)}$ denote an n -dimensional vector with 1 in the i^{th} entry and 0's elsewhere. Clearly, we can select $\delta > 0$ such that

$$\mathbf{q} = \mathbf{p} + \delta[\mathbf{e}^{(1)} - \mathbf{e}^{(n)}] > 0$$

and $d(\mathbf{p} + \delta[\mathbf{e}^{(1)} - \mathbf{e}^{(n)}] | \mathbf{p}) < \varepsilon$, so

$$\mathcal{A}(\varepsilon; f) \geq \sum_{i=1}^n q_i f_i = \delta(f_1 - f_n) > 0$$

which contradicts our initial assumption that $\mathcal{A}(\varepsilon; f) = 0$. It follows that f must equal 0.

Since $\mathcal{A}(\varepsilon; f)$ is a measure of deviation, so too is $\mathcal{S}(\varepsilon; f)$.

$\mathcal{S}_{\mathbf{p}}(\mathbf{f})$ is a generalized measure of deviation. We first establish the growth condition on $\mathcal{A}(\varepsilon; f)$. Suppose there is a constant $k > 0$ such that the continuity condition (3.4) holds for every $\Delta \in \mathbb{R}^n$ satisfying $\mathbf{1}'\Delta = 0$. Given f and $\varepsilon > 0$, let $\mathbf{q}(\varepsilon)$ denote the solution of the worst-case problem (1.1). Since $d(\mathbf{q}(\varepsilon) | \mathbf{p}) = \varepsilon$, it follows that $\mathbf{q}(\varepsilon) - \mathbf{p} = O(\varepsilon^{\frac{1}{k}})$ so

$$\mathcal{A}(\varepsilon; f) = \sum_{i=1}^n q_i(\varepsilon) f_i \sim O(\varepsilon^{\frac{1}{k}}).$$

We now show that $\mathcal{S}_{\mathbf{p}}(\mathbf{f})$ is a generalized measure of risk. Clearly, $\mathcal{S}_{\mathbf{p}}(\mathbf{f})$ satisfies conditions (2) and (3) of Definition 3.1. As for Condition (1), it is easy to show that $\mathcal{S}_{\mathbf{p}}(\mathbf{f}) \geq 0$ and that $\mathcal{S}_{\mathbf{p}}(\mathbf{f}) = 0$ if f is equal to 0; these properties follow from those of $\mathcal{A}(\varepsilon; f)$. To prove the converse, suppose that

$$\mathcal{S}_{\mathbf{p}}(\mathbf{f}) := \lim_{\varepsilon \downarrow 0} \frac{\mathcal{A}(\varepsilon; f)}{\varepsilon^{\frac{1}{k}}} = 0 \tag{A.1}$$

but that f is not equal to 0. Since $d(\mathbf{q}|\mathbf{p})$ is continuous in \mathbf{q} , there is a constant $\delta > 0$ such that

$$\mathbf{q}(\varepsilon) \equiv (q_1(\varepsilon), \dots, q_n(\varepsilon))^\top := \mathbf{p} + \varepsilon^{\frac{1}{k}} \delta [\mathbf{e}^{(1)} - \mathbf{e}^{(n)}] > 0$$

and

$$d(\mathbf{p} + \varepsilon^{\frac{1}{k}} \delta [\mathbf{e}^{(1)} - \mathbf{e}^{(n)}] | \mathbf{p}) \leq \varepsilon$$

for all $\varepsilon > 0$ sufficiently small. In particular, $q_1(\varepsilon) = p_1 + \varepsilon^{\frac{1}{k}} \delta > 0$, $q_n(\varepsilon) = p_n - \varepsilon^{\frac{1}{k}} \delta > 0$, and $q_i(\varepsilon) = p_i(\varepsilon)$ for $i = 2, \dots, n-1$. Since $\mathbf{q}(\varepsilon) \in \mathcal{Q}(\varepsilon)$, it follows that

$$0 < \sum_{i=1}^n q_i(\varepsilon) f_i = \varepsilon^{\frac{1}{k}} \delta (f_1 - f_n) \leq \mathcal{A}(\varepsilon)$$

so

$$\frac{\mathcal{A}(\varepsilon)}{\varepsilon^{\frac{1}{k}}} \geq \delta (f_1 - f_n) > 0.$$

This contradicts (A.1).

APPENDIX B. PROOFS FROM SECTION 4

B.1. ϕ -divergence uncertainty sets: Proof of Proposition 4.2.

$$\begin{aligned} V(\varepsilon) &= \min_{\delta \geq 0, c} \frac{1}{\delta} \sum_{i=1}^n p_i \max_{q_i} \left\{ \frac{q_i}{p_i} \delta (f_i + c) - \phi\left(\frac{q_i}{p_i}\right) \right\} + \frac{\varepsilon}{\delta} - c \\ &= \sum_{i=1}^n p_i f_i + \min_{\delta \geq 0, c} \frac{1}{\delta} \left\{ \sum_{i=1}^n p_i \left[\phi^*\left(\delta (f_i + c)\right) - \delta (f_i + c) \right] + \varepsilon \right\} \end{aligned}$$

where

$$\phi^*(\zeta) = \max_z \{ \zeta z - \phi(z) \}$$

is the convex conjugate of $\phi(z)$ and

$$q_i = p_i [\phi']^{-1}(\delta (f_i + c)) = \arg \max_{q_i} \left\{ \frac{q_i}{p_i} \delta (f_i + c) - \phi\left(\frac{q_i}{p_i}\right) \right\} \quad (\text{B.1})$$

is the optimizer in the first equality. The worst-case measure $\mathbf{q}(\varepsilon) = (q_1(\varepsilon), \dots, q_n(\varepsilon))^\top$ is obtained by substituting the optimizers over δ and c .

Under the assumptions about $\phi(z)$ from [16], $\phi^*(\zeta)$ is convex and twice continuously differentiable with

$$\phi^*(\zeta) = \zeta + \frac{\zeta^2}{2\phi''(1)} + o(\zeta^2) \quad (\text{B.2})$$

Differentiating with respect to c and δ , the first order conditions are

$$\sum_{i=1}^n p_i [\phi^*]'(\delta[f_i + c]) = 1 \quad (\text{B.3})$$

$$\sum_{i=1}^n p_i \left\{ \phi^*(\delta[f_i + c]) - [\phi^*]'(\delta[f_i + c]) \delta[f_i + c] \right\} + \varepsilon = 0. \quad (\text{B.4})$$

Clearly $c(0) = -\mathbb{E}_{\mathbf{p}}(\mathbf{f})$. We can apply the Implicit Function Theorem to show that $c(\delta)$ is continuously differentiable in δ in the neighborhood of $\delta = 0$ so it follows from (B.2) and (B.3) that

$$c(\delta) = -\mathbb{E}_{\mathbf{p}}(\mathbf{f}) + O(\delta).$$

Together with the expansion of $\phi^*(\zeta)$, (B.4) becomes

$$\frac{\delta^2}{2\phi''(1)} \sum_{i=1}^n p_i (f_i - \mathbb{E}_{\mathbf{p}}(\mathbf{f}))^2 + o(\delta^2) = \varepsilon.$$

The Implicit Function Theorem can again be used to show that $\delta(\varepsilon)$ is continuously differentiable on some open interval³ $(0, b)$ and given by (4.3). It follows that $c(\varepsilon)$ is given by (4.2).

Finally, the worst-case distribution is obtained by substituting (4.2) and (4.3) into (B.1).

B.2. Proof of Lemma 4.6.

Proof. The worst-case objective is written by the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{q}}{\text{maximize}} && \mathbf{f}^\top \mathbf{q} \\ & \text{subject to} && \mathbf{1}^\top \mathbf{q} = 1, \\ & && \mathbf{q} \geq 0, \\ & && \mathbf{1}^\top |\mathbf{p} - \mathbf{q}| \leq \varepsilon, \end{aligned}$$

Note that if $\varepsilon > 0$ is sufficiently small, the nonnegativity condition can be omitted (as long as $p_i > 0$ for all i). In addition, let us introduce nonnegative vectors $\mathbf{u}, \mathbf{v} \geq 0$ such that

³Consider

$$\frac{y}{2\phi''(1)} \sum_{i=1}^n p_i (f_i - \mathbb{E}_{\mathbf{p}}(\mathbf{f}))^2 + o(y) = \varepsilon$$

Observe that $y = 0$ when $\varepsilon = 0$. The Implicit Function Theorem implies that $y(\varepsilon)$ is continuously differentiable on an open interval containing $\varepsilon = 0$. Continuous differentiability of $\delta(\varepsilon)$ on some open interval $(0, b)$ follows from the observation that $\delta(\varepsilon) = \sqrt{y(\varepsilon)}$ for $\varepsilon > 0$.

$\mathbf{u} - \mathbf{v} = \mathbf{q} - \mathbf{p}$. (Note that $\mathbf{1}^\top \mathbf{p} = 1$.) So, under the condition of small ε , the worst-case objective can be rewritten by

$$\begin{aligned} & \underset{\mathbf{u}, \mathbf{v}}{\text{maximize}} && \mathbf{f}^\top \mathbf{u} - \mathbf{f}^\top \mathbf{v} + \mathbf{p}^\top \mathbf{f} \\ & \text{subject to} && \mathbf{1}^\top \mathbf{u} - \mathbf{1}^\top \mathbf{v} = 0, \\ & && \mathbf{1}^\top \mathbf{u} + \mathbf{1}^\top \mathbf{v} \leq \varepsilon, \\ & && \mathbf{u}, \mathbf{v} \geq 0. \end{aligned}$$

The dual LP problem is derived as

$$\begin{aligned} & \underset{\lambda \geq 0, \theta}{\text{minimize}} && \varepsilon \lambda + \mathbf{p}^\top \mathbf{f} \\ & \text{subject to} && \mathbf{1} \theta + \mathbf{1} \lambda \geq \mathbf{f}, \\ & && -\mathbf{1} \theta + \mathbf{1} \lambda \geq -\mathbf{f}, \\ & && \lambda \geq 0, \end{aligned}$$

which can be reduced to

$$\underset{\theta}{\text{minimize}} \quad \varepsilon \max\{|f_1 - \theta|, \dots, |f_n - \theta|\} + \mathbf{p}^\top \mathbf{f}.$$

It is easy to see that its optimality attained at $\theta = \frac{f_{(1)} + f_{(n)}}{2}$ with the optimal value being $\frac{\varepsilon(f_{(1)} - f_{(n)})}{2} + \mathbf{p}^\top \mathbf{f}$.

By the strong duality theorem, the worst-case objective turns out to be

$$\mathbf{p}^\top \mathbf{f} + \frac{\varepsilon(f_{(1)} - f_{(n)})}{2} = \mathbb{E}_{\mathbf{p}}(\mathbf{f}) + \frac{\varepsilon(\max(\mathbf{f}) - \min(\mathbf{f}))}{2},$$

which is attained by the solution (4.11). Consequently, we complete the proof of Lemma 4.6. \square

B.3. Proof of Proposition 4.8.

Proof. Solving LP (4.16) (by a greedy algorithm), a worst-case distribution is given as

$$q_{(i)} = \begin{cases} (1 + \varepsilon)p_{(i)}, & i = 1, \dots, k, \\ 1 - (1 + \varepsilon) \sum_{i=1}^k p_{(i)}, & i = k + 1, \\ 0, & i = k + 2, \dots, n, \end{cases} \quad (\text{B.5})$$

where $q_{(i)}$ denotes the probability mass for the i -th largest cost $f_{(i)}$. Putting (B.5) into the objective of (4.16) ($\equiv \sum_{i=1}^n q_{(i)} f_{(i)}$), we have

$$\begin{aligned} V_{\text{CVaR}}(\varepsilon) &= \text{CVaR}_{\mathbf{p}, \frac{\varepsilon}{1+\varepsilon}}(\mathbf{f}) \\ &= (1+\varepsilon) \sum_{i=1}^k p_{(i)} f_{(i)} + \left(1 - (1+\varepsilon) \sum_{i=1}^k p_{(i)}\right) f_{(k+1)}. \end{aligned}$$

With this formula,

$$\begin{aligned} &V_{\text{CVaR}}(\varepsilon + \Delta) - V_{\text{CVaR}}(\varepsilon) \\ &= \text{CVaR}_{\mathbf{p}, \frac{\varepsilon+\Delta}{1+\varepsilon+\Delta}}(\mathbf{f}) - \text{CVaR}_{\mathbf{p}, \frac{\varepsilon}{1+\varepsilon}}(\mathbf{f}) \\ &= (1+\varepsilon+\Delta) \sum_{i=1}^k p_{(i)} f_{(i)} + \left\{1 - (1+\varepsilon+\Delta) \sum_{i=1}^k p_{(i)}\right\} f_{(k+1)} \\ &\quad - (1+\varepsilon) \sum_{i=1}^k p_{(i)} f_{(i)} - \left\{1 - (1+\varepsilon) \sum_{i=1}^k p_{(i)}\right\} f_{(k+1)} \\ &= \frac{\Delta}{1+\varepsilon} \left[(1+\varepsilon) \sum_{i=1}^k p_{(i)} f_{(i)} + \left\{1 - (1+\varepsilon) \sum_{i=1}^k p_{(i)}\right\} f_{(k+1)} - \underbrace{f_{(k+1)}}_{\text{VaR}_{\mathbf{p}, \frac{\varepsilon}{1+\varepsilon}}(\mathbf{f})} \right] \quad (\text{B.6}) \\ &= \frac{\Delta}{1+\varepsilon} \left(\text{CVaR}_{\mathbf{p}, \frac{\varepsilon}{1+\varepsilon}}(\mathbf{f}) - \text{VaR}_{\mathbf{p}, \frac{\varepsilon}{1+\varepsilon}}(\mathbf{f}) \right), \end{aligned}$$

which proves (4.19). \square

B.4. Proof of Proposition 4.12.

Proof. Note that the derivation of the tight inequality (4.25) is equivalent to showing the reciprocal of the minimum of the following optimization problem is equal to $1/C_{\alpha,n}$

$$\frac{1}{C_{\alpha,n}} = \left| \begin{array}{ll} \underset{\mathbf{f}}{\text{minimize}} & \frac{\sqrt{\mathbb{V}_{\mathbf{p}}(\mathbf{f})}}{\text{CVaR}_{\mathbf{p},\alpha}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f})} \\ \text{subject to} & \mathbf{f} \neq C\mathbf{1} \text{ for any } C, \end{array} \right. \quad (\text{B.7})$$

where $\mathbf{p} = 1/n$. Noting that both the denominator and numerator are positively homogeneous and that CVaR is translation invariant and, thus, $\text{CVaR}_{\mathbf{p},\alpha}(\mathbf{f}) - \mathbb{E}_{\mathbf{p}}(\mathbf{f}) = \text{CVaR}_{\mathbf{p},\alpha}(\mathbf{f} - \mathbb{E}_{\mathbf{p}}(\mathbf{f})\mathbf{1})$, the fractional program (B.7) is equivalently rewritten as

$$\frac{1}{C_{\alpha,n}^2} = \left| \begin{array}{ll} \underset{\mathbf{f}, \mathbf{z}}{\text{minimize}} & \frac{1}{n} \mathbf{z}^\top \mathbf{z} \\ \text{subject to} & \text{CVaR}_{\mathbf{p},\alpha}(\mathbf{z}) = 1, \\ & \mathbf{z} = \mathbf{f} - \frac{1}{n} \mathbf{1}\mathbf{1}^\top \mathbf{f}, \mathbf{f} \neq C\mathbf{1} \text{ for any } C, \end{array} \right.$$

where the objective function is squared for the convenience. Noting that the final constraint implies $1^\top \mathbf{z} = 0$, we consider the following relaxed optimization problem:

$$\left| \begin{array}{ll} \underset{\mathbf{z}}{\text{minimize}} & \frac{1}{n}(z_1^2 + \cdots + z_n^2) \\ \text{subject to} & \frac{1}{\kappa}\{z_1 + \cdots + z_k + (\kappa - k)z_{k+1}\} = 1, \\ & z_1 \geq z_2 \geq \cdots \geq z_n, \\ & z_1 + \cdots + z_n = 0, \end{array} \right. \quad (\text{B.8})$$

where $\kappa := n(1 - \alpha)$ and $k := \lfloor \kappa \rfloor$.

Let us tentatively consider to remove the inequality constraints from (B.8):

$$\left| \begin{array}{ll} \underset{\mathbf{z}}{\text{minimize}} & \frac{1}{n}(z_1^2 + \cdots + z_n^2) \\ \text{subject to} & z_1 + \cdots + z_k + (\kappa - k)z_{k+1} = \kappa, \\ & (1 - \kappa + k)z_{k+1} + z_{k+2} + \cdots + z_n = -\kappa. \end{array} \right.$$

The Lagrangian of this is given by

$$L(\mathbf{z}, \lambda, \mu) = z_1^2 + \cdots + z_n^2 - 2\lambda\{z_1 + \cdots + z_k + (\kappa - k)z_{k+1} - \kappa\} - 2\mu\{(1 - \kappa + k)z_{k+1} + z_{k+2} + \cdots + z_n + \kappa\}$$

and the optimality condition is given by

$$\left\{ \begin{array}{ll} \frac{\partial L}{\partial z_j}(\mathbf{z}, \lambda, \mu) = 0, & j = 1, \dots, n, \\ \frac{\partial L}{\partial \lambda}(\mathbf{z}, \lambda, \mu) = 0 & \leftrightarrow z_1 + \cdots + z_k + (\kappa - k)z_{k+1} = \kappa, \\ \frac{\partial L}{\partial \mu}(\mathbf{z}, \lambda, \mu) = 0 & \leftrightarrow (1 - \kappa + k)z_{k+1} + z_{k+2} + \cdots + z_n = -\kappa. \end{array} \right.$$

Here the first condition is

$$\frac{\partial L}{\partial z_j}(\mathbf{z}, \lambda, \mu) = \left\{ \begin{array}{ll} 2z_j - 2\lambda = 0, & 1 \leq j \leq k, \\ 2z_{k+1} - 2(\kappa - k)\lambda - 2(1 - \kappa + k)\mu = 0, & j = k + 1, \\ 2z_j - 2\mu = 0, & k + 2 \leq j \leq n. \end{array} \right.$$

Substituting z_j 's derived from this to the last two conditions, we have

$$\left\{ \begin{array}{ll} \{k + (\kappa - k)^2\}\lambda + (\kappa - k)(1 - \kappa + k)\mu & = \kappa, \\ (\kappa - k)(1 - \kappa + k)\lambda + \{n - k - 1 + (1 - \kappa + k)^2\}\mu & = -\kappa, \end{array} \right.$$

and

$$\left\{ \begin{array}{l} \lambda = \frac{\kappa(n - \kappa)}{n\{k + (\kappa - k)^2\} - \kappa^2} \\ \mu = \frac{-\kappa^2}{n\{k + (\kappa - k)^2\} - \kappa^2}, \end{array} \right.$$

so

$$\begin{cases} z_j &= \frac{\kappa(n - \kappa)}{n\{k + (\kappa - k)^2\} - \kappa^2}, & j = 1, \dots, k, \\ z_{k+1} &= \frac{-\kappa^2 + n\kappa(\kappa - k)}{n\{k + (\kappa - k)^2\} - \kappa^2}, \\ z_j &= \frac{-\kappa^2}{n\{k + (\kappa - k)^2\} - \kappa^2}, & j = k + 2, \dots, n, \end{cases} \quad (\text{B.9})$$

and the objective value is then

$$\frac{1}{n} \mathbf{z}^\top \mathbf{z} = \frac{\kappa^2}{n\{k + (\kappa - k)^2\} - \kappa^2}. \quad (\text{B.10})$$

The vector $\mathbf{z} = (z_1, \dots, z_n)^\top$ given by (B.9) satisfies $z_1 \geq \dots \geq z_n$, and turns out to be optimal to (B.8) as well. Furthermore, for the vector \mathbf{z} given by (B.9) and any constant C , we can reproduce a nonconstant vector \mathbf{f} as $\mathbf{f} = \mathbf{z} + C\mathbf{1}$, which satisfies $\mathbf{z} = \mathbf{f} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top \mathbf{f}$. Consequently, the square root of (B.10) is the optimal value of (B.7). \square

B.5. Wasserstein uncertainty sets.

Preliminaries: Some useful results from convex duality. We summarize general properties of the dual problem and the relationship between optimal dual variables and super-gradients of the value function for the optimization problem

$$V(\varepsilon) := \max_{x \in \Omega} F(x) \text{ subject to: } G(x) \leq \varepsilon \quad (\text{B.11})$$

which we apply to (4.27). Here, we assume that $F : X \rightarrow \mathbb{R}$ and $G : X \rightarrow \mathbb{R}$, where X is a vector space and Ω is a convex subset of \mathcal{X} . The associated dual problem is

$$D(\varepsilon) := \min_{\lambda \geq 0} \max_{x \in \mathcal{X}} F(x) + \lambda\{\varepsilon - G(x)\} \equiv \max_{\lambda \geq 0} \left\{ H(\lambda) + \lambda\varepsilon \right\} \quad (\text{B.12})$$

where $\lambda \in \mathbb{R}$ is the Lagrange multiplier and

$$H(\lambda) := \max_{x \in \Omega} \left\{ F(x) - \lambda G(x) \right\}.$$

Note that $H(\lambda)$ is convex in λ .

We denote the derivative of H at λ by $H'(\lambda)$, and the directional derivatives

$$\begin{aligned} H'(\lambda^+) &:= \lim_{\delta \downarrow 0} \frac{H(\lambda + \delta) - H(\lambda)}{\delta} \\ H'(\lambda^-) &:= \lim_{\delta \downarrow 0} \frac{H(\lambda - \delta) - H(\lambda)}{\delta} \end{aligned}$$

(We refer to $H'(\lambda^+)$ and $H'(\lambda^-)$ as the right and left derivative of H at λ , respectively). $H(\lambda)$ is convex in λ so $H'(\lambda^+)$ is increasing in λ and $H'(\lambda^-)$ is decreasing in λ .

Lemma B.1. *Consider the optimization problem (B.11). Assume that $F(x)$ is concave, that $G(x)$ is convex and non-negative, that there is an $x \in \Omega$ such that $G(x) = 0$, and that $V(\varepsilon)$ is finite for every ε . Then $V(\varepsilon)$ is concave and increasing in $\varepsilon \geq 0$, and differentiable at almost every $\varepsilon > 0$. When $\varepsilon > 0$, strong duality holds ($V(\varepsilon) = D(\varepsilon)$) and the maximum of the dual problem is achieved. If $\lambda(\varepsilon)$ is a solution of the dual problem corresponding to $\varepsilon > 0$, then $\lambda(\varepsilon)$ is a super-gradient of V at ε . If V is differentiable at $\varepsilon > 0$, then the dual problem has a unique solution and $V'(\varepsilon) = \lambda(\varepsilon)$.*

Proof. When $\varepsilon > 0$ the Lagrange Duality Theorem [24] implies that strong duality holds and that there exists a solution $\lambda(\varepsilon)$ of the dual problem. \square

Proposition B.2. *Assume that $F(x)$ is concave, that $G(x)$ is convex and non-negative, that there is an $x \in \Omega$ such that $G(x) = 0$, and that $V(\varepsilon)$ is finite for every ε . Let $\{\varepsilon_i\}$ be sequence of positive numbers such that $\varepsilon_i \downarrow \varepsilon$. Suppose that V is differentiable at ε_i , strong duality holds at ε_i , and there is a solution $\lambda_i \equiv \lambda(\varepsilon_i)$ of the dual problem at ε_i , for every i . Then $\lambda_i \equiv \lambda(\varepsilon_i)$ is increasing. If $\lambda_i \uparrow \lambda^* < \infty$ when $\varepsilon_i \downarrow \varepsilon$, then λ^* is a solution of the dual problem at ε .*

Proof. Since strong duality holds at ε_i

$$V(\varepsilon_i) = \min_{\lambda \geq 0} \{H(\lambda) + \lambda \varepsilon_i\} = H(\lambda(\varepsilon_i)) + \lambda(\varepsilon_i) \varepsilon_i,$$

where $H(\lambda)$ is a convex function of $\lambda \geq 0$, and hence is differentiable at almost every $\lambda \geq 0$.

Consider first the case that $\varepsilon_i \downarrow \varepsilon$ where $\varepsilon > 0$. Since $\lambda_i \equiv \lambda(\varepsilon_i)$ is a solution of the dual problem

$$\begin{aligned} H'(\lambda_i^+) + \varepsilon_i &\geq 0 \\ H'(\lambda_i^-) - \varepsilon_i &\geq 0. \end{aligned} \tag{B.13}$$

Since $\varepsilon_i \downarrow \varepsilon$ and $\lambda_i \uparrow \lambda^*$ as $i \rightarrow \infty$, and the right derivative is increasing in λ , it follows that

$$H'(\lambda^{*+}) + \varepsilon \geq \lim_{i \rightarrow \infty} \{H'(\lambda_i^+) + \varepsilon_i\} \geq 0.$$

On the other hand, the left derivative $H'(\lambda^-)$ is left-continuous in λ so

$$H'(\lambda^{*-}) - \varepsilon = \lim_{i \rightarrow \infty} \{H'_-(\lambda(\varepsilon_i)) - \varepsilon_i\} \geq 0.$$

It follows that λ^* is optimal for the dual problem at ε .

When $\varepsilon = 0$, we need only consider the derivative from the right at $\lambda(0)$. In particular, it follows from (B.13) that

$$H'(\lambda^{*+}) \geq \lim_{\varepsilon_i \downarrow 0} \left\{ H'_+(\lambda(\varepsilon_i)) - \varepsilon_i \right\} \geq 0.$$

□

B.5.1. Proof of Proposition 4.14.

Proof. For every $\lambda \geq 0$ we have

$$\sum_{i=1}^k p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\} \geq 0.$$

We characterize the optimal dual variables by finding the values of λ such that the lower bound is attained.

If

$$\lambda \geq \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}$$

it follows that

$$f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \geq 0$$

for every i , with equality when $z_i = Y_i$. It follows that

$$\sum_{i=1}^k p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\} = 0$$

so λ is a solution of the dual problem. If

$$\lambda < \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}$$

then for every i such that

$$\lambda < \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}$$

we can find z_i such that $f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p > 0$. It follows that

$$\sum_{i=1}^k p_i \max_{z_i} \left\{ f(z_i) - f(Y_i) - \lambda \|z_i - Y_i\|_p \right\} > 0$$

so λ is not a solution of the dual problem.

□

B.5.2. *Proof of Proposition 4.15.*

Proof. Suppose that the right derivative of V at $\varepsilon = 0$ is finite (i.e., $|V'(0^+)| < \infty$). Let $\{\varepsilon_i\}$ be any sequence such that $\varepsilon_i > 0$, $\varepsilon_i \downarrow 0$ and V is differentiable at ε_i . The Lagrange Duality Theorem [24] implies that strong duality holds and that there exists a unique solution $\lambda_i \equiv \lambda(\varepsilon_i)$ of the dual problem for each ε_i . Since $\{\lambda(\varepsilon_i)\}$ is an increasing sequence that is bounded above by $0 \leq V'(0^+) < \infty$ (we have selected $\{\varepsilon_i\}$ such that V is differentiable at ε_i , so $\lambda(\varepsilon_i) = V'(\varepsilon_i)$; in addition, $0 \leq V'(\varepsilon_i) \leq V'(0^+)$ by the concavity of $V(\varepsilon)$), it converges to a limit that is also bounded by $V'_+(0)$

$$\lambda^* = \lim_{i \rightarrow \infty} \lambda(\varepsilon_i) = \lim_{i \rightarrow \infty} V'(\varepsilon_i) \leq V'(0^+) < \infty.$$

By Proposition B.2, the limit λ^* is a solution of the dual problem at $\varepsilon = 0$. Since

$$\left[\max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}, \infty \right)$$

is the set of Lagrange multipliers that solve the dual problem when $\varepsilon = 0$,

$$V'(0^+) \geq \lambda^* \geq \max_{i=1, \dots, n} \max_{z_i} \frac{f(z_i) - f(Y_i)}{\|z_i - Y_i\|_p}$$

Together with (4.29), it now follows that worst-case sensitivity satisfies (4.30). \square

APPENDIX C. COMMENTS ON WASSERSTEIN DRO IN EXAMPLE 5.2

For the Wasserstein DRO, [29] shows that when the transportation cost is given by $c((\mathbf{x}_i, y_i), (\mathbf{x}_j, y_j)) = \|\mathbf{x}_i - \mathbf{x}_j\|_2 + \kappa|y_i - y_j|$ for a constant $\kappa \geq 0$, the robust counterpart can be reduced to the following convex optimization problem:

$$\begin{aligned} & \underset{\lambda \geq 0, \mathbf{s}, \mathbf{w}}{\text{minimize}} && \varepsilon \lambda + \frac{1}{n} \sum_{i=1}^n s_i \\ & \text{subject to} && s_i \geq \ln \left(1 + \exp \left(-y_i (\mathbf{x}_i^\top \mathbf{w}) \right) \right), \quad i = 1, \dots, n, \\ & && s_i \geq \ln \left(1 + \exp \left(y_i (\mathbf{x}_i^\top \mathbf{w}) \right) \right) - \kappa \lambda, \quad i = 1, \dots, n, \\ & && \lambda \geq \|\mathbf{w}\|_2. \end{aligned}$$

Here κ balance the cost of the covariates and that of the binary label. If the user is not interested in the mislabeling, s/he can set as $\kappa = \infty$, so that the formulation results in the regularized logistic regression:

$$\underset{\lambda \geq 0, \mathbf{s}, \mathbf{w}}{\text{minimize}} \quad \varepsilon \|\mathbf{w}\|_2 + \frac{1}{n} \sum_{i=1}^n \ln \left(1 + \exp \left(-y_i (\mathbf{x}_i^\top \mathbf{w}) \right) \right).$$

Then the worst-case sensitivity is given by $\mathcal{S}_p(f) = \|\mathbf{w}_{SAA}\|_2$ where \mathbf{w}_{SAA} is a solution to the ordinary logistic regression.