

Fast Computation of Superquantile-Constrained Optimization Through Implicit Scenario Reduction

Jake Roth

ROTH0674@UMN.EDU

*Department of Industrial and Systems Engineering
University of Minnesota
Minneapolis, MN 55414*

Ying Cui

YINGCUI@BERKELEY.EDU

*Department of Industrial Engineering and Operations Research
University of California, Berkeley
Berkeley, CA 94720*

Editor: Zhihua Zhang

Abstract

Superquantiles have recently gained significant interest as a risk-aware metric for addressing fairness and distribution shifts in statistical learning and decision making problems. This paper introduces a fast, scalable and robust second-order computational framework to solve large-scale optimization problems with superquantile-based constraints. Unlike empirical risk minimization, superquantile-based optimization requires ranking random functions evaluated across all scenarios to compute the tail conditional expectation. While this tail-based feature might seem computationally unfriendly, it provides an advantageous setting for a semismooth-Newton-based augmented Lagrangian method. The superquantile operator effectively reduces the dimensions of the Newton systems since the tail expectation involves considerably fewer scenarios. Notably, the extra cost of obtaining relevant second-order information and performing matrix inversions is often comparable to, and sometimes even less than, the effort required for gradient computation. Our developed solver is particularly effective when the number of scenarios substantially exceeds the number of decision variables. In synthetic problems with linear and convex diagonal quadratic objectives, numerical experiments demonstrate that our method outperforms existing approaches by a large margin: It achieves speeds more than 750 times faster for linear and quadratic objectives than the alternating direction method of multipliers as implemented by OSQP for computing low-accuracy solutions. Additionally, it is up to 25 times faster for linear objectives and 70 times faster for quadratic objectives than the commercial solver Gurobi, and 20 times faster for linear objectives and 30 times faster for quadratic objectives than the Portfolio Safeguard optimization suite for high-accuracy solution computations. For the quantile regression problem involving over 30 million scenarios, our method computes solution paths up to 20 times faster than Gurobi. The Julia implementation of the solver is available at <https://github.com/jacob-roth/superquantile-opt>.

Keywords: superquantile; semismooth Newton method; sparsity; scenario reduction

1. Introduction

In this paper, we develop a fast, scalable and robust second-order computational solver for the following superquantile-constrained convex optimization problem:

$$\begin{aligned} & \underset{x \in X \subseteq \mathbb{R}^n}{\text{minimize}} && f(x) + \text{CVaR}_{\tau_0}(G^0(x; \omega^{[m]})) \\ & \text{subject to} && \text{CVaR}_{\tau_\ell}(G^\ell(x; \omega^{[m]})) \leq 0, \ell \in \{1, 2, \dots, L\}, \end{aligned} \quad (1)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex and continuously differentiable function, $X \subseteq \mathbb{R}^n$ is a closed convex set with a strongly semismooth and easily computable Euclidean projection (such as a box), and for each $\ell \in \{0, 1, \dots, L\}$, $G^\ell(x; \omega^{[m]}) := \{g^\ell(x; \omega^j)\}_{j=1}^m$ is a collection of random convex continuously differentiable functions $g^\ell(\cdot; \omega)$ evaluated at scenarios $\{\omega^j\}_{j=1}^m$, and

$$\text{CVaR}_{\tau_\ell}(G^\ell(x; \omega^{[m]})) := \inf_t \left\{ t + \frac{1}{(1 - \tau_\ell)m} \sum_{j=1}^m \max(g^\ell(x; \omega^j) - t, 0) \right\} \quad (2)$$

denotes the empirical superquantile (also known as the conditional value-at-risk) of $G^\ell(x; \omega^{[m]})$ at the confidence level $\tau_\ell \in (0, 1)$ (Rockafellar and Uryasev, 1999). In the remainder of this paper, we assume that τ_ℓ is chosen so that $k(\tau_\ell) := (1 - \tau_\ell)m$ is an integer. Since the superquantile operator preserves the convexity, one can easily see that (1) is a convex optimization problem.

The optimization model (1) is flexible and broadly applicable. The minimization of a superquantile objective can be interpreted as distributionally robust optimization with a φ -divergence ambiguity set (Bertsimas and Brown, 2009), and varying the confidence level τ_0 from 0 to 1 in (1) provides a mechanism to interpolate between (risk-neutral) stochastic programs and robust optimization formulations. There is extensive literature on using the superquantile objectives and constraints to model risk-averse preferences and ensure decisions against adverse uncertainty; for an early example, see Pflug (2000). In engineering design, superquantiles help identify potential vulnerabilities (Airaudo et al., 2023), while in statistics, they can be used to formulate the quantile regression problem (Royset and Wets, 2021). Recently, risk-aware models have attracted attention in the machine learning community as alternatives to empirical risk minimization. These models are especially valuable in situations where decision-makers aim to mitigate distribution shifts between training and test datasets (Laguel et al., 2021; Robey et al., 2022), or prioritize performance on certain regions of the distribution, *e.g.*, losses over the left tail to prevent over-fitting or the right tail to exclude outliers (Yuan et al., 2020; Peng et al., 2022). Meanwhile, superquantile constraints have been used in engineering to certify solutions against adverse events as well as adverse computational outcomes (Basova et al., 2011; Chaudhuri et al., 2022), in operations research as the convex, conservative approximation of chance constraints (Nemirovski and Shapiro, 2006; Chen et al., 2010), and in machine learning to promote fairness (Williamson and Menon, 2019; Leqi et al., 2019; Fröhlich and Williamson, 2022). Interested readers are referred to the recent survey of Royset (2022) for a comprehensive overview of risk-aware optimization. Despite the promising numerical results demonstrated by superquantile models in various applications, the absence of a highly stable, efficient, and scalable numerical solver has somewhat limited their widespread adoption.

1. For simplicity of notation, we assume that scenarios $\{\omega_j\}_{j=1}^m$ are the same for all $\ell \in \{0, 1, \dots, L\}$.

In fact, by introducing auxiliary variables $\{u_{\ell j}\}_{j=1,\dots,m}^{\ell=0,\dots,L} \subseteq \mathbb{R}$, one can rewrite (1) in the form of the standard nonlinear optimization problem:

$$\begin{aligned}
 & \underset{\substack{x \in X \subseteq \mathbb{R}^n \\ \{t_\ell\}_{\ell=0,\dots,L}, \{u_{\ell j}\}_{j=1,\dots,m}^{\ell=0,\dots,L}}}{\text{minimize}} && f(x) + t_0 + \frac{1}{(1 - \tau_0)m} \sum_{j=1}^m u_{0j} \\
 & \text{subject to} && t_\ell + \frac{1}{(1 - \tau_\ell)m} \sum_{j=1}^m u_{\ell j} \leq 0, \ell \in \{1, \dots, L\} \\
 & && u_{\ell j} \geq g^\ell(x; \omega^j) - t_\ell, u_{\ell j} \geq 0, j \in \{1, \dots, m\}, \ell \in \{0, \dots, L\}.
 \end{aligned} \tag{3}$$

However, when the number of scenarios m is large, such as in the millions, the above problem is of large-scale with $O(n + mL)$ number of variables and $O(mL)$ number of constraints. Consequently, standard off-the-shelf interior-point-based solvers may not be able to handle it. In such cases, one may directly tackle the compact form of problem (1) with n variables through (stochastic) subgradient-based first-order methods. In this approach, only scenarios from the $(1 - \tau_\ell)$ -tail of $\{g^\ell(x; \omega^j)\}_{j=1}^m$ contribute a non-zero term to the subgradient of $\text{CVaR}_{\tau_\ell}(G^\ell(\cdot; \omega^{[m]}))$. This differs from the risk-neutral expected value of $\{g^\ell(x; \omega^j)\}_{j=1}^m$, where any batch of scenarios can be used to construct an unbiased estimate of the subgradient. To address these computational hurdles, recent research has suggested to take larger mini-batches to increase the likelihood of including a tail scenario in the batch estimate (Yang, 2021), or to adopt a biased sampling strategy (Tamar et al., 2015; Levy et al., 2020). While these methods maintain scalability, they do not fully exploit the inherent sparsity structure offered by superquantile operators, which suggest that only a small subset of scenarios is essential for computation. This issue similarly applies to the smoothing approach of the superquantile operators, where the nonsmooth term $\max(\cdot, 0)$ in (2) is approximated by a smooth function such as the log-sum-exponential. Thus, without care, some smoothing approaches can make all scenarios relevant in the subgradient computation of $\text{CVaR}_{\tau_\ell}(G^\ell(\cdot; \omega^{[m]}))$, detracting from the sparsity that otherwise could be leveraged.

Another line of research for solving problem (1) may be classified as *scenario reduction* strategies. Recall that m represents the total number of collected scenarios. The concept of scenario reduction emerges from the observation that (1) should have an equivalent formulation with a reduced number of scenarios $m_\ell^* \leq m$ for each constraint. Here, the m_ℓ^* number of “effective” scenarios—those that make up the $(1 - \tau_\ell)$ -tail of $G^\ell(x^*; \omega^{[m]})$ at an optimal solution x^* of (1) (e.g., see Rahimian et al., 2018), are not known prior to solving problem (1). This observation leads to a fixed-point approach for iteratively refining the set of active scenarios by excluding those outside the tail, thus leveraging the sparsity inherent in the superquantile operator. Methods that eliminate inactive scenarios based on a candidate solution are known as scenario reduction methods and resemble active-set methods in the optimization literature (Basova et al., 2011; Arp n et al., 2018). The *outer approximation* (OA) scheme, as outlined in Royset and Wets (2021, Chapter 6.C), provides a structured framework for these strategies, allowing the dynamic expansion and pruning of the candidate active scenario set. An advantage of these methods is that the subproblems can be solved in a significantly reduced space, potentially small enough to facilitate the use of off-the-shelf interior-point-based solvers. However, a major challenge in implementing effective scenario strategies lies in the iterative update of the active scenario set, which, based on our experience, often requires manual tuning to achieve good performance.

The goal of this paper is to introduce a scalable and robust second-order computational solver specifically designed to efficiently solve problem (1) at both low and high accuracy levels. Our solver implements a semismooth-Newton-based augmented Lagrangian method (ALM) with a carefully designed linear system solver. We have opted for this method for two primary reasons. (i) Our algorithm is designed to calculate the superquantile terms $\{\text{CVaR}_{\tau_\ell}(G^\ell(x; \omega^{[m]}))\}_{\ell=0}^L$ at every iteration, instead of relying on a subset of scenarios for estimation. This precise computation is especially critical when τ is close to 1 to prevent biased estimations. However, this computation, which involves the evaluation of all $\{g^\ell(x; \omega^j)\}_{j=1}^m$, could be expensive when m is large. To reduce the number of such evaluations, we adopt a second-order method so that the total number of iterations is usually small. (ii) To identify critical scenarios that belong to the $(1 - \tau_\ell)$ -tail of $G^\ell(x; \omega^{[m]})$, it is sufficient to compute the function values of $\{g^\ell(x; \omega^j)\}_{j=1}^m$ rather than their more computationally demanding gradients or Hessians. In each subproblem, the generalized Jacobian matrix in the semismooth Newton method use these function values to identify scenarios likely to influence the optimal $(1 - \tau_\ell)$ -tail. This number may be slightly larger than $(1 - \tau_\ell)m$ since there may exist ties for the $(1 - \tau_\ell)m$ -th largest value of $\{g^\ell(x; \omega^j)\}_{j=1}^m$, but it is usually much smaller than the total number of scenarios m . The inversion of the (generalized) Hessian matrix can thus be done in a small dimensional space using the Sherman-Morrison-Woodbury formula. **As a consequence, our method can be viewed as an implicit scenario reduction approach that uses second-order variational analysis to perform the tuning automatically.**

Our method draws inspiration from recent advances in second-order methods for large-scale optimization, where similar frameworks have been utilized to exploit various forms of structured sparsity, such as in risk-neutral and risk-averse linear regression using ℓ_1 -norm or ℓ_1 - ℓ_2 norm regularization (Li et al., 2018b,a; Zhang et al., 2020; Wu et al., 2023). These studies employ the ALM on the dual problem, leveraging the structured sparsity of the solution to reduce computational costs. In contrast, our method applies the ALM to the primal problem, utilizing the sparsity introduced by superquantiles to achieve efficiency.

Numerical experiments on randomly generated problems with affine random functions $G^\ell(\cdot; \omega^{[m]})$ indicate that the ALM outperforms general-purpose solvers such as Gurobi and OSQP, and a specialized OA method for small instances. In larger scenarios with multiple superquantile constraints or high risk aversion, the performance improvements are substantial, sometimes exceeding 50 times of the aforementioned solvers. In addition, we find that the ALM compares favorably to the specialized Portfolio Safeguard suite in terms of both reliability and speed. A significant advantage of the ALM framework is its ability to efficiently utilize warm-start solutions, enabling a highly efficient process for solving a series of similar problems, such as computing the solution path of (1) across various confidence levels $\{\tau_\ell\}_{\ell=0}^L$. In the context of large-scale quantile regression, ALM proves adept at predicting the full distribution of a response variable across a spectrum of confidence levels. This approach effectively reuses prior solutions, with many effective scenarios in one problem remaining so in subsequent ones, leading to improvements of up to 20 times over existing techniques.

The rest of the paper is organized as follows. In Section 2, we introduce some concepts from nonsmooth analysis and review known results related to the projection onto the top- k -sum constraint. In Section 3, we outline our computational framework to solve (1) using the augmented Lagrangian method, and discuss how to solve the subproblems via the semismooth Newton method. In Section 4, we describe the structure of the resulting linear systems from the semismooth New-

ton method and examine the computational cost associated with solving these systems. Section 5 presents extensive numerical results from both synthetic data and real datasets. The paper ends with a concluding section that summarizes our findings and discusses future research directions.

NOTATION

For a matrix A , the submatrix formed by the rows in an index set \mathcal{I} and the columns in an index set \mathcal{J} is denoted $A_{\mathcal{I},\mathcal{J}}$. We use “:” to denote MATLAB notation for index sets, *e.g.*, $A_{\mathcal{I},:}$ represents the submatrix of A formed by rows \mathcal{I} and all columns. The cardinality of a set \mathcal{I} is denoted by $|\mathcal{I}|$. The vector of all ones in dimension m is denoted by $\mathbf{1}_m$; for an index set $\mathcal{I} \subseteq \{1, \dots, m\}$, $\mathbf{1}_{\mathcal{I}}$ denotes the vector with ones in the indices corresponding to \mathcal{I} and zeros otherwise when the ambient dimension is clear; and e^i denotes the i -th standard basis vector. For a vector $x \in \mathbb{R}^m$, x_i denotes the i -th element of x , $x_{(i)}$ denotes the i -th largest element of x , and $\tilde{x} = x_{\pi}$ denotes a nonincreasing rearrangement of the elements of x according to permutation π with the convention that $\tilde{x}_0 := +\infty$ and $\tilde{x}_{m+1} := -\infty$. For any $x \in \mathbb{R}^m$ and positive integer $1 \leq k \leq m$, there exist integers k_0, k_1 satisfying $0 \leq k_0 \leq k-1$ and $k \leq k_1 \leq m$ such that

$$\tilde{x}_1 \geq \dots \geq \tilde{x}_{k_0} > \tilde{x}_{k_0+1} = \dots = \tilde{x}_k = \dots = \tilde{x}_{k_1} > \tilde{x}_{k_1+1} \geq \dots \geq \tilde{x}_m, \quad (4)$$

with the convention that $k_0 = 0$ if $\tilde{x}_1 = \tilde{x}_k$ and $k_1 = m$ if $\tilde{x}_m = \tilde{x}_k$. We call (k_0, k_1) satisfying (4) the *index-pair of x associated with k* and define the corresponding partition of $\{1, \dots, m\}$ as

$$\alpha := \{1, \dots, k_0\}, \quad \beta := \{k_0 + 1, \dots, k_1\}, \quad \gamma := \{k_1 + 1, \dots, m\}. \quad (5)$$

For any positive integer $k(\leq m)$, we denote the *top- k -sum operator* of $x \in \mathbb{R}^m$ as

$$\mathsf{T}_k(x) := x_{(1)} + \dots + x_{(k)}.$$

In addition, we denote the set of all non-positive top- k -sum vectors as

$$\mathcal{B}_k := \{x \in \mathbb{R}^m : \mathsf{T}_k(x) \leq 0\}. \quad (6)$$

For a closed convex set $S \subseteq \mathbb{R}^n$, its indicator function at a point $x \in \mathbb{R}^n$ is denoted as $\delta_S(x)$, which takes the value 0 if $x \in S$ and $+\infty$ otherwise; the Euclidean projection of x onto S is written as $\text{proj}_S(x) = \arg \min_{y \in S} \|x - y\|_2$. The normal cone of a convex set S at $x \in S$ is denoted by $\mathcal{N}_S(x) := \{d : \langle d, x - y \rangle \geq 0, \forall y \in S\}$. The Fenchel conjugate of a function f is denoted by f^* .

2. Preliminaries

The purpose of this section is to review two known results that will be used in our subsequent development of the ALM framework. The first result addresses the computation of the projection onto the set \mathcal{B}_k defined in (6). The second result further characterizes the Fréchet differentiability of this projection, based on which we are able to derive elements from its generalized Jacobian.

We begin by reviewing some concepts from variational analysis. Interested readers are referred to Facchinei and Pang (2007, Chapters 4 & 7) for more detail. Let $\mathcal{O} \subseteq \mathbb{R}^m$ be an open set. For

any locally Lipschitz function $F : \mathcal{O} \rightarrow \mathbb{R}^n$, it follows from the Rademacher theorem that F is Fréchet-differentiable almost everywhere. We denote the set of Fréchet-differentiable points of F as $\mathcal{D}_F := \{x \in \mathcal{O} : F \text{ is Fréchet-differentiable at } x\}$, and the Jacobian of F at $x \in \mathcal{D}_F$ by $JF(x)$. In addition, we write the Bouligand subdifferential of F at $x \in \mathcal{O}$ as

$$\partial_B F(x) := \{V \in \mathbb{R}^{n \times m} : V = \lim_{k \rightarrow \infty} JF(x^k), x^k \in \mathcal{D}_F, x^k \rightarrow x\},$$

and the Clarke generalized Jacobian of F at $x \in \mathcal{O}$ as $\partial F(x) := \text{conv}\{\partial_B F(x)\}$, i.e., the convex hull of the Bouligand subdifferential at x . The function F is said to be directionally differentiable at $x \in \mathcal{O}$ along the direction d if

$$F'(x; d) := \lim_{t \downarrow 0} \frac{F(x + td) - F(x)}{t}$$

exists. We say F is directionally differentiable at x if it is directionally differentiable along any direction. Under this condition, if it further holds that for any $\Delta x \rightarrow 0$ and $V \in \partial F(x + \Delta x)$, $F(x + \Delta x) - F(x) - V\Delta x = o(\|\Delta x\|_2)$ ($O(\|\Delta x\|_2^2)$), we say F is semismooth (strongly semismooth) at $x \in \mathcal{O}$. Finally, F is (strongly) semismooth on \mathcal{O} if it is (strongly) semismooth for any $x \in \mathcal{O}$.

2.1 Projection onto the set \mathcal{B}_k

Recall the definition of the set \mathcal{B}_k in (6). It is easy to see that when $(1 - \tau_\ell)m$ is an integer, the constraint $\text{CVaR}_{\tau_\ell}(G^\ell(x; \omega^{[m]})) \leq 0$ in (1) can be equivalently written as $G^\ell(x; \omega^{[m]}) \in \mathcal{B}_{k_\ell}$ where $k_\ell = (1 - \tau_\ell)m$. Given a vector $y \in \mathbb{R}^m$, its projection onto \mathcal{B}_k , denoted by

$$\tilde{y} := \text{proj}_{\mathcal{B}_k}(y) = \underset{y' \in \mathbb{R}^m}{\text{argmin}} \left\{ \frac{1}{2} \|y' - y\|_2^2 : y' \in \mathcal{B}_k \right\}, \quad (7)$$

can be computed in $O(m \log m)$ operations (Wu et al., 2014; Davis, 2015; Li and Li, 2021; Roth and Cui, 2025). For special cases when $k = 1$ or $k = m$, the projection has a closed form and can be computed in $O(m)$ operations. Excluding the two special cases, all of these computations use the fact that the top- k -sum operator T_k is permutation-invariant. Therefore, one may instead first sort y to y_π such that $(y_\pi)_1 \geq \dots \geq (y_\pi)_m$, and then solve the following problem:

$$\bar{y} := \underset{y \in \mathbb{R}^m}{\text{argmin}} \left\{ \frac{1}{2} \|y - y_\pi\|_2^2 : y \in \mathcal{B}_k \right\} = \underset{y \in \mathbb{R}^m}{\text{argmin}} \left\{ \frac{1}{2} \|y - y_\pi\|_2^2 : y_1 \geq y_2 \geq \dots \geq y_m, \sum_{i=1}^k y_i \leq 0 \right\}. \quad (8)$$

In this case, the solution \bar{y} must satisfy the following Karush–Kuhn–Tucker (KKT) optimality conditions for some unique indices \bar{k}_0 and \bar{k}_1 and some multipliers $(\bar{\mu}, \bar{\lambda}, \bar{\theta}) \in \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$:

$$\begin{aligned} \bar{y}_{\bar{\alpha}} &= \bar{y}_{\bar{\alpha}} - \bar{\lambda} \mathbf{1}_{\bar{\alpha}} & \bar{\mu}_{\bar{\alpha}} &= \mathbf{1}_{\bar{\alpha}} \\ \bar{y}_{\bar{\beta}} &= \bar{y}_{\bar{\beta}} - \bar{\lambda} \bar{\mu}_{\bar{\beta}} = \bar{\theta} \mathbf{1}_{\bar{\beta}} & \bar{\mu}_{\bar{\beta}} &\in \left\{ \mu \in \mathbb{R}^{|\bar{\beta}|} : 0 \leq \mu \leq \mathbf{1}, \mathbf{1}^\top \mu = k - \bar{k}_0 \right\} \\ \bar{y}_{\bar{\gamma}} &= \bar{y}_{\bar{\gamma}} & \bar{\mu}_{\bar{\gamma}} &= 0 \\ 0 &= \mathbf{1}_{\bar{\alpha}}^\top \bar{y}_{\bar{\alpha}} - \bar{k}_0 \bar{\lambda} + (k - \bar{k}_0) \bar{\theta} & \bar{\lambda} &> 0, \quad \bar{y}_{\bar{k}_0} > \bar{\theta} > \bar{y}_{\bar{k}_1+1}. \end{aligned} \quad (9)$$

where the sets $\{\bar{\alpha}, \bar{\beta}, \bar{\gamma}\}$ are defined according to (5) based on the index-pair (\bar{k}_0, \bar{k}_1) . The pivoting procedures in Roth and Cui (2025) compute a (sorted) solution to (9) by identifying the optimal index-pair (\bar{k}_0, \bar{k}_1) satisfying the order structure (4). It was shown there that the total cost to compute \bar{y} is $O(m + \bar{k}_1 \log(m))$, which is achieved by integrating the sorting procedure into the pivoting procedure after an initial heapsort construction, allowing for the dynamic computation of \bar{k}_1 . This approach makes use of the max-heap data structure (*e.g.*, see Williams, 1964), which can be constructed with a cost of $O(m)$, and offers substantial speed advantages over a full sort especially when $k \ll m$.

2.2 Differentiability of the projection map $\text{proj}_{\mathcal{B}_k}$

In this section, we discuss the differentiability of the projection map $\text{proj}_{\mathcal{B}_k}$, which is needed in the construction of the generalized Jacobian of this projection. Our derivations are adapted from the work Wu et al. (2014), which focused on the projection onto the vector k -norm ball.

Based on the definition of the Fréchet differentiability, the projection map $\text{proj}_{\mathcal{B}_k}$ is differentiable at a point $y \in \mathbb{R}^m$ when its one-sided directional derivative $\text{proj}'_{\mathcal{B}_k}(y; \cdot)$ is linear in the second argument. The latter directional derivative can be calculated by projecting the direction onto the critical cone of \mathcal{B}_k at y (Facchinei and Pang, 2007, Theorem 4.1.1). Our subsequent derivations are based on these observations.

For convenience, we will consider a sorted point $y \in \mathbb{R}^m$ with the permutation π such that $y = y_\pi = \tilde{y}$. In general, one may first obtain the permutation π and then apply the following derivations. In the sorted case, the critical cone of \mathcal{B}_k at y is given by

$$\mathcal{C}_{\mathcal{B}_k}(y) := \mathcal{T}_{\mathcal{B}_k}(\bar{y}) \cap (\bar{y} - y)^\perp \quad \text{with } \bar{y} = \text{proj}_{\mathcal{B}_k}(y),$$

where $\mathcal{T}_{\mathcal{B}_k}(\bar{y})$ is the tangent cone of \mathcal{B}_k at \bar{y} and v^\perp represents the orthogonal space of a vector v . The KKT conditions (9) of the sorted projection problem (8) gives

$$(\bar{y} - y)^\perp = \{d \in \mathbb{R}^m : \mathbf{1}_{\bar{\alpha}}^\top d_{\bar{\alpha}} + \bar{\mu}_{\bar{\beta}}^\top d_{\bar{\beta}} = 0\}, \quad \bar{\mu}_{\bar{\beta}} \in [0, 1]^{|\bar{\beta}|}, \quad \mathbf{1}^\top \bar{\mu}_{\bar{\beta}} = k - k_0,$$

where $\{\bar{\alpha}, \bar{\beta}, \bar{\gamma}\}$ are the partition of the indices associated with \bar{y} in (5); cf. eq. (44) in Wu et al. (2014). This holds since $\bar{y} = y - \bar{\lambda} \bar{\mu}$ for $\bar{\mu} \in \partial \mathcal{T}_k(\bar{y})$ and $\bar{\lambda} > 0$. By direct computation, the tangent cone at any $\bar{y} \in \mathcal{B}_k$ with indices $\{\bar{\alpha}, \bar{\beta}, \bar{\gamma}\}$ is given by

$$\mathcal{T}_{\mathcal{B}_k}(\bar{y}) = \begin{cases} \mathbb{R}^m, & \text{if } \mathcal{T}_k(\bar{y}) < 0, \\ \{d \in \mathbb{R}^m : \mathbf{1}_{\bar{\alpha}}^\top d_{\bar{\alpha}} + \mathcal{T}_{k-\bar{k}_0}(d_{\bar{\beta}}) \leq 0\}, & \text{if } \mathcal{T}_k(\bar{y}) = 0; \end{cases}$$

cf. eq. (46) in Wu et al. (2014). Therefore, when $\mathcal{T}_k(y) > 0$ so that $\mathcal{T}_k(\bar{y}) = 0$, the critical cone can be expressed as

$$\begin{aligned} \mathcal{C}_{\mathcal{B}_k}(y) &= \{d \in \mathbb{R}^m : \mathbf{1}_{\bar{\alpha}}^\top d_{\bar{\alpha}} + \mathcal{T}_{k-\bar{k}_0}(d_{\bar{\beta}}) \leq 0, \mathbf{1}_{\bar{\alpha}}^\top d_{\bar{\alpha}} + \bar{\mu}_{\bar{\beta}}^\top d_{\bar{\beta}} = 0\} \\ &= \{d \in \mathbb{R}^m : \mathcal{T}_{k-\bar{k}_0}(d_{\bar{\beta}}) \leq \bar{\mu}_{\bar{\beta}}^\top d_{\bar{\beta}}, \mathbf{1}_{\bar{\alpha}}^\top d_{\bar{\alpha}} + \bar{\mu}_{\bar{\beta}}^\top d_{\bar{\beta}} = 0\}. \end{aligned} \quad (10)$$

When $\mathcal{T}_k(y) = 0$ so that $y = \bar{y}$, we have $y - \bar{y} = \{0\}$ and $\mathcal{C}_{\mathcal{B}_k}(y)$ coincides with the tangent cone $\mathcal{T}_{\mathcal{B}_k}(\bar{y})$. The next result states the conditions under which the metric projector is differentiable, i.e., when the critical cone $\mathcal{C}_{\mathcal{B}_k}(y)$ is a linear subspace. It is readily obtained from conditions (i), (iii), and (iv) of Wu et al. (2014, Theorem 4.2).

Lemma 1 *Let $1 \leq k \leq m$, $y = \tilde{y} \in \mathbb{R}^m$ be a given sorted vector, and \bar{y} be its projection with associated index-pair (\bar{k}_0, \bar{k}_1) in (4) and index-sets $\{\bar{\alpha}, \bar{\beta}, \bar{\gamma}\}$ in (5). Let $(\bar{y}, \bar{\mu}, \bar{\lambda}, \bar{\theta})$ satisfy the optimality conditions in (9). Then $\text{proj}_{\mathcal{B}_k}(\cdot)$ is differentiable at y if and only if y satisfies one of the following three properties:*

- (a) $\mathsf{T}_k(y) < 0$;
- (b) $\mathsf{T}_k(y) > 0$, $k < \bar{k}_1$, and $y_{\bar{k}_0+1} - \bar{\lambda} < \bar{y}_k \equiv \bar{\theta} < y_{\bar{k}_1}$;
- (c) $\mathsf{T}_k(y) > 0$ and $k = \bar{k}_1$.

Next we use Lemma 1 to compute $J_{\mathcal{B}_k}$, the Jacobian of the map $\text{proj}_{\mathcal{B}_k}$ at differentiable points corresponding to the three cases in the above lemma. For simplicity, we assume that $y = \tilde{y}$ is sorted. For any positive integer p , we write I_p as the $p \times p$ identity matrix. The proof of the following lemma, which is adapted from Wu et al. (2014, eq. (68) in Section 4.3.4), is given in Appendix A.

Lemma 2 *Let $1 \leq k \leq m$, $y = \tilde{y} \in \mathbb{R}^m$ be a given sorted vector and \bar{y} be its projection with associated index-pair (\bar{k}_0, \bar{k}_1) in (4) and index-sets $\{\bar{\alpha}, \bar{\beta}, \bar{\gamma}\}$ in (5).*

- (a) *If $\mathsf{T}_k(y) < 0$, then $\text{proj}'_{\mathcal{B}_k}(y; d) = d$ for any $d \in \mathbb{R}^m$ and $J_{\mathcal{B}_k} \in \mathbb{R}^{m \times m}$ is the identity matrix.*
- (b) *If $\mathsf{T}_k(y) > 0$, $k < \bar{k}_1$ and $y_{\bar{k}_0+1} - \bar{\lambda} < \bar{\theta} < y_{\bar{k}_1}$. Then*

$$J_{\mathcal{B}_k} = \begin{bmatrix} I_{|\bar{\alpha}|+|\bar{\beta}|} - B^\top(BB^\top)^{-1}B & \\ & I_{|\bar{\gamma}|} \end{bmatrix}, \quad (11)$$

where

$$B := \begin{bmatrix} \mathbf{1}_{\bar{\alpha}}^\top & \bar{\mu}_{\bar{\beta}}^\top \\ 0_{(|\bar{\beta}|-1), |\bar{\alpha}|} & C \end{bmatrix} \in \mathbb{R}^{(|\bar{\alpha}|+|\bar{\beta}|-1) \times (|\bar{\alpha}|+|\bar{\beta}|)}, \quad C := \begin{bmatrix} 1 & -1 & & \\ & \ddots & \ddots & \\ & & 1 & -1 \end{bmatrix} \in \mathbb{R}^{(|\bar{\beta}|-1) \times |\bar{\beta}|} \quad (12)$$

and the matrix BB^\top is invertible.

- (c) *If $\mathsf{T}_k(y) > 0$ and $k = \bar{k}_1$, then*

$$J_{\mathcal{B}_k} = \begin{bmatrix} I_{|\bar{\alpha}|+|\bar{\beta}|} - \frac{cc^\top}{\|c\|_2^2} & \\ & I_{|\bar{\gamma}|} \end{bmatrix} \quad \text{with } c := \mathbf{1}_{(\bar{\alpha}, \bar{\beta})}.$$

Finally, since the locally Lipschitz continuous map $\text{proj}_{\mathcal{B}_k}$ is differentiable almost everywhere, any point $y \in \mathbb{R}^m$ can be viewed as the limit of a sequence of differentiable points. Hence, one can further obtain an element from the generalized Jacobian of $\text{proj}_{\mathcal{B}_k}$ at any point y based on Lemmas 1 and 2. See Table 1 for an explicit construction, which relies on the following cases.

When $\mathsf{T}_k(y) = 0$: The map $\text{proj}_{\mathcal{B}_k}$ is differentiable for case 1 in Table 1. Observe that cases 2 and 4 cannot occur since $\mathsf{T}_k(y) = 0$ and $(<, <)$ pattern is impossible.² For cases 3 and 5, (a) can be

2. If $\mathsf{T}_k(y) = 0$, then $\bar{\lambda} = 0$ since otherwise $\sum_{i=1}^k \bar{y} < r = \sum_{i=1}^k y_i$, contradicting optimality. This implies that $\bar{y} = y$. Since $k \in \{\bar{k}_0 + 1, \dots, \bar{k}_1\}$ by construction, $\bar{\theta} = \bar{y}_k = \bar{y}_{\bar{k}_0+1} = \bar{y}_{\bar{k}_1}$, so strict inequalities cannot hold.

obtained by a sequence $y^\nu \uparrow y$ with $y_i^\nu \leq y_i$ and at least one i' with $y_{i'}^\nu < y_{i'}'$ for $i \in \{1, \dots, k\}$ that preserves sorting for all i as $\nu \rightarrow \infty$.

When $T_k(y) > 0$: The map $\text{proj}_{\mathcal{B}_k}$ is differentiable for cases 6, 8, and 9 in Table 1. For case 7, (c) can be obtained by a sequence $y^\nu \rightarrow y$ that differs only in the $\bar{\beta}$ components with $y_i^\nu - \epsilon_i$ for $i \in \bar{\beta} \equiv \{\bar{k}_0 + 1, \dots, \bar{k}_1\}$ satisfying a strictly decreasing ordering $\epsilon_i < \epsilon_j$ for $i < j \in \bar{\beta}$. For $\max_i \epsilon_i$ small enough, this ensures that $\bar{k}_1 = k$ by breaking any existing ties.

#	$T_k(y) \ ? \ 0$	$k \ ? \ \bar{k}_1$	$y_{\bar{k}_0+1} - \bar{\lambda} \ ? \ \bar{y}_k \ ? \ y_{\bar{k}_1}$	active selection
1 [*]	<	—	—	(a)
2 [†]	=	<	<, <	—
3	=	<	$\neg(<, <)$	(a)
4 [†]	=	=	<, <	—
5	=	=	$\neg(<, <)$	(a)
6	>	<	<, <	(b)
7	>	<	$\neg(<, <)$	(c)
8 [*]	>	=	<, <	(c)
9 [*]	>	=	$\neg(<, <)$	(c)

^{*} differentiable
[†] cannot occur

Table 1: Selection from the set of generalized Jacobian for different situations in Algorithm 1.

3. The algorithmic framework

In this section, we introduce our algorithmic framework to solve the superquantile-based problem (1). Since one can always move the superquantile term $\text{CVaR}_{\tau_0}(G^0(x; \omega^{[m]}))$ in the objective function to the constraint set using the epigraphic transformation, we assume for simplicity that $G^0(x; \omega^{[m]}) = \{0\}$ in (1) so that the objective function only consists of the convex smooth term $f(x)$. By introducing auxiliary variables $\{y^\ell\}_{\ell=1}^L \subseteq \mathbb{R}^m$ and $z \in \mathbb{R}^n$, one can reformulate problem (1) into

$$\begin{aligned}
 \min_{x, y, z} \quad & f(x) \\
 \text{s.t.} \quad & y^\ell \geq (g^\ell(x; \omega^1), \dots, g^\ell(x; \omega^m))^\top, \quad y^\ell \in \mathcal{B}_{k_\ell}, \quad \ell \in \{1, \dots, L\}, \\
 & z = x, \quad z \in X,
 \end{aligned} \tag{13}$$

where $k_\ell = (1 - \tau_\ell)m$ and \mathcal{B}_k is the non-positive top- k -sum set defined in (6). For exposition, we consider $X = [p, q] \subseteq \mathbb{R}^n$ to be a box with $p \leq q \in \{\mathbb{R} \cup \pm\infty\}^n$. Our overall algorithm has double loops, where the outer loop uses the ALM to deal with the constraints, and the inner loop adopts the semismooth Newton method to solve the ALM subproblems.

3.1 An augmented Lagrangian framework to solve problem (13)

For notational simplicity, we write $y = ((y^1)^\top, \dots, (y^L)^\top)^\top \in \mathbb{R}^{mL}$,

$$G(x) := \left(\underbrace{g^1(x; \omega^1), \dots, g^1(x; \omega^m)}_{G^1(x; \omega^{[m]})}, \dots, \underbrace{g^L(x; \omega^1), \dots, g^L(x; \omega^m)}_{G^L(x; \omega^{[m]})} \right)^\top : \mathbb{R}^n \rightarrow \mathbb{R}^{mL},$$

and $\bar{\mathcal{B}} := \mathcal{B}_{k_1} \times \dots \times \mathcal{B}_{k_L} \subseteq \mathbb{R}^{mL}$. Let $\lambda(\geq 0) \in \mathbb{R}^{mL}$ be the multiplier of the inequality constraint $y \geq G(x)$ and $\mu \in \mathbb{R}^n$ be the multiplier of the equality constraint $z = x$. The augmented Lagrangian function associated with problem (13) for a parameter $\sigma > 0$ takes the form of

$$L_\sigma(x, y, z; \lambda, \mu) := f(x) + \frac{\sigma}{2} \|\max(G(x) + \lambda/\sigma - y, 0)\|_2^2 + \frac{\sigma}{2} \|x + \mu/\sigma - z\|_2^2 + \text{const}$$

for $x \in \mathbb{R}^n$, $y \in \bar{\mathcal{B}}$ and $z \in X$, where **const** is a constant depending on the multipliers. To make the subproblems easier to solve, at the ν -th iteration, we may also add a proximal term $\frac{1}{2\sigma} \|x - x^\nu\|_M^2 := \frac{1}{2\sigma} (x - x^\nu)^\top M (x - x^\nu)$ for some positive semidefinite matrix M to regularize the augmented Lagrangian function. The overall algorithm is summarized in Algorithm 1.

Algorithm 1 The ALM framework for problem (13)

Initialize: Choose $x^0 = z^0 \in \mathbb{R}^n$, $y^0, \lambda^0 \in \mathbb{R}^{mL}$, and $\mu \in \mathbb{R}^n$. Choose $\sigma_0 > 0$ and a positive semidefinite matrix M . Set iteration $\nu = 0$, and repeat the following steps until a proper stopping criterion is satisfied.

1: Primal update:

$$(x^{\nu+1}, y^{\nu+1}, z^{\nu+1}) \approx \arg \min_{x, y \in \bar{\mathcal{B}}, z \in X} L_{\sigma_\nu}(x, y, z; \lambda^\nu, \mu^\nu) + \frac{1}{2\sigma_\nu} \|x - x^\nu\|_M^2. \quad (14)$$

2: Dual update:

$$\begin{aligned} \lambda^{\nu+1} &:= \max(\lambda^\nu + \sigma_\nu(G(x^{\nu+1}) - y^{\nu+1}), 0) \\ \mu^{\nu+1} &:= \mu^\nu + \sigma_\nu(x^{\nu+1} - z^{\nu+1}). \end{aligned}$$

3: Check termination criteria and update σ_ν .

Let $JG(x) \in \mathbb{R}^{(mL) \times n}$ denote the Jacobian matrix of G at x . The KKT optimality condition of problem (13) is

$$\begin{cases} 0 = \nabla f(x) + JG(x)^\top \lambda + \mu \\ \lambda \in \mathcal{N}_{\bar{\mathcal{B}}}(y), \quad \mu \in \mathcal{N}_X(z) \\ 0 \leq \lambda \perp y - G(x) \geq 0, \quad x - z = 0. \end{cases} \quad (15)$$

The above condition is necessary and sufficient to characterize a solution to the convex problem (13) when the Slater condition holds.

It is known that under mild conditions, any accumulation point of the sequence $\{(x^\nu, y^\nu, z^\nu, \lambda^\nu, \mu^\nu)\}$ generated by the ALM converges to a tuple satisfying the KKT condition (15) (*e.g.*, see Rockafellar, 1976a). Conditions ensuring the local superlinear convergence rate of ALM are also well studied in

the literature; see the early seminal works by Rockafellar (1976a,b) and the later refinement by Cui et al. (2019). For example, consider the following constraint non-degeneracy condition at a tuple $(\bar{x}, \bar{y}, \bar{z})$ due to Robinson (1984, 1987):

$$0 \in \text{int}\{F(\bar{x}, \bar{y}, \bar{z}) + JF(\bar{x}, \bar{y}, \bar{z}) (\mathbb{R}^n \times \mathbb{R}^{mL} \times \mathbb{R}^n) - \mathbb{R}_+^{mL} \times 0_n \times \bar{\mathcal{B}} \times X\},$$

where $F(x, y, z) := (y - G(x), z - x, y, z)$, int denotes the interior of a set and \mathbb{R}_+^{mL} denotes the (mL) -dimensional non-negative orthant. See Appendix B for a discussion of the non-degeneracy assumption based on a simplified form of problem (13). The convergence result in Rockafellar (1976a) implies that if this condition is satisfied at an optimal solution (x^*, y^*, z^*) , the sequence of dual variables $\{(\lambda^\nu, \mu^\nu)\}$ generated by the ALM converges at an asymptotically superlinear rate to the unique multiplier when subproblems are solved to a sufficient accuracy. A proper termination condition to ensure the rate is detailed in the numerical experiments (Section 5).

3.2 A semismooth Newton method for subproblem (14)

The effectiveness of the ALM for solving problem (1) hinges on the efficient computation of the subproblem (14). In this section, we present a semismooth Newton method to solve this subproblem. Despite the seemingly high computational cost typically associated with such procedures, the actual cost for the superquantile-constrained problem is in fact quite low. This efficiency is largely attributed to the unique structural properties brought by the projection onto the non-positive top- k -sum set $\bar{\mathcal{B}}$.

For exposition, we begin by considering the augmented Lagrangian subproblem associated with the (possibly nonconvex) equality constraint $y = G(x)$ before returning to the convex inequality case. For fixed λ^ν and μ^ν , consider the subproblem

$$(x^{\nu+1}, y^{\nu+1}, z^{\nu+1}) \in \arg \min_{x, y \in \bar{\mathcal{B}}, z \in X} \left\{ f(x) + \frac{\sigma_\nu}{2} \|y - (G(x) + \sigma_\nu^{-1} \lambda^\nu)\|_2^2 + \frac{\sigma_\nu}{2} \|z - (x + \sigma_\nu^{-1} \mu^\nu)\|_2^2 + \frac{1}{2\sigma_\nu} \|x - x^\nu\|_M^2 \right\} \quad (16)$$

for some positive semidefinite matrix M which may help to regularize the subproblem. Observe that the unique solution in terms of (y, z) satisfies

$$y^{\nu+1} = \text{proj}_{\bar{\mathcal{B}}}(G(x^{\nu+1}) + \sigma_\nu^{-1} \lambda^\nu), \quad z^{\nu+1} = \text{proj}_X(x^{\nu+1} + \sigma_\nu^{-1} \mu^\nu). \quad (17)$$

Substituting the above solutions into (16) for y and z gives the problem in terms of x alone:

$$x^{\nu+1} \in \arg \min_x \left\{ \varphi_{\sigma_\nu}(x) := f(x) + \frac{\sigma_\nu}{2} \|G(x) + \sigma_\nu^{-1} \lambda^\nu - \text{proj}_{\bar{\mathcal{B}}}(G(x) + \sigma_\nu^{-1} \lambda^\nu)\|_2^2 + \frac{\sigma_\nu}{2} \|x + \sigma_\nu^{-1} \mu^\nu - \text{proj}_X(x + \sigma_\nu^{-1} \mu^\nu)\|_2^2 + \frac{1}{2\sigma_\nu} \|x - x^\nu\|_M^2 \right\}. \quad (18)$$

Therefore, the function φ_{σ_ν} may serve as the subproblem objective, which has a Lipschitz continuous gradient that takes the following form:

$$\begin{aligned} \nabla \varphi_{\sigma_\nu}(x) &= \nabla f(x) + \sigma_\nu JG(x)^\top [G(x) + \sigma_\nu^{-1} \lambda^\nu - \text{proj}_{\bar{\mathcal{B}}}(G(x) + \sigma_\nu^{-1} \lambda^\nu)] \\ &\quad + \sigma_\nu [x + \sigma_\nu^{-1} \mu^\nu - \text{proj}_X(x + \sigma_\nu^{-1} \mu^\nu)] + \sigma_\nu^{-1} M(x - x^\nu). \end{aligned} \quad (19)$$

Since $\bar{\mathcal{B}}$ and X are polyhedral sets, the projections onto them are strongly semismooth (Shapiro, 1988). Therefore, one can apply the semismooth Newton method to solve the nonsmooth equation $\nabla \varphi_{\sigma_\nu}(x) = 0$ in (19) and use the line search for globalization.

The following lemma, whose proof is deferred to Appendix A, shows that we may perform the same procedure to solve the augmented Lagrangian subproblem associated with *inequality* constraints.

Lemma 3 *Let $g \in \mathbb{R}^m$ be given. Let \mathcal{B}_k be the non-positive top- k -sum set defined in (6). Consider the set*

$$\mathcal{P} := \left\{ y^* : y^* \in \arg \min_y \left\{ \|\max(\tilde{g} - y, 0)\|_2^2 : y \in \mathcal{B}_k \right\} \right\} \quad (20)$$

and the element

$$\bar{y} = \arg \min_y \left\{ \|\tilde{g} - y\|_2^2 : y \in \mathcal{B}_k \right\}. \quad (21)$$

Then it holds that $\bar{y} \in \mathcal{P}$. In fact, it holds that $\tilde{y} \in \mathcal{P}$ for any \tilde{y} such that $(\tilde{y})_{1:\bar{k}_1} = (\bar{y})_{1:\bar{k}_1}$ with $(\tilde{y})_{\bar{k}_1+1:m} \leq (\bar{y})_{\bar{k}_1+1:m}$, where \bar{k}_1 associated with \bar{y} is defined in (4).

Therefore, given λ^ν and μ^ν , instead of (18) for the equality constrained case, we may consider the following objective function

$$\begin{aligned} \tilde{\varphi}_{\sigma_\nu}(x) := & f(x) + \frac{\sigma_\nu}{2} \left\| \max(G(x) + \sigma_\nu^{-1} \lambda^\nu - \text{proj}_{\bar{\mathcal{B}}}(G(x) + \sigma_\nu^{-1} \lambda^\nu), 0) \right\|_2^2 \\ & + \frac{\sigma_\nu}{2} \left\| (x + \sigma_\nu^{-1} \mu^\nu) - \text{proj}_X(x + \sigma_\nu^{-1} \mu^\nu) \right\|_2^2 + \frac{1}{2\sigma_\nu} \|x - x^\nu\|_M^2. \end{aligned}$$

Observe that if $x^+ \in \arg \min_x \tilde{\varphi}_{\sigma_\nu}(x)$, we may take

$$y^+ := \text{proj}_{\bar{\mathcal{B}}}(G(x^+) + \sigma_\nu^{-1} \lambda^\nu), \quad z^+ := \text{proj}_X(x^+ + \sigma_\nu^{-1} \mu^\nu)$$

to obtain a solution to the ALM subproblem

$$\begin{aligned} (x^+, y^+, z^+) \in & \arg \min_{x, y \in \bar{\mathcal{B}}, z \in X} f(x) + \frac{\sigma_\nu}{2} \left\| \max(G(x) + \sigma_\nu^{-1} \lambda^\nu - y, 0) \right\|_2^2 \\ & + \frac{\sigma_\nu}{2} \left\| x + \sigma_\nu^{-1} \mu^\nu - z \right\|_2^2 + \frac{1}{2\sigma_\nu} \|x - x^\nu\|_M^2. \end{aligned}$$

Though there may exist multiple solutions y^+ , the value of the subproblem objective coincides with $\tilde{\varphi}_{\sigma_\nu}(x^+)$ for any optimal (y^+, z^+) . Therefore, the ALM framework can readily handle inequality constraints. In the inequality constrained case, the first- and second-order information of the function φ can be computed analogous to (19) and (23) by replacing $JG(x)$ with $I_{G^+} \cdot JG(x)$, where I_{G^+} is the diagonal matrix with $(I_{G^+})_{ii} = 1$ when $G_i(x) - y_i > 0$ and zero otherwise. In addition, the semismooth Newton method summarized in Algorithm 2 can be used to obtain x^+ , and as a result, y^+ and z^+ .

Below we outline the semismooth Newton algorithm with line search to solve the nonsmooth equation (16), which is the optimality condition of the ν -th subproblem of the ALM. Throughout the statement of this algorithm, we maintain ν as the fixed outer iterate of the ALM, and write the semismooth Newton iteration as $t \geq 0$.

Algorithm 2 A semismooth Newton method for the ν -th subproblem (16)

Initialize: Set $t = 0$, choose an initial point $x^{\nu,0}$, and choose backtracking line search parameters $\delta, c \in (0, 1)$. Perform the following steps until a convergence criterion is reached.

- 1: Semismooth Newton direction: using direct factorization or an indirect method, obtain an approximate solution d^t to the linear system

$$Vd + \nabla\varphi_{\sigma_\nu}(x^{\nu,t}) = 0, \quad (22)$$

where $V \in \partial(\nabla\varphi_{\sigma_\nu}(x^{\nu,t}))$.

- 2: Line search (backtracking): Set $\rho_t = \rho$, where ρ is the smallest nonnegative integer such that

$$\varphi_{\sigma_\nu}(x^{\nu,t} + \delta^\rho d^t) \leq \varphi_{\sigma_\nu}(x^{\nu,t}) + c \delta^\rho \nabla\varphi_{\sigma_\nu}(x^{\nu,t})^\top d^t.$$

- 3: Update solution:

$$x^{\nu,t+1} = x^{\nu,t} + \delta^{\rho_t} d^t.$$

In fact, it is not easy to obtain an element from $\partial(\nabla\varphi_{\sigma_\nu}(x))$ in the equation (22) even if $G(x)$ is an affine function, see, for example, the discussion in Dolgopolik (2024). One remedy to overcome this difficulty is to construct an element V in the following way:

$$V := \nabla^2 f(x) + \sum_{i=1}^{mL} \nabla^2 G_i(x) \cdot \left(G(x) + \sigma_\nu^{-1} \lambda^\nu - \text{proj}_{\bar{B}}(G(x) + \sigma_\nu^{-1} \lambda^\nu) \right)_i + \sigma_\nu JG(x)^\top \cdot (I - J_{\bar{B}}) \cdot JG(x) + \sigma_\nu (I - J_X) + \sigma_\nu^{-1} M \quad \Bigg\}, \quad (23)$$

$$\text{with } J_{\bar{B}} \in \partial \text{proj}_{\bar{B}}(G(x) - \sigma_\nu^{-1} \lambda), \quad J_X \in \partial \text{proj}_X(x - \sigma_\nu^{-1} \mu).$$

An explicit form for $J_{\bar{B}}$ is given in Lemma 2, and the matrix J_X is a diagonal matrix with the i -th diagonal element given by

$$[J_X]_{ii} = \begin{cases} 1, & \text{if } x_i - \mu_i/\sigma_\nu \in [p_i, q_i]; \\ 0, & \text{otherwise.} \end{cases}$$

Although such a matrix V may not belong to $\partial(\nabla\varphi_{\sigma_\nu}(x))$, following the proof of Li et al. (2018b), it can be shown that $Vd \in \widehat{\partial}(\nabla\varphi_{\sigma_\nu}(x))(d)$ for some upper-semicontinuous and convex-valued set-valued map $\widehat{\partial}(\nabla\varphi_{\sigma_\nu}(\cdot))$, which is a superset of $\partial(\nabla\varphi_{\sigma_\nu}(x))$ at any x . Furthermore, the gradient function $\nabla\varphi_{\sigma_\nu}$ is semismooth with respect to $\widehat{\partial}(\nabla\varphi_{\sigma_\nu}(\cdot))$. Therefore, one can safely use the matrix V in (23) in the semismooth Newton equation (22). The details of how to efficiently solve this linear equation are deferred to the next section.

Before ending this section, we highlight a few extensions of problem (1) that can be similarly handled by our proposed semismooth-Newton based ALM. Consider a problem involving only a single superquantile objective. Let $\omega^j := \{(A_{j,:}, b_j)\}_{j=1}^m$, where in the context of the supervised learning problem, we interpret $b_j \in \mathbb{R}$ as the response and $A_{j,:}^\top \in \mathbb{R}^n$ as the covariates of the j -th observation. The risk-averse regression problem takes the form of

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \text{CVaR}_\tau \left(\psi(A_{j,:}^\top x - b_j) \right)_{j=1}^m,$$

where $\psi : \mathbb{R} \rightarrow \mathbb{R}$ measures the individual loss for each covariate-response pair, such as $\psi(t) = t^2$ for the regression problem. This problem is known as the performance at the top or hard example mining for (imbalanced) classification in machine learning (Lyu et al., 2020; Shrivastava et al., 2016; Bengio et al., 2009; Yang, 2022). Our proposed framework can directly handle convex continuously differentiable loss functions. In some cases, it is preferable to consider nonsmooth loss functions such as the absolute loss $\psi(t) = |t|$ or the hinge loss $\psi(t) = \max(1 - t, 0)$. It is then convenient to consider a redefined risk set associated with a composite (empirical) superquantile operator $\tilde{B}_{k(\tau),r} := \{y \in \mathbb{R}^m : \sum_{i=1}^{k(\tau)} \psi(y_i) \leq r\}$ for some $r \geq 0$, where $k(\tau) = (1 - \tau)m$. The reason that we consider the shifted set with $r \geq 0$ is because the newly defined operator CVaR_τ may not be translation equivariant. The projections onto these modified sets and their Jacobians need to be computed accordingly, but they retain similar sparsity structures to those detailed below in Section 4.1. In particular, for the absolute value loss, the composite superquantile operator corresponds to the vector- k -norm, where an explicit form for the derivative of the associated projection operator is given in Wu et al. (2014). We omit the detail here for brevity.

4. Implicit scenario reduction via the semismooth Newton equation

In this section, we explore how the inherent nonsmoothness of the superquantile operator enables implicit scenario reduction within the augmented Lagrangian subproblems, resulting in significant computational efficiencies in solving the linear equations of the semismooth Newton method. For simplicity, we focus our discussion throughout this section to a single superquantile constraint (i.e., $L = 1$) and omit the superscript ℓ . Each $g(x; \omega^j) := A_{j,:}x + b_j$ is an affine function with $\omega^j := \{(A_{j,:}, b_j)\}$ so that the constraint $\text{CVaR}_\tau(G(x; \omega^{[m]})) \leq 0$ can be written as $\mathbb{T}_k(Ax + b) \leq 0$ for $k = (1 - \tau)m$, or equivalently, $Ax + b \in \mathcal{B}_k$. Extending the discussion to multiple superquantile constraints with $L > 1$ is straightforward, using the appropriate concatenations of A^ℓ , b^ℓ , and the Cartesian product of the sets $\{\mathcal{B}_{k_\ell}\}$.

4.1 Structured sparsity of the linear system (22)

Since the function G is assumed to be affine in x , we consider the objective function of the ALM subproblem to be φ_σ as in (18) to simplify the discussion. The matrix V in (23) then takes the form of

$$V = \nabla^2 f(x) + \sigma_\nu A^\top (I - J_{\mathcal{B}_k}) A + \sigma_\nu (I - J_X) + \sigma_\nu^{-1} M. \quad (24)$$

After computing the gradient $\nabla \varphi_\sigma$ at an iterate \bar{x} (which requires the sorting of $\{A_{j,:}^\top \bar{x} + b_j\}_{j=1}^m$), one can efficiently compute the generalized Jacobian of $\nabla \varphi_\sigma$ based on the already sorted values of $\{A_{j,:}^\top \bar{x} + b_j\}_{j=1}^m$, as specified in Lemma 2. Below we will show that the extra cost for this computation of the generalized Jacobian is actually very low even if m is large.

Let $\bar{\alpha}, \bar{\beta}, \bar{\gamma}$ be the associated index-sets of $\bar{y} := \text{proj}_{\mathcal{B}_k}(A\bar{x} + b - \sigma\bar{\lambda})$ defined in (4) for $k = (1 - \tau)m$. When \bar{x} and $\bar{\lambda}$ are such that the projection onto \mathcal{B}_k is differentiable at $A\bar{x} + b - \sigma\bar{\lambda}$, we use the Jacobian formula given in Lemma 2. Otherwise, a sequence of differentiable points converging to $A\bar{x} + b - \sigma\bar{\lambda}$ can be constructed, and we use the limit of the gradients at this sequence based on whether $\bar{k}_1 = k$ (Case 3) or not (Case 2) from Lemma 2.

In Case 1 of Lemma 2 (when $\mathsf{T}_k(A\bar{x} + b - \sigma^{-1}\bar{\lambda}) < 0$), then the matrix in linear system (22) has no contribution from the metric projection and can be constructed easily.

In Case 2, we have $I - J_{\mathcal{B}_k} = \begin{bmatrix} Q & \\ & 0_{\bar{\gamma}} \end{bmatrix}$ where $Q = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} \in \mathbb{R}^{\bar{k}_1 \times \bar{k}_1}$ is a symmetric projection matrix that is computable in closed form as a function of the index-sets. In particular, for $\bar{\rho} := k^2 - 2k\bar{k}_0 + \bar{k}_0\bar{k}_1$, direct computation shows that

$$\begin{aligned} Q_{11} &= \bar{\rho}^{-1}(\bar{k}_1 - \bar{k}_0)\mathbf{1}_{\bar{\alpha}}\mathbf{1}_{\bar{\alpha}}^\top, & Q_{12} &= \bar{\rho}^{-1}(k - \bar{k}_0)\mathbf{1}_{\bar{\alpha}}\mathbf{1}_{\bar{\beta}}^\top, \\ Q_{21} &= Q_{12}^\top, & Q_{22} &= I - \bar{\rho}^{-1}\bar{k}_0\mathbf{1}_{\bar{\beta}}\mathbf{1}_{\bar{\beta}}^\top. \end{aligned}$$

To ease notation, let $\bar{\kappa} := (\bar{\alpha}, \bar{\beta})$ with $|\bar{\kappa}| = \bar{k}_1$. Since Q is a symmetric projection matrix, we have

$$\sigma A^\top (I - J_{\mathcal{B}_k}) A = \sigma \cdot (A_{\bar{\kappa},:})^\top Q A_{\bar{\kappa},:} = \sigma \cdot (A_{\bar{\kappa},:})^\top Q^\top Q A_{\bar{\kappa},:} = \sigma T^\top T, \quad T := Q A_{\bar{\kappa},:}. \quad (25)$$

Observe that $T \in \mathbb{R}^{|\kappa| \times n}$ can be computed as follows:

$$\begin{aligned} T_{i,:} &= \bar{\rho}^{-1}(\bar{k}_1 - \bar{k}_0)\mathbf{1}_{\pi(\bar{\alpha})}^\top A_{\pi(\bar{\alpha}),:} + \bar{\rho}^{-1}(k - \bar{k}_0)\mathbf{1}_{\pi(\bar{\beta})}^\top A_{\pi(\bar{\beta}),:} & i \in \bar{\alpha}, \\ T_{i,:} &= A_{\pi(i),:} + \bar{\rho}^{-1}(k - \bar{k}_0)\mathbf{1}_{\pi(\bar{\alpha})}^\top A_{\pi(\bar{\alpha}),:} - \bar{\rho}^{-1}\bar{k}_0\mathbf{1}_{\pi(\bar{\beta})}^\top A_{\pi(\bar{\beta}),:} & i \in \bar{\beta}, \end{aligned}$$

where $\pi(\mathcal{I})$ denotes the \mathcal{I} components of π . Since the rows of T corresponding to $\bar{\alpha}$ are identical, we may construct $T^\top T = \tilde{T}^\top \tilde{T}$ for some \tilde{T} in a reduced space where we define, for any $i \in \bar{\alpha}$,

$$\tilde{T} := \begin{bmatrix} \sqrt{|\bar{\alpha}|} T_{i,:} \\ T_{\bar{\beta},:} \end{bmatrix} \in \mathbb{R}^{(|\bar{\beta}|+1) \times n}. \quad (26)$$

We further illustrate the special structured pattern of the generalized Jacobian for this case in Fig. 1. This depiction highlights how the Newton matrix incorporates only a subset of the scenarios, resulting from the form of the non-positive top- k -sum projection in Section 4.1. Section 4.1 demonstrates that the projection matrix onto \mathcal{B}_k is composed of four components: (i) the $(\bar{\alpha}, \bar{\alpha})$ block, which averages the scenarios within the index set $\bar{\alpha}$; (ii) the $(\bar{\alpha}, \bar{\beta})$ block and its transpose, both of which average the scenarios between $\bar{\alpha}$ and $\bar{\beta}$; (iii) the $(\bar{\beta}, \bar{\beta})$ block, which adjusts the $\bar{\beta}$ scenarios; and (iv) the $(\bar{\gamma}, \bar{\gamma})$ block, which excludes scenarios in $\bar{\gamma}$.

In Case 3 of Lemma 2, we have $Q = \begin{bmatrix} (|\bar{\alpha}| + |\bar{\beta}|)^{-1} \mathbf{1}_{\bar{\kappa}} \mathbf{1}_{\bar{\kappa}}^\top & \\ & I_{|\bar{\gamma}|} \end{bmatrix}$. Hence, $A^\top (I - J_{\mathcal{B}_k}) A = A^\top Q A = \tilde{T} \tilde{T}^\top$ where $\tilde{T} = \mathbf{1}_{\bar{\kappa}}^\top A_{\bar{\kappa},:} / \sqrt{|\bar{\alpha}| + |\bar{\beta}|}$, i.e., the scaled average of the $\bar{\alpha}$ and $\bar{\beta}$ rows of A .

Note that the Sherman-Morrison-Woodbury (SMW) identity can be used to reduce the computational cost of solving the linear system (22) as follows:

$$V^{-1} = (D + \tilde{T}^\top (\sigma I) \tilde{T})^{-1} = D^{-1} - D^{-1} \tilde{T}^\top (\sigma^{-1} I + \tilde{T} D^{-1} \tilde{T}^\top)^{-1} \tilde{T} D^{-1}, \quad (27)$$

where $D := \nabla^2 f(\bar{x}) + \sigma^{-1} M + \sigma(I - J_X)$ is assumed to be invertible.

Table 2 summarizes the extra computational cost of obtaining first-vs-zeroth-order and second-vs-first-order information in our method. After computing the value $\varphi_\sigma(\bar{x})$, the incremental cost of the gradient computation is due to a matrix-vector product of the form $A^\top v$, where $v = \bar{y} - \text{proj}_{\mathcal{B}_k}(\bar{y})$




Figure 1: (a) Schematic of the structured sparsity (with affine $G(x) := Ax + b$ and a single superquantile constraint) for constructing the generalized Jacobian at a sorted point $\tilde{y}^0 \in \mathbb{R}^m$ based on (b) sorted projection $\bar{y} \in \mathcal{B}_k$ with index sets $\bar{\alpha}, \bar{\beta}, \bar{\gamma}$. Only the rows $A_{j,:}$ for $j \in \bar{\alpha} \cup \bar{\beta}$ are relevant in the semismooth Newton equation. See (23) for the case of the nonlinear function G .

Terms	Computational cost
Objective function value (18)	$O(m \log(m) + mn)$
Gradient (19)	$O(n \cdot \bar{k}_1)$ (after computing the objective function (18))
Generalized Jacobian matrix (24)	$O(\overbrace{n \cdot \bar{k}_1}^{\tilde{T}} + \min\{\overbrace{n \bar{\beta} ^2 + n^3}^{\text{direct}}, \overbrace{n \bar{\beta} (n + \bar{\beta}) + \bar{\beta} ^3}^{\text{SMW}}\})$
and semismooth Newton direction	(after computing the gradient (19))

Table 2: Computational cost for ALM subproblem (14) when the objective is linear or diagonal convex quadratic. The regularization matrix M is assumed to be diagonal. The cost for the generalized Jacobian is according to the most expensive situation in Lemma 2 (Case 2). Note that calculating \tilde{T} can be faster than generic $n \bar{k}_1$ multiplication, since the former involves summing contiguous ranges that can exploit SIMD.

is a sparse vector with \bar{k}_1 nonzero elements, as can be seen from Fig. 1. After evaluating $\nabla\varphi$, the extra cost to compute its generalized Jacobian is contributed by two steps: (i) the computation of the matrix for the semismooth Newton equation, which involves summing across $A_{j,:}$ for $j \in \bar{\alpha} \cup \bar{\beta}$ and can be performed more efficiently than a sparse matrix-vector product of the same theoretical

cost in the gradient computation; and (ii) solving the linear equation, which can be negligible in the large- m case. The first line in Table 2 indicates that the function evaluation depends explicitly on the number of scenarios, m , via the sorting, linear transformation, and projection procedures, whereas the cost of computing the gradient, generalized Jacobian, and solving the linear equations is not directly dependent on m .

When $|\bar{\beta}| \ll m$ at the current iterate $(\bar{x}, \bar{\lambda})$, Table 2 shows that obtaining the generalized Jacobian of $\nabla \varphi_\sigma(\bar{x})$ and inverting the matrix incurs little extra cost after knowing its function value and the gradient, indicating the scalability of Algorithm 1 with respect to the number of scenarios. In particular, due to the low rank structure of T and additional simplifications arising when f is linear or separable quadratic, the Newton system in (22) may be computed in n -dimensions or $|\bar{\beta}| + 1$ -dimensions due to the Sherman-Morrison-Woodbury identity, where a scalar parameter $\theta > 0$ may be used to tune this trade-off. Numerical experiments on synthetic data support the conclusion that in certain instances, the incremental cost of obtaining the Newton direction may be smaller than the incremental cost of evaluating the gradient; see Table 4 for detailed timing results. Note that $JG^\top Q JG$ still may be computed with the reduced dimension $|\bar{\beta}|$ when G is nonlinear.

4.2 Implicit vs explicit scenario reduction

As discussed in the last section, our semismooth-Newton-based ALM can be viewed as an implicit scenario reduction approach. We will now compare this approach to an explicit scenario reduction strategy used in the outer-approximation (OA) scheme in Royset and Wets (2021, Chapter 6.C). This comparison highlights the potential advantages of our approach.

The OA method leverages the intuition that many $u_{\ell j} = 0$ at the solution to the reformulation (3) when τ_ℓ is close to 1, and in its simplest form, proceeds by maintaining active sets \mathcal{A}_ℓ^ν of candidate effective scenarios with $\mathcal{A}_\ell^\nu = \{1, \dots, m\} \setminus \{j : u_{\ell j} = 0\}$ with the property that $\mathcal{A}_\ell^\nu \subseteq \mathcal{A}_\ell^{\nu+1}$ for iteration ν . The set for the next iteration is obtained by $\mathcal{A}_\ell^{\nu+1} = \mathcal{A}_\ell^\nu \cup J_\ell^\nu$ where J_ℓ^ν contains the indices of a fraction of the $(1 - \tau_\ell)m$ -largest elements of the values $\{u_{\ell j}^\nu - g^\ell(x^\nu; \omega^j) + t_\ell^\nu\}_{j=1}^m$. The method terminates in iteration $\nu+1$ when $\mathcal{A}_\ell^\nu = \mathcal{A}_\ell^{\nu+1}$. Achieving optimal performance with the OA method necessitates fine-tuning the update strategy for \mathcal{A}_ℓ^ν and often involves solving subproblems ν to progressively tighter tolerances. In fact, the problem of adding an appropriate amount of additional scenarios is challenging even when the candidate solution is near the true solution set, as there is a trade-off between number of subproblems solved and the difficulty of solving each subproblem. In addition, the set of effective scenarios at the solution may be larger than $(1 - \tau_\ell)m$, but this number is not known *a priori*.

In contrast, our ALM implicitly maintains an approximation to the optimal set of effective scenarios through the partition set $\bar{\alpha} \cup \bar{\beta}$, and it uses an inherent update rule for \mathcal{A}_ℓ^ν that is guided by the generalized Jacobian of $\text{proj}_{\bar{g}}$. Consequently, the ALM effectively prunes and approximates this set in each iteration, refining the pool of scenarios considered.

4.3 Comparison to existing semismooth-Newton ALMs for structured sparsity

A number of existing works use the semismooth-Newton based ALM to introduce structured sparsity (Li et al., 2018a,b; Li and Li, 2021; Wu et al., 2023) in the corresponding linear system to compute the semismooth Newton direction. Among them, Wu et al. (2023) is most closely related to our superquantile problem (1), but a major distinction lies in how the algorithm leverages the sparsity. The optimization model in the latter paper explicitly added an ℓ_1 regularizer to the objective function to promote solution sparsity. In the case when $G^0(x; \omega^{[m]}) := \{|A_{i,:}^\top x + b_i|\}_{i=1}^m$ as in Wu et al. (2023), the ALM subproblems for the dual formulation require the linear solution of a system in the form of $APA^\top + Q$, where P is a sparse diagonal matrix based on the generalized Jacobian of the proximal mapping of the ℓ_1 regularizer, and Q corresponds to the generalized Jacobian of the proximal operator for the superquantile. However, when the sparse regularizer is not present, the system instead requires the solution of a linear system involving a matrix of the form $AA^\top + Q$, which can be significantly more costly.

In contrast, the corresponding linear systems in our formulation take the form $A^\top QA + \nabla^2 f(x)$ with $g = \delta_{\mathcal{B}_{(k)}}^\circ$. In this way, our framework can employ careful numerical linear algebra to reduce the computational cost of the linear solve without the use of a sparse decision-space regularizer in the objective. Instead, we interpret the matrix Q as promoting a form of constraint sparsity, i.e., *implicit* scenario reduction. This is made clear by the form of Q , which subsets certain rows (scenarios) of the matrix A , leading to a principled approach for adjusting the active $(1 - \tau_0)$ -tail scenarios rather than an ad hoc, manually-tuned procedure.

Since our technique scales in the number of scenarios, we expect our method to be complementary to the decision-space scalability developed in Wu et al. (2023).

5. Experiments

To study the performance of our proposed framework, we conduct two simulation studies: the first uses synthetic data (Section 5.2), and the second demonstrates its practical application through quantile regression on a realistic dataset containing more than 30 million records (Section 5.3). The implementation in Julia is available at <https://github.com/jacob-roth/superquantile-opt>.

We compare the performance of our semismooth-Newton-based ALM implemented in Julia v1.9.3 (Bezanson et al., 2017) against several other approaches: (i) The interior point methods for solving the nonlinear reformulation (3) by Gurobi v11.0 (Gurobi Optimization, LLC, 2023); (ii) The alternating direction method of multipliers implemented by OSQP v0.6.3 (Stellato et al., 2020); (iii) the Portfolio Safeguard v3.2 (American Optimal Decisions, Inc., 2011); and (iv) an implementation of the outer approximation algorithm based on Royset and Wets (2021, Chapter 6.C) that calls Gurobi as a subproblem solver. For simplicity, we refer to these methods as (a) GRB, (b) OSQP, (c) PSG, and (d) G-OA respectively in the later discussion.

5.1 Implementation details

Before showing the numerical results, we first provide some implementation details of our approach. At a candidate primal-dual tuple (x, y, z, λ, μ) for problem (13) and a given tolerance ϵ , we stop the algorithm if $\eta := \max\{\eta_p, \eta_d, \eta_r\} \leq \epsilon$, where

$$\begin{aligned}\eta_p &:= \max\left\{\left\|\max(Ax + b - y, 0)\right\|/(1 + \|b\|), \left\|y - \text{proj}_{\bar{B}}(y)\right\|/(1 + \|y\|), \right. \\ &\quad \left. \|x - z\|/(1 + \|z\|), \left\|z - \text{proj}_X(z)\right\|/(1 + \|z\|)\right\}, \\ \eta_d &:= \|A^\top \lambda + \mu + \nabla f(x)\|/(1 + \|\nabla f(x)\|), \\ \eta_r &:= |\text{obj}_p - \text{obj}_d|/(1 + |\text{obj}_p|), \\ &\quad \text{with } \text{obj}_p := f(x) \text{ and } \text{obj}_d := -f^*(-A^\top \lambda - \mu) - \delta_{\bar{B}}^*(-\lambda) - \delta_X^*(-\mu) + b^\top \lambda.\end{aligned}$$

The above conditions ensure the KKT conditions (15) hold when $\epsilon = 0$. For the ALM subproblems (14), we use the following termination criteria:

$$\|\nabla \varphi_{\sigma_\nu}(x^{\nu+1})\|_2 \leq \frac{\epsilon_\nu}{\sigma_\nu}, \quad \epsilon_\nu \geq 0, \quad \sum_{\nu=0}^{\infty} \epsilon_\nu < +\infty \quad (29)$$

$$\|\nabla \varphi_{\sigma_\nu}(x^{\nu+1})\|_2 \leq \frac{\delta_\nu}{\sigma_\nu} \|(x^{\nu+1}, y^{\nu+1}, z^{\nu+1}) - (x^\nu, y^\nu, z^\nu)\|_M, \quad 0 \leq \delta_\nu < 1, \quad \sum_{\nu=0}^{\infty} \delta_\nu < +\infty, \quad (30)$$

where M is the positive semidefinite matrix chosen at iteration ν in Algorithm 1 and φ_{σ_ν} is defined in (18). These conditions are given in Li et al.; Wu et al. (2023) and are based on the guidelines developed by Rockafellar (1976a) for the proximal point algorithm.

Note that if the subproblem is solved exactly with $M \equiv 0$, then the resulting iterate is dual feasible, i.e., satisfies $\nabla f(x) + A^\top \lambda + \mu = 0$. This holds because the ALM updates give

$$\begin{aligned}0 &= \nabla f(x^{\nu+1}) + \sigma_\nu A^\top (Ax^{\nu+1} + b + \sigma_\nu^{-1} \lambda^\nu - y^{\nu+1}) + \sigma_\nu (x^{\nu+1} + \sigma_\nu^{-1} \mu^\nu - z^{\nu+1}), \\ \lambda^{\nu+1} &= \lambda^\nu - \sigma_\nu (y^{\nu+1} - Ax^{\nu+1} - b), \\ \mu^{\nu+1} &= \mu^\nu - \sigma_\nu (z^{\nu+1} - x^{\nu+1}),\end{aligned}$$

which in turn gives

$$\nabla f(x^{\nu+1}) = -\sigma_\nu A^\top (Ax^{\nu+1} + b + \sigma_\nu^{-1} \lambda^\nu - y^{\nu+1}) - \sigma_\nu (x^{\nu+1} + \sigma_\nu^{-1} \mu^\nu - z^{\nu+1}) = -A^\top \lambda^{\nu+1} - \mu^{\nu+1}.$$

Next, an important aspect of the numerical implementation involves mitigating the computational burden associated with sorting operations, as a sorting step is executed in each line search iteration. Although the optimal index \bar{k}_1 cannot be determined *a priori*, an existing permutation may already be optimal, based on the observation that the projection onto the top- k -sum constraint set only requires the elements in the set $\bar{\alpha} \cup \bar{\beta}$ to be in sorted order. At the ν -th iteration of the ALM, within the line search procedure, we find that the candidate solution's current $\bar{k}_1^{(\nu)}$ augmented by a small buffer provides an effective estimate for the next iterate's index $\tilde{k}_1^{(\nu+1)}$. We leverage this observation to perform a partial quick-sort of the top $\tilde{k}_1^{(\nu+1)}$ elements and thereafter compute the projection

based on this partial permutation. A projection \tilde{z} computed from the partial permutation of an initial point $z^0 \in \mathbb{R}^n$ can be certified as correct by verifying that $z_i^0 > \tilde{z}_{\tilde{k}_1}$ for $i \in \{\tilde{k}_1, \tilde{k}_1 + 1, \dots, n\}$. Effectively, this checks whether or not the projection \tilde{z} is sorted and obeys the KKT conditions $\tilde{z}_{\tilde{k}_1+1} < \tilde{z}_{\tilde{k}_1}$ and $\tilde{z}_i \geq \tilde{z}_{i+1}$ for $i \in \{\tilde{k}_1, \tilde{k}_1 + 1, \dots, n\}$. We repeatedly solve the projection problem associated with progressively larger partial-sorts until the above conditions are satisfied. In practice, we find that re-sorts are required infrequently (approximately 5% of iterations necessitate this adjustment) and can be efficiently warm-started.

Through experimentation, we find that the OA method performs relatively well by targeting ≈ 10 total subproblems and solving subproblem ν to progressively tighter tolerances when using Gurobi's barrier solver. Gurobi has no warm-start feature for the barrier method, but we find that performance did not improve using more frequent warm-starts with primal/dual simplex methods. In tuning the $L = 1$, $k = 1\%m$ case, we set J^ν in iteration ν to be the largest $3k$ elements of $(Ax + b) - \mathbb{1}t - u$ in an effort to perform an aggressive scan of the risk space (where $u_j = 0$ for $j \notin \mathcal{A}^\nu$). Note that the size of $\mathcal{A}^{\nu+1}$ may increase modestly relative to $3k$ if many active elements in iteration ν remain active in $\nu + 1$, which is often the case in later iterations.

5.2 Results on synthetic data

In this section, we generate synthetic data for the superquantile-constrained problem (1). Problems instances are randomly generated according to the following procedure, which ensures that a solution exists. The box constraint is $X := \{x \in \mathbb{R}^n : -\mathbb{1} \leq x \leq \mathbb{1}\}$. For each $\ell \in \{1, \dots, L\}$, constraint data $A_{:,j}^\ell \sim \text{Normal}(0, 10^2)$ with unit column infinity-norm (i.e., $A_{:,j}^\ell$ is normalized by $\|A_{:,j}^\ell\|_\infty$), and $\tilde{b}^\ell \sim \text{Normal}(0, 1)$. To ensure that the problem feasible, we define b^ℓ to be a shift of \tilde{b}^ℓ via the following procedure: for $p \in \{1, \dots, \lceil \log(N) \rceil\}$, generate feasible points $\tilde{x}^p \in X$ uniformly by $\tilde{x}^p := -\mathbb{1} + 2u$ for $u \sim \text{Uniform}(0, 1)^n$. Subsequently, for each $\ell \in \{1, \dots, L\}$, define $\tilde{y}^{\ell,p} := A^\ell \tilde{x}^p + \tilde{b}^\ell$ and set $b^\ell := \tilde{b}^\ell - \min_p (\mathbf{T}_{k_\ell}(\tilde{y}^{\ell,p}))$ to generate difficult instances. We consider two types of objective functions: (i) diagonal (strongly) convex quadratic function $f(x) = \frac{1}{2}x^\top Cx + c^\top x$ where $C_{ii} \sim |\text{Normal}(0, 1)|$ and $c_i \sim \text{Normal}(0, 1)$; (ii) linear function $f(x) = c^\top x$, where c is generated in the same way as (i). The ALM is initialized at the feasible point $x^0 = 0$, $\mu = 0$, and $\lambda = 0$, and we find that the method performs similarly for initial points with elements drawn uniformly within the bounds of X .

The experiments are conducted on a Dell workstation running a Windows environment, equipped with 8 physical processors and 32GB of RAM. We compare our ALM with GRB, G-OA, and PSG for computing highly accurate solutions (where $\epsilon = 10^{-8}$), and with OSQP for computing solutions with lower accuracy (where $\epsilon = 10^{-3}$). While both ALM and Gurobi barrier methods are able to use multiple threads (via BLAS/LAPACK routines and parallelized matrix factorization, respectively), establishing a fair comparison is somewhat ambiguous since Gurobi's procedure is not well documented. Since OSQP's default solver is single-threaded, we limit GRB, G-OA, and ALM to one execution thread in our experiments. PSG's multithreading capabilities are not documented, and we use the default settings. In addition, we turn off Gurobi's presolve option and OSQP's polishing procedure, as our findings indicate that these features do not enhance performance.

In scenarios where $m \gg n$, we find that Gurobi's barrier method is more efficient than either primal simplex or dual simplex; when $m \approx n$, the barrier method is comparable to the dual simplex for less

stringent risk tolerance $(1 - k(\tau)/m)$ and may be slower for more stringent risk tolerance, but both are more efficient than primal simplex; when $m \ll n$, the dual simplex is most efficient and may significantly outperform the barrier method, as well as ALM. Since the focus of the comparison is on instances where $m \gg n$, we report solve times using the minimum of Gurobi’s barrier and dual-simplex methods in a grid of instances with $m \geq n$ when reporting GRB solve statistics. We set the feasibility and optimality tolerance to 10^{-8} . To obtain OSQP solutions with $\epsilon = 10^{-3}$ precision, we set the relative and absolute tolerances to 10^{-5} and 10^{-3} , respectively. In the PSG suite, we set the precision to 10^{-8} . For PSG, we use the HELI solver, which uses Gurobi as a subproblem solver, for both linear and quadratic problems. The HELI solver managed to solve instances to the specified tolerance more reliably than the PSG suggested solvers,³ CAR and VAN, as well as their Gurobi-based variants CARGRB and VANGRB.

Finally, problems are generated for a variety of parameters: $m \in \{2^{14}/L, 2^{15}/L, \dots, 2^{20}/L\}$, $n \in \{2^7, 2^8, \dots, 2^{13}\}$, $L \in \{1, 10\}$, and $k = \{0.01, 0.1\} \times m$. However, some combinations of m , n , and L exceed the 32GB RAM capacity, so we exclude them in the experiments. In addition, PSG’s Gurobi-based solvers are constrained by the number of decision variables they can handle. Consequently, we only select instances that are solvable by PSG.

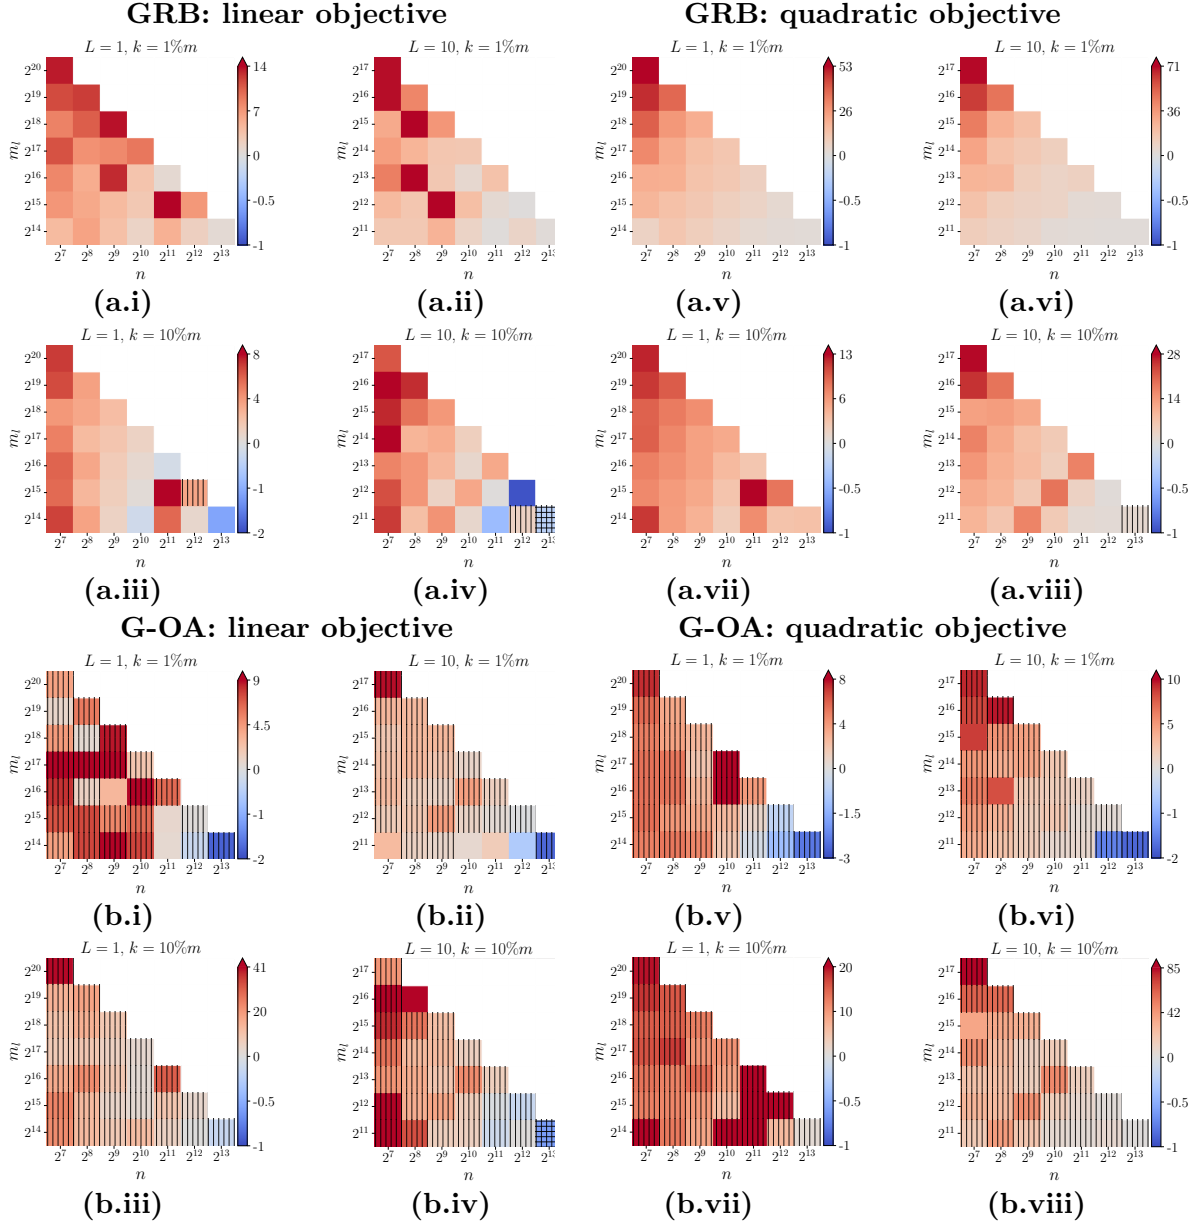
Fig. 2, which splits over the next couple of pages, shows the relative solve time of method i versus ALM with $i \in \{a = \text{GRB}, b = \text{G-OA}, c = \text{PSG}, d = \text{OSQP}\}$ for a range of problems across parameters m , n , L , and k for both linear and diagonal quadratic objectives.

The first row of Fig. 2a shows that across linear and quadratic objectives, ALM was never slower than GRB for instances with high risk aversion ($k = 1\%m$). We highlight that ALM’s relative performance generally improves both as m increases within each subplot and as L increases (between consecutive subplots). In the largest- m instances, ALM demonstrates improvement over GRB of at least 20x and up to 70x when both trends are maximized. Similar trends are observed in the second row of Fig. 2a, corresponding to lower risk aversion instances (larger k). In these instances, ALM’s relative performance degrades due to the increased cost of constructing \tilde{T} associated with larger k but maintains an advantage of at least 5x and up to 30x on large- m problems. For relatively small- m and large- n problems, GRB’s dual simplex can be very efficient, beating ALM by up to 4x on problems with a linear objective, but losing its advantage for problems with quadratic objective.

In Fig. 2b, G-OA obtained solutions satisfying primal and dual feasibility in each instance, but objective gap tolerances were violated by a factor of ≈ 10 , contributing to the vertical shading. The top row of Fig. 2b shows that ALM retains a solve-time advantage of about 5-10x over G-OA in the large- m cases for the high risk-aversion instances. The bottom row shows that ALM retains a more significant advantage in lower risk-aversion instances. In contrast to GRB, this difference is due to the fact that G-OA must scan through a larger set of effective scenarios, degrading its performance as the size of the OA subproblem grows. This issue could perhaps be overcome by using adaptive scenario pruning strategies, which indicates the importance of adaptive scenario selection in ALM for problems with a larger effective scenario tail.

Fig. 2c shows that in some instances, PSG can obtain solutions to the large- m cases very fast, almost matching ALM on problems with linear objective (left two columns). However, PSG failed

3. In fact, the quadratic linear (quadratic) solvers suggested in the PSG documentation failed to find solutions with the desired tolerance in 75% (96%) of the instances tested below.



to obtain solutions with satisfactory primal feasibility in many instances. Finally, Fig. 2d shows that ALM ($\epsilon = 10^{-3}$) was at least 50x faster, and sometimes more than 750x faster, than OSQP ($\epsilon = 10^{-3}$) in large- m instances.

Table 3 provides additional detail on the comparison of ALM against other methods across the three instances with the largest m . It shows that PSG can outperform GRB and GRB-OA in certain instances, but PSG can also fail and return solutions with significant primal feasibility violation (*e.g.*, the third subproblem in each group). In addition, the table indicates that G-OA finds feasible solutions that are nearly optimal.

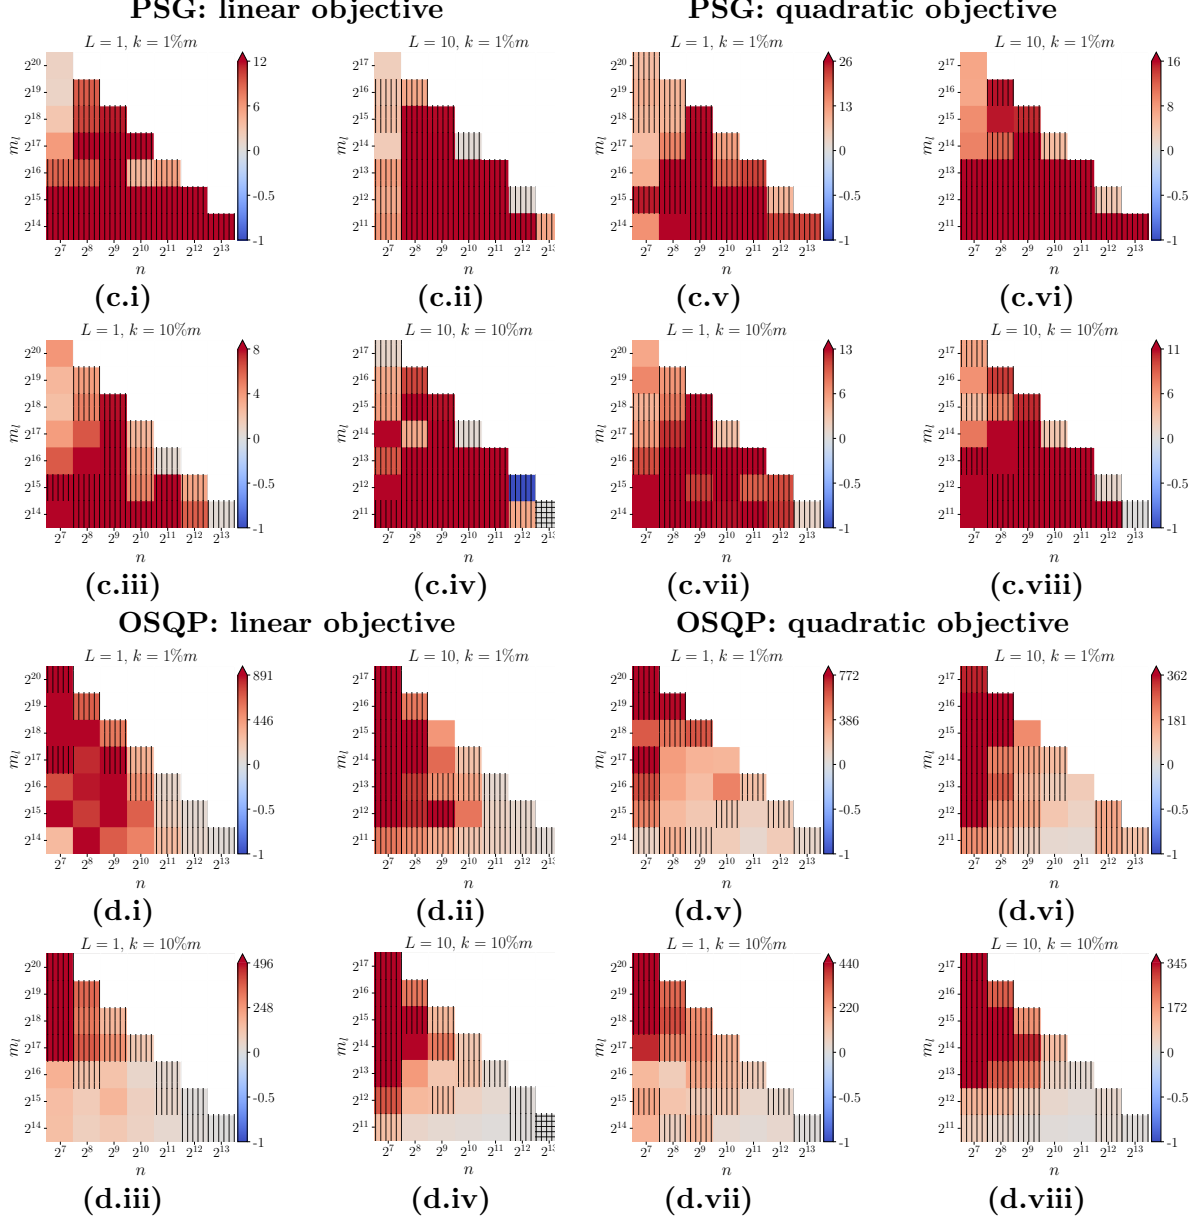


Figure 2: Heatmap of relative solve time (method i - ALM)/ALM, for $i \in \{a = \text{GRB}, b = \text{G-OA}, c = \text{PSG}, d = \text{OSQP}\}$. Subfigures (i)–(iv) ((v)–(vii)) correspond to instances with linear (separable quadratic) objective. Red shading (positive) indicates that ALM was faster than method i ; blue shading (negative) indicates that method i was faster than ALM. Shading intensity represents relative solve time. Vertical (horizontal) hatching indicates that method i (ALM) failed to obtain a satisfactory solution within 1hr. For convenience: $2^{20} = 1.0 \times 10^6$, $2^{17} = 1.3 \times 10^5$, and $2^{13} = 0.8 \times 10^4$.

Detailed timing statistics for ALM on the ten largest- m instances (top left region of subplots in Fig. 2) are summarized in Table 4. Notably, the incremental cost of obtaining second-order information (after computing $\nabla \varphi_\sigma$) can be 3-7x *cheaper* than the incremental cost of obtaining first-

order information (after computing φ_σ) when $m \gg n$ and $k \ll m$. However, this cost advantage lessens as the ratio m/n decreases.

In addition, comparison of the three timing columns indicates several key insights. First, solutions with the accuracy around $\eta = 10^{-3}$ can be obtained in approximately the same amount of time for both linear and quadratic objectives. Second, ALM is highly effective at refining a solution from moderate precision ($\eta = 10^{-6}$) to high precision ($\eta = 10^{-8}$). This stands in contrast to barrier methods, which often encounter difficulties in the later stages of solving large-scale instances due to poor numerical conditioning. The efficiency of ALM can be attributed to the small cardinality of $\tilde{\beta}$ near a solution, as noted in the last column, which enhances performance particularly when the set of effective variables remains consistent, thereby minimizing the need for frequent sorting. The penultimate column highlights that the matrix $\tilde{T}^\top \tilde{T}$ is constructed in a space considerably smaller than m , and the Newton system can be solved in such a small dimensional space.

5.3 A quantile regression experiment on real data

In this section, we consider an application of the superquantile minimization problem to estimating the τ -conditional quantile under a linear model (Royset and Wets, 2021, Chapter 8.C). Given input-output data $\{(a^i, b_i)\}_{i=1}^m \subseteq \mathbb{R}^n \times \mathbb{R}$ and a scalar $\tau \in (0, 1)$ such that $k(\tau) = (1 - \tau) \cdot m$ is an integer, the problem takes the following form:

$$\begin{aligned} & \underset{x, t}{\text{minimize}} && t + x^\top \left(\frac{1}{m} \sum_{i=1}^m a^i \right) \\ & \text{subject to} && \mathsf{T}_{k(\tau)}(\{b_i - x^\top a^i - t\}_{i=1}^m) \leq 0. \end{aligned} \tag{31}$$

See Appendix C for the detailed derivation of the above formulation.

Using ride-share data from the NYC Taxi and Limousine Commission (2023) and weather data from OpenWeather (2023) for October 2022 through September 2023, we construct a dataset containing $m = 32,534,601$ trips with $n = 19$ features to predict the tip received by the driver for rides paying with credit card. The full dataset $\{(a^i, b_i)\}_{i=1}^m$ is available at <https://doi.org/10.5281/zenodo.11183368>.

Based on the numerical results in Section 5.2, we have narrowed the focus of our study to the most efficient methods: ALM, GRB, and G-OA for the superquantile formulation (31). Additionally, we use Gurobi's barrier method to solve the following linear reformulation of quantile regression (31):

$$\begin{aligned} & \underset{x, (z^+, z^-) \in \mathbb{R}^{2m}}{\text{minimize}} && \frac{1}{m} \sum_{i=1}^m (1 - \tau) \cdot z_i^- + \tau \cdot z_i^+ \\ & \text{subject to} && z_i^+ - z_i^- = b_i - x^\top a^i, \quad z_i^+ \geq 0, \quad z_i^- \geq 0, \quad i \in \{1, \dots, m\}, \end{aligned} \tag{32}$$

which we refer to as G-QR. For all of the above mentioned methods, we set a time limit of 3,600s.

We conduct two experiments. For the first experiment, we predict the τ -th top-quantile for $\tau \in \{1\%, 10\%, 50\%\}$ using all methods at a precision of $\epsilon = 10^{-4}$. Notably, $\tau = 50\%$ presents the most computational challenge for ALM since for $\tau > 0.5$, the superquantile formulation (31) can be inverted to $\bar{\tau} = 1 - \tau < 0.5$ by reversing the sign of $\{b_i\}_{i=1}^m$ and $\{a^i\}_{i=1}^m$. Secondly, we employ ALM

(a) Linear objective

Problem (m, n)	Relative Pobj				Pinfeas				Dinfeas				Relgap				Walltime [s]					
	GRB	G-OA	PSG	OSQP	ALM	GRB	G-OA	PSG	OSQP	ALM	GRB	G-OA	OSQP	ALM	GRB	G-OA	OSQP	ALM	GRB	G-OA	PSG	OSQP
$L = 1, k = 1\%m$																						
($2^{20}, 2^7$)	-2.6e-11	3.6e-8	5.3e-9	—	3.3e-9	5.2e-12	0.0e0	7.0e-9	—	3.3e-11	2.8e-14	1.3e-14	—	7.3e-11	-1.9e-12	2.2e-6	—	2.6e1	3.8e2	1.3e2	6.9e1	3.6e3
($2^{17}, 2^7$)	4.1e-10	-2.5e-6	4.0e-8	-7.8e-4	2.3e-9	0.0e0	7.2e-6	7.4e-9	2.3e-3	8.2e-10	1.1e-14	1.4e-14	1.9e-4	1.8e-10	1.8e-9	5.4e-7	-7.5e-4	1.6e0	1.9e1	2.8e1	1.1e1	8.8e2
($2^{17}, 2^{10}$)	8.6e-10	-2.7e-6	2.3e-5	-4.0e-4	9.4e-9	0.0e0	3.5e-5	4.4e-8	5.4e-3	1.1e-10	5.4e-11	1.3e-10	2.3e-4	-4.3e-10	7.5e-10	-1.6e-6	-3.9e-4	2.3e1	2.3e2	7.0e1	6.9e2	3.6e3
$L = 1, k = 10\%m$																						
($2^{20}, 2^7$)	6.0e-9	-4.6e-6	1.1e-9	—	7.5e-9	0.0e0	8.0e-6	7.8e-9	—	2.1e-10	2.9e-14	1.4e-14	—	-1.3e-9	5.2e-9	-3.8e-6	—	2.7e1	2.2e2	1.1e3	1.4e2	3.6e3
($2^{17}, 2^7$)	6.4e-9	-6.7e-5	1.4e-8	-3.2e-3	6.7e-11	0.0e0	1.3e-4	5.0e-9	6.0e-3	7.5e-11	7.4e-15	4.5e-15	7.5e-5	1.5e-12	6.3e-12	-6.4e-5	-3.1e-3	2.9e0	1.7e1	2.5e1	1.4e1	5.5e2
($2^{17}, 2^{10}$)	7.7e-10	-9.3e-8	-2.8e-1	1.0e-2	9.6e-9	3.3e-12	7.8e-7	2.8e0	4.6e-3	7.0e-11	2.0e-12	3.2e-11	1.9e-4	2.8e-11	1.4e-9	6.7e-8	1.0e-2	5.6e1	1.1e2	2.3e2	2.0e2	3.6e3
$L = 10, k = 1\%m$																						
($2^{17}, 2^7$)	5.5e-11	-4.6e-8	9.4e-9	—	2.0e-9	1.4e-11	1.4e-7	8.6e-10	—	3.9e-10	1.0e-12	1.8e-13	—	3.9e-10	5.3e-10	-4.4e-8	—	2.5e1	6.7e2	7.0e2	1.0e2	3.6e3
($2^{14}, 2^7$)	3.8e-9	-2.3e-6	1.3e-7	-2.5e-4	4.3e-7	0.0e0	8.6e-6	1.7e-14	2.3e-3	8.7e-12	2.2e-12	1.8e-13	2.8e-4	-1.2e-9	2.2e-9	-2.3e-6	-2.8e-4	4.6e0	6.1e1	3.7e1	2.0e1	9.9e2
($2^{14}, 2^{10}$)	-7.8e-11	3.3e-8	8.0e-3	-5.8e-5	4.5e-9	0.0e0	2.0e-7	8.7e-1	3.0e-3	1.7e-11	1.4e-12	4.0e-12	6.9e-4	1.6e-10	1.7e-10	5.3e-8	-5.0e-5	4.2e1	2.8e2	1.6e2	7.6e1	3.6e3
$L = 10, k = 10\%m$																						
($2^{17}, 2^7$)	5.0e-10	-1.5e-7	1.1e-3	-9.1e-3	6.1e-9	1.5e-9	3.4e-7	6.9e-9	2.2e-2	3.7e-11	7.9e-14	2.7e-14	5.8e-4	-3.8e-10	2.7e-9	-1.2e-7	-3.0e-3	6.0e1	4.9e2	1.0e3	9.5e1	3.6e3
($2^{14}, 2^7$)	1.6e-8	-4.4e-8	3.2e-7	-1.2e-3	5.8e-10	0.0e0	1.6e-7	9.2e-9	3.8e-3	4.2e-9	7.7e-13	1.8e-13	2.7e-4	4.3e-11	1.2e-10	3.8e-8	-1.1e-3	4.8e0	4.9e1	1.0e2	3.4e1	1.9e3
($2^{14}, 2^{10}$)	4.1e-9	3.9e-7	-1.1e-1	-5.6e-4	8.7e-9	3.6e-12	0.0e0	2.0e0	1.8e-2	8.0e-11	7.0e-13	9.7e-11	4.4e-4	-7.8e-10	4.0e-9	9.5e-7	-5.1e-4	8.2e1	2.0e2	4.6e2	1.1e2	3.6e3

(b) Diagonal quadratic objective

Problem	Relative Pobj				Pinfeas				Dinfeas				Relgap				Walltime [s]					
	GRB	G-OA	PSG	OSQP	ALM	GRB	G-OA	PSG	OSQP	ALM	GRB	G-OA	OSQP	ALM	GRB	G-OA	OSQP	ALM	GRB	G-OA	PSG	OSQP
(m, n)																						
$L = 1, k = 1\%m$																						
$(2^{20}, 2^7)$	3.2e-10	2.0e-7	7.3e-11	—	1.4e-9	0.0e0	0.0e0	1.4e-8	—	1.2e-11	4.9e-14	8.3e-15	—	-1.0e-10	4.2e-10	3.9e-7	—	1.1e1	5.9e2	9.2e1	9.5e1	3.6e3
$(2^{17}, 2^7)$	3.7e-9	4.7e-7	3.0e-9	-5.3e-4	9.7e-9	0.0e0	0.0e0	6.1e-14	2.0e-3	5.5e-10	8.3e-15	1.0e-14	1.5e-5	-1.2e-9	4.9e-9	9.2e-7	-5.2e-4	9.5e-1	3.0e1	6.1e0	8.1e0	4.4e2
$(2^{17}, 2^{10})$	9.9e-10	4.2e-8	-5.0e-2	-5.3e-5	1.1e-9	0.0e0	0.0e0	4.5e0	7.1e-4	4.1e-11	5.1e-13	6.3e-10	7.2e-6	-7.5e-10	4.6e-10	8.3e-8	-5.1e-5	1.2e1	1.5e2	2.6e2	1.5e2	2.4e3
$L = 1, k = 10\%m$																						
$(2^{20}, 2^7)$	3.2e-9	1.0e-5	1.7e-9	—	7.7e-9	0.0e0	0.0e0	8.2e-9	—	2.6e-11	2.8e-14	1.5e-14	—	-1.9e-9	2.4e-9	2.0e-5	—	1.8e1	2.4e2	3.7e2	1.2e2	3.6e3
$(2^{17}, 2^7)$	4.7e-9	8.0e-7	4.4e-10	-3.2e-4	2.1e-9	0.0e0	0.0e0	1.2e-8	7.4e-4	1.6e-11	8.9e-15	8.6e-15	4.2e-5	-4.3e-10	8.3e-9	1.6e-6	-3.3e-4	2.0e0	2.1e1	3.3e1	1.4e1	4.0e2
$(2^{17}, 2^{10})$	5.5e-9	3.5e-7	-7.4e-2	-1.7e-4	6.8e-9	0.0e0	0.0e0	3.4e0	1.7e-3	9.7e-10	3.7e-14	2.1e-14	1.1e-5	-2.6e-9	5.7e-9	6.8e-7	-1.7e-4	2.7e1	1.7e2	3.1e2	1.2e2	3.2e3
$L = 10, k = 1\%m$																						
$(2^{17}, 2^7)$	4.5e-10	9.1e-8	3.2e-10	-1.1e-2	3.3e-9	0.0e0	0.0e0	8.2e-14	5.6e-2	5.6e-11	8.2e-13	1.4e-14	1.3e-3	-3.0e-10	1.3e-9	5.9e-7	-2.1e-3	1.5e1	1.1e3	1.5e2	1.2e2	3.6e3
$(2^{14}, 2^7)$	9.0e-10	2.2e-6	2.1e-9	-1.7e-3	4.6e-9	0.0e0	0.0e0	1.3e-14	1.4e-2	1.5e-10	8.5e-13	1.5e-14	2.8e-4	-2.5e-10	2.2e-9	7.0e-6	-1.5e-3	1.6e0	5.2e1	9.7e0	1.7e1	5.4e2
$(2^{14}, 2^{10})$	3.2e-9	1.9e-8	-5.2e-2	-9.6e-5	6.1e-9	0.0e0	0.0e0	6.5e0	9.2e-3	1.2e-9	1.7e-13	2.3e-14	1.3e-4	-3.1e-9	4.7e-11	3.6e-8	-9.4e-5	2.3e1	2.8e2	7.1e1	1.3e2	2.0e3
$L = 10, k = 10\%m$																						
$(2^{17}, 2^7)$	1.5e-9	1.6e-8	2.1e-3	-4.7e-3	6.3e-9	0.0e0	0.0e0	1.1e-8	1.7e-2	2.3e-11	5.2e-14	4.3e-14	1.0e-4	-1.3e-9	6.9e-10	7.2e-8	-6.5e-4	2.0e1	5.6e2	1.7e3	1.1e2	3.6e3
$(2^{14}, 2^7)$	4.6e-10	4.0e-8	1.5e-10	-7.2e-4	1.2e-9	0.0e0	0.0e0	2.2e-9	2.3e-3	1.4e-10	1.7e-14	8.3e-15	3.0e-4	-1.4e-10	4.4e-9	3.9e-7	-6.7e-4	2.2e0	4.0e1	1.1e2	1.8e1	1.4e3
$(2^{14}, 2^{10})$	2.6e-9	1.3e-8	-6.4e-2	-6.3e-4	7.9e-10	0.0e0	0.0e0	3.6e0	2.0e-2	9.4e-11	1.6e-14	4.6e-15	1.3e-4	-2.7e-10	5.3e-9	2.8e-8	-5.6e-4	4.9e1	3.3e2	8.2e2	1.7e2	3.6e3

Table 3: Performance of different methods for synthetic data. *Relative Pobj* := $(\text{Pobj} - \text{Pobj}_{\text{ALM}}) / |\text{Pobj}_{\text{ALM}}|$, and a positive value indicates that the method achieved a higher objective value than ALM. Tolerance for OSQP is set at $\epsilon = 10^{-3}$, while for all others are set to 10^{-8} . Note: PSG does not provide metrics for dual infeasibility or relative gap.

Problem (m, n)	KKT residual		Time _a : $\epsilon = 10^{-3}$		Time _b : $\epsilon = 10^{-6}$		Time _c : $\epsilon = 10^{-8}$		% Time: Sort		% Time: $V^{-1}v$		% Time: $\nabla\varphi$		ALM SSN iter		avg $ \beta $		final $ \beta $	
	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.	lin.	quad.
$L = 1, k = 10\%m$																				
(20, 27)	3.3e-9	1.4e-9	4.0e0	4.7e0	1.0e1	8.8e0	2.6e1	1.1e1	42	16	4	6	12	20	18	100	17	70	380	290
(210, 27)	6.7e-9	2.0e-9	1.9e0	1.9e0	7.3e0	3.8e0	1.3e1	4.4e0	34	14	5	7	16	21	19	130	17	60	270	200
(218, 27)	1.5e-10	2.6e-9	9.6e-1	1.1e0	2.5e0	1.9e0	5.5e0	2.1e0	39	14	5	7	14	21	18	90	16	50	210	150
(217, 27)	2.3e-9	9.7e-9	4.5e-1	4.7e-1	1.4e0	8.7e-1	1.6e0	9.5e-1	22	10	8	8	19	22	15	70	15	50	200	130
(218, 28)	7.0e-9	1.2e-9	5.4e0	4.7e0	1.2e1	7.6e0	1.7e1	8.9e0	18	10	8	9	24	27	18	120	15	70	350	230
(218, 28)	6.3e-9	9.9e-9	2.1e0	2.6e0	6.8e0	4.1e0	9.8e0	4.4e0	21	11	9	10	23	26	19	140	14	70	300	190
(217, 28)	4.6e-10	4.3e-9	1.4e0	1.1e0	2.8e0	1.7e0	4.0e0	1.9e0	18	7	11	11	23	27	20	130	13	60	270	170
(218, 29)	4.0e-10	9.0e-10	6.3e0	6.0e0	1.5e1	8.4e0	1.8e1	9.1e0	9	5	16	13	29	31	20	160	13	80	500	290
(217, 29)	9.3e-9	7.1e-9	4.0e0	3.3e0	7.4e0	4.7e0	8.4e0	5.0e0	8	4	21	17	27	30	17	140	13	90	480	290
(217, 210)	9.4e-9	1.1e-9	1.1e1	9.1e0	2.1e1	1.1e1	2.3e1	1.2e1	3	2	38	24	24	30	16	170	13	100	900	500
$L = 1, k = 10\%m$																				
(220, 27)	7.5e-9	7.7e-9	7.3e0	8.2e0	2.7e1	1.5e1	2.7e1	1.8e1	19	17	30	31	15	14	18	110	14	80	450	400
(219, 27)	3.5e-9	3.7e-9	4.7e0	3.9e0	1.1e1	6.9e0	1.3e1	8.4e0	20	17	29	31	14	15	20	100	15	80	310	270
(218, 27)	7.2e-9	2.4e-9	2.2e0	2.2e0	7.8e0	4.0e0	8.1e0	4.5e0	16	17	32	31	17	15	19	160	15	80	240	190
(217, 27)	7.5e-11	2.1e-9	1.0e0	9.8e-1	2.9e0	1.7e0	2.9e0	2.0e0	22	16	27	30	15	15	18	100	15	70	200	160
(219, 28)	1.7e-9	4.9e-9	1.0e1	1.1e1	3.0e1	1.7e1	3.1e1	1.8e1	13	11	37	37	18	18	19	160	14	100	370	280
(218, 28)	2.7e-9	8.6e-9	6.3e0	5.5e0	1.3e1	8.3e0	1.4e1	8.9e0	14	13	36	36	17	18	18	140	14	100	320	220
(217, 28)	9.9e-10	1.7e-9	3.0e0	2.5e0	7.7e0	3.7e0	8.0e0	4.1e0	16	7	35	40	19	19	17	160	14	100	290	190
(218, 29)	5.5e-9	1.9e-9	2.2e1	1.5e1	3.4e1	2.1e1	2.2e1	2.2e1	10	6	42	43	19	20	22	220	16	130	520	330
(217, 29)	6.3e-9	1.8e-9	1.1e1	8.3e0	1.7e1	1.1e1	1.8e1	1.1e1	6	6	46	44	19	20	21	200	15	140	500	280
(217, 210)	9.6e-9	6.8e-9	3.4e1	2.0e1	5.3e1	2.6e1	5.6e1	2.7e1	4	3	56	50	17	20	22	280	15	160	940	570
$L = 10, k = 1\%m$																				
(217, 27)	2.0e-9	3.3e-9	7.1e0	1.1e1	1.9e1	1.4e1	2.5e1	1.5e1	27	14	3	3	18	24	18	110	16	90	160	110
(216, 27)	7.3e-10	2.9e-9	2.9e0	4.2e0	1.0e1	5.7e0	1.2e1	6.3e0	27	15	3	3	17	22	16	100	15	70	160	100
(215, 27)	4.0e-10	6.9e-10	3.3e0	2.7e0	6.8e0	3.4e0	7.8e0	3.6e0	20	15	5	4	22	23	18	170	16	80	150	80
(214, 27)	4.3e-7	4.6e-9	1.4e0	1.0e0	3.0e0	1.5e0	4.6e0	1.6e0	11	13	5	6	17	23	24	160	15	70	310	100
(216, 28)	8.7e-11	3.0e-9	1.1e1	1.0e1	2.3e1	1.3e1	2.9e1	1.3e1	17	8	5	6	26	30	19	190	14	100	240	160
(215, 28)	6.8e-9	2.3e-9	5.9e0	5.0e0	1.0e1	6.3e0	1.2e1	6.9e0	12	8	9	7	27	30	21	170	14	100	250	140
(214, 28)	2.6e-9	1.2e-9	2.9e0	2.9e0	5.3e0	3.7e0	5.6e0	3.9e0	12	9	11	9	28	30	15	160	13	120	250	140
(215, 29)	1.0e-9	5.4e-9	1.4e1	1.3e1	3.1e1	1.7e1	3.1e1	1.8e1	8	5	14	11	31	34	20	240	14	150	470	260
(214, 29)	8.8e-9	2.4e-9	7.1e0	6.2e0	1.2e1	9.0e0	1.4e1	9.3e0	9	5	19	14	28	32	22	210	15	140	480	270
(214, 210)	4.5e-9	6.1e-9	2.7e1	1.8e1	3.9e1	2.2e1	4.2e1	2.3e1	3	3	35	21	26	33	21	270	14	180	890	500
$L = 10, k = 10\%m$																				
(217, 27)	6.1e-9	6.3e-9	1.0e1	1.0e1	2.7e1	1.7e1	6.0e1	2.0e1	38	18	11	17	13	19	17	190	15	90	220	170
(216, 27)	1.0e-11	7.9e-9	5.1e0	4.9e0	1.2e1	8.3e0	1.8e1	9.1e0	31	21	10	11	15	19	17	130	14	90	170	110
(215, 27)	3.2e-9	4.8e-9	3.5e0	3.3e0	7.3e0	5.0e0	9.6e0	9.3e0	30	21	10	9	17	23	18	170	15	220	150	90
(214, 27)	4.2e-9	1.2e-9	1.4e0	1.5e0	3.9e0	2.1e0	4.8e0	2.2e0	33	19	6	8	17	22	16	170	14	100	130	80
(216, 28)	2.7e-9	6.3e-9	1.5e1	1.4e1	3.2e1	1.8e1	4.1e1	2.0e1	16	14	19	13	22	25	20	230	15	120	280	170
(215, 28)	4.8e-10	1.4e-9	8.7e0	7.5e0	1.7e1	1.1e1	1.9e1	1.2e1	19	15	16	16	23	25	17	220	15	150	270	160
(214, 28)	3.4e-9	1.5e-9	1.3e1	4.3e0	1.3e1	5.4e0	2.2e1	5.7e0	25	15	16	15	21	25	21	460	15	140	240	160
(215, 29)	9.7e-9	6.8e-10	3.1e1	2.0e1	5.0e1	3.4e1	5.2e1	3.5e1	12	8	27	27	25	27	17	320	16	240	480	260
(214, 29)	8.1e-10	4.0e-9	1.5e1	1.1e1	2.5e1	1.6e1	2.9e1	1.6e1	11	8	29	27	24	27	22	350	15	220	470	280
(214, 210)	8.7e-9	7.9e-10	4.7e1	3.6e1	7.5e1	4.8e1	8.2e1	4.9e1	5	4	48	39	21	25	19	420	15	310	900	500

Table 4: Detailed ALM performance on synthetic instances with linear (lin.) and convex quadratic (quad.) objective. The column %T_c indicates the proportion of total solve time; columns $\nabla\varphi$ and $V^{-1}v$ report extra time, assuming that φ and $\nabla\varphi$ have already been evaluated, respectively.

to approximate the full distribution over a region of interest by solving prediction problems for the following 22 quantiles: $\{10^{-3}, 10^{-2}, 0.025, 0.05, \dots, 0.475, 0.5\}$, each with a tolerance of $\epsilon = 10^{-4}$. With the exception of the first quantile, each subsequent instance is warm-started using the solution from the preceding smaller quantile. Experiments are conducted in a Linux environment on a Dell workstation with eight physical Intel(R) Xeon(R) W-2145 CPUs @ 3.70GHz with two threads per core and 125GB of RAM. Gurobi-based methods required about 85GB of RAM, and limited the number of features used for prediction, while ALM used about 40GB of RAM.

Numerical experiments are summarized in Tables 5 and 6. Table 5 shows that ALM is up to $\approx 20\times$ faster than all three of the Gurobi-based methods for small (or large, by using $\bar{\tau}$) quantile problems. For the most difficult case when $\tau = 0.5$, ALM required 1.5x the time needed by the most efficient Gurobi method (G-QR). When solving instances from Table 5, the Gurobi methods were able to find feasible primal and dual solutions to numerical precision, but the error η is determined by the duality gap. In large-scale problems, the linear systems solved by the barrier method can become prohibitively ill-conditioned. Even the G-OA method encountered this problem, as the set of active scenarios could change significantly from one iteration to the next. In contrast, ALM has handled these challenging instances more reliably.

Table 6 summarizes the results from the second experiment, demonstrating that warm-starting can substantially decrease the computational time for challenging larger- τ cases. For instance, when warm-started, ALM can solve the $\tau = 0.5$ case approximately 1.4 times faster than the best-performing Gurobi method. Overall, ALM can obtain the solution path for ≈ 40 quantiles ranging from $\tau = 10^{-3}$ to $\tau = 1 - 10^{-3}$ with a precision of $\epsilon = 10^{-4}$ in roughly the same duration that the most efficient Gurobi method takes to solve two problems in the $\tau \approx 1\%$ range.

$k = \tau\%m$	Time [s]				η			
	ALM	GRB	G-OA	G-QR	ALM	GRB	G-OA	G-QR
1	1.3e2	3.6e3	6.2e3	3.6e3	4.4e-9	2.5e-1	2.9e-1	2.3e1
10	2.1e2	3.6e3	4.2e3	3.6e3	2.5e-9	2.7e-5	7.3e-3	1.1e-2
25	6.4e2	7.6e2	4.2e3	2.5e3	2.5e-9	4.3e-9	1.6e-5	8.4e-9
50	3.6e3	8.2e2	3.6e3	1.7e3	3.6e-4	1.4e-9	6.9e-4	9.6e-8

Table 5: Results for the computation of the solution path for the cold-start quantile regression with an accuracy of $\epsilon = 10^{-8}$. The time limit is one hour.

6. Conclusions

In this paper, we introduced a semismooth-Newton-based augmented Lagrangian method for the superquantile-constrained optimization problem (1) when the uncertain parameter ω is supported on a finite set of scenarios with equal probability. The proposed approach ensures that the generalized Jacobian in the semismooth Newton method can be obtained cheaply via the structured sparsity inherent to the superquantile operator. In numerical experiments conducted on synthetic and realistic datasets where the number of scenarios greatly exceeds the number of decision variables ($m \gg n$), we find that the ALM improves upon off-the-shelf solvers such as Gurobi and OSQP as well as manually-tuned active set outer-approximation methods, and can be faster and more reliable than the specialized software suite PSG.

$k = \tau\%m$	T := Time [s]	η	ALM iter	% T: $V^{-1}v$	% T: $\nabla\varphi$	% T: sort	% T: $\text{proj}_{\bar{B}}$
1.00e-3	1.3e2	6.5e-5	10	5	10	14	11
1.00e-2	1.3e2	3.2e-5	4	1	5	53	12
2.50e-2	1.4e2	3.7e-5	4	2	4	55	12
5.00e-2	1.7e2	4.2e-5	5	4	4	54	11
7.50e-2	1.5e2	8.5e-5	3	4	3	61	11
1.00e-1	1.5e2	4.6e-5	3	5	3	60	10
1.25e-1	1.5e2	6.7e-5	2	4	2	64	11
1.50e-1	1.9e2	6.0e-5	4	7	3	56	10
1.75e-1	2.0e2	5.4e-5	4	7	3	57	10
2.00e-1	1.8e2	8.1e-5	3	7	2	59	9
2.25e-1	1.8e2	8.7e-5	3	7	2	60	9
2.50e-1	1.9e2	8.2e-5	3	8	2	59	9
2.75e-1	2.2e2	5.7e-5	3	8	2	60	9
3.00e-1	2.2e2	3.4e-5	3	8	2	60	9
3.25e-1	2.0e2	5.5e-5	2	7	1	64	9
3.50e-1	1.9e2	5.1e-5	2	8	2	62	9
3.75e-1	2.3e2	3.9e-5	3	9	2	59	8
4.00e-1	2.3e2	4.7e-5	3	10	2	58	8
4.25e-1	2.2e2	8.7e-5	2	7	1	64	9
4.50e-1	2.1e2	6.2e-5	2	8	1	62	8
4.75e-1	2.7e2	3.3e-5	3	10	1	60	8
5.00e-1	3.3e2	7.3e-5	4	12	1	57	8
Total time:	4.3e3	—	—	(avg) 7	(avg) 2	(avg) 58	(avg) 9

Table 6: Results for the computation of the solution path with an accuracy of $\epsilon = 10^{-4}$. The last four columns show the proportion of time spent computing $V^{-1}v$, $\nabla\varphi$, sorting, and $\text{proj}_{\bar{B}}$ relative to the total runtime. The final row presents the average proportion across all instances.

Finally, when the uncertain function $g^\ell(\cdot; \omega^j)$ is nonlinear, the ALM is still expected to exhibit improved performance relative to the generic nonlinear solvers. Underlying this anticipation is the form of the generalized Jacobian

$$\sum_{i=1}^{mL} \nabla^2 G_i(x) \cdot (G(x) - \text{proj}_{\bar{B}}(G(x)))_i + JG(x)^\top \cdot (I - J_{\bar{B}}) \cdot JG(x),$$

which only depends on the Hessian of the effective tail scenarios. Second-order solvers that are not able to exploit this sparsity must evaluate the Hessian on the entire scenario set, which could become prohibitively expensive. Further improvements are expected when the function G is nonlinear but with a simple structure, since the corresponding Hessian $\nabla^2 G_i(x)$ in the first term may be computed cheaply and the second term is always cheap to compute due to the special sparsity structure of $(I - J_{\bar{B}})$, whereas standard convex solvers may have to resort to extended reformulations.

Acknowledgements

The authors are partially supported by NSF grants CCF-2416172 and DMS-2416250, and NIH grant R01CA287413. The authors thank Eric Sager Luxenberg and Johannes Royset for helpful discussion of the algorithms presented in the manuscript.

References

- F. N. Airaud, H. Antil, R. Lohner, and U. Rakhimov. On the use of risk measures in digital twins to identify weaknesses in structures, 2023. [arXiv:2311.12206](#).
- American Optimal Decisions, Inc. Portfolio safeguard (PSG) in windows shell environment, 2011. <http://www.aorda.com/>.
- S. Arpòn, T. Homem-de-Mello, and B. Pagnoncelli. Scenario reduction for stochastic programs with conditional value-at-risk. *Math. Program., Ser. B*, 170:327–356, 2018.
- H. G. Basova, R. T. Rockafellar, and J. O. Royset. A computational study of the buffered failure probability in reliability-based design optimization. In *Proc. 11th Int. Conf. Appl. Stat. Probab. Civ. Eng.*, Zurich, Switzerland, 2011.
- Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proc. 26th Int. Conf. Mach. Learn.*, pages 41–48, 2009.
- D. Bertsimas and D. B. Brown. Constructing uncertainty sets for robust linear optimization. *Oper. Res.*, 57(6):1483–1495, 2009.
- J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM Rev.*, 59(1):65–98, 2017.
- A. Chaudhuri, B. Kramer, M. Norton, J. O. Royset, and K. Willcox. Certifiable risk-based engineering design optimization. *Am. Inst. Aeronaut. Astronaut. J.*, 60(2):551–565, 2022.
- W. Chen, M. Sim, J. Sun, and C.-P. Teo. From CVaR to uncertainty set: Implications in joint chance-constrained optimization. *Oper. Res.*, 58(2):470–485, 2010.
- Y. Cui, D. Sun, and K.-C. Toh. On the R-superlinear convergence of the KKT residuals generated by the augmented Lagrangian method for convex composite conic programming. *Math. Program. Ser. A*, 178:381–415, 2019.
- D. Davis. An $O(n \log(n))$ algorithm for projecting onto the ordered weighted ℓ_1 norm ball, 2015. [arXiv:1505.00870](#).
- C. Ding. Variational analysis of the Ky Fan k -norm. *Set-Valued Var. Anal.*, 25:265–296, 2016.
- M. V. Dolgopolik. A note on the generalised Hessian of the least squares associated with systems of linear inequalities, 2024. [arXiv:2404.15571](#).
- F. Facchinei and J.-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer Science & Business Media, New York, 2007.
- C. Fröhlich and R. C. Williamson. Risk measures and upper probabilities: Coherence and stratification, 2022. [arXiv:2206.03183](#).
- Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2nd edition, 2013.

- R. Koenker and G. Bassett, Jr. Regression quantiles. *Econometrica*, 46(1):33–50, 1978.
- Y. Laguel, K. Pillutla, J. Malick, and Z. Harchaoui. Superquantiles at work: Machine learning applications and efficient subgradient computation. *Set-Valued Var. Anal.*, 29:967–996, 2021.
- L. Leqi, A. Prasad, and P. K. Ravikumar. On human-aligned risk minimization. In *Proc. 32nd Adv. Neural Inf. Process. Syst.*, pages 15029–15038, 2019.
- D. Levya, Y. Carmon, J. C. Duchi, and A. Sidford. Large-scale methods for distributionally robust optimization. In *Proc. 33rd Adv. Neural Inf. Process. Syst.*, pages 8847–8860, 2020.
- Q. Li and X. Li. Fast projection onto the ordered weighted ℓ_1 norm ball. *Sci. China Math.*, 65: 869–886, 2021.
- X. Li, D. Sun, and K.-C. Toh. An asymptotically superlinearly convergent semismooth Newton augmented Lagrangian method for linear programming. *SIAM J. Optim.*, 30(3):2410–2440.
- X. Li, D. Sun, and K.-C. Toh. On efficiently solving the subproblems of a level-set method for fused lasso problems. *SIAM J. Optim.*, 28(2):1842–1866, 2018a.
- X. Li, D. Sun, and K.-C. Toh. A highly efficient semismooth Newton augmented Lagrangian method for solving Lasso problems. *SIAM J. Optim.*, 28(1):433–458, 2018b.
- S. Lyu, Y. Fan, Y. Ying, and B.-G. Hu. Average top-k aggregate loss for supervised learning. *IEEE Trans. on Pattern Anal. and Mach. Intell.*, 44(1):76–86, 2020.
- A. Nemirovski and A. Shapiro. Convex approximations of chance constrained programs. *SIAM J. Optim.*, 17(4):969–996, 2006.
- NYC Taxi and Limousine Commission. NYC TLC Trip Record Data, 2023. <https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- OpenWeather. OpenWeather API, 2023. <https://openweathermap.org/>.
- L. Peng, Y. Travadi, R. Zhang, Y. Cui, and J. Sun. Imbalanced classification in medical imaging via regrouping, 2022. [arXiv:2210.12234](https://arxiv.org/abs/2210.12234).
- G. C. Pflug. *Probabilistic Constrained Optimization. Nonconvex Optimization and Its Applications*, chapter: Some Remarks on the Value-at-Risk and the Conditional Value-at-Risk. Springer, Boston, MA, 2000.
- H. Rahimian, G. Bayraksan, and T. Homem-de-Mello. Identifying effective scenarios in distributionally robust stochastic programs with total variation distance. *Math. Program., Ser. A*, 173: 393–430, 2018.
- A. Robey, L. Chamon, G. J. Pappas, and H. Hassani. Probabilistically robust learning: Balancing average and worst-case performance. In *Proc. 39th Int. Conf. Mach. Learn.*, pages 18667–18686, 2022.
- S. M. Robinson. Local structure of feasible sets in nonlinear programming, part II: Nondegeneracy. *Math. Program. Study* 22, pages 217–230, 1984.

- S. M. Robinson. Local structure of feasible sets in nonlinear programming, part III: Stability and sensitivity. *Math. Program. Study* 30, pages 45–66, 1987.
- R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- R. T. Rockafellar. Augmented Lagrangians and applications of the proximal point algorithm in convex programming. *Math. Oper. Res.*, 1(2):97–116, 1976a.
- R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM J. Control and Optim.*, 14(5), 1976b.
- R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *J. of Risk*, 2(3):21–41, 1999.
- J. Roth and Y. Cui. On $O(n)$ algorithms for projection onto the top- k -sum sublevel set. *Math. Prog. Comput.*, 17, 2025.
- J. O. Royset. Risk-adaptive approaches to learning and decision making: A survey, 2022. [arXiv:2212.00856](#).
- J. O. Royset and R. J.-B. Wets. *An Optimization Primer*. Springer Nature Switzerland, 2021.
- A. Shapiro. Sensitivity analysis of nonlinear programs and differentiability properties of metric projections. *SIAM J. Control and Optim.*, 26(3):628–645, 1988.
- A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 761–769, 2016.
- B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: An operator splitting solver for quadratic programs. *Math. Program. Comput.*, 12(4):637–672, 2020.
- A. Tamar, Y. Glassner, and S. Mannor. Optimizing the CVaR via sampling. In *Proc. 29th AAAI Conf. Artif. Intell.*, pages 2993–2999. AAAI, 2015.
- J. W. J. Williams. Heapsort. *Comm. of the ACM*, 7, 1964.
- R. Williamson and A. Menon. Fairness risk measures. In *Proc. 36th Int. Conf. Mach. Learn.*, volume 97, pages 6786–6797. PMLR, 2019.
- B. Wu, C. Ding, D. Sun, and K.-C. Toh. On the Moreau–Yosida regularization of the vector k -norm related functions. *SIAM J. Optim.*, 24(2):766–794, 2014.
- C. Wu, Y. Cui, D. Li, and D. Sun. Convex and nonconvex risk-based linear regression at scale. *INFORMS J. Comput.*, 35(4):797–816, 2023.
- D. Yang. Robust machine learning using superquantiles, 2021. Naval Postgraduate School.
- T. Yang. Algorithmic foundation of deep X-risk optimization, 2022. [arXiv:2206.00439](#).
- P. Yuan, X. You, H. Chen, Q. Peng, Y. Zhao, Z. Xu, X.-Y. Jing, and Z. He. Group sparse additive machine with average top- k loss. *Neurocomputing*, 395:1–14, 2020. ISSN 0925-2312.

Y. Zhang, N. Zhang, D. Sun, and K.-C. Toh. An efficient Hessian based algorithm for solving large-scale sparse group lasso problems. *Math. Program., Ser. A*, 179:223–263, 2020.

J. Zowe and S. Kurcyusz. Regularity and stability for the mathematical programming problem in Banach spaces. *Appl. Math. Optim.*, 5:49–62, 1979.

Appendix A. Proofs of Lemma 2 and Lemma 3

Proof of Lemma 2. Case 1 is clear. Case 2 relies on the characterization of the critical cone given in (10). The corresponding projection problem can be stated as follows. For any $y^0 \in \mathbb{R}^m$ with $\mathcal{T}_k(y^0) > r$, the directional derivative of $\text{proj}_{\mathcal{B}_k}$ along direction (u^0, v^0, w^0) is given by projecting (u^0, v^0, w^0) onto the critical cone $\bar{\mathcal{D}}$ at $\bar{y} = \text{proj}_{\mathcal{B}_k}(\bar{y}^0)$. In Case 2, the projection is onto a linear subspace, and hence the linear operator can be recovered by inspecting the projection onto the critical cone. It can be shown (Wu et al., 2014, Section 4.3.4) that for $u \in \mathbb{R}^{|\bar{\alpha}|}$, $v \in \mathbb{R}^{|\bar{\beta}|}$, $w \in \mathbb{R}^{|\bar{\gamma}|}$, this problem can be expressed as

$$\begin{aligned} \min_{u,v} \quad & \frac{1}{2} \|u - u^0\|_2^2 + \frac{1}{2} \|v - v^0\|_2^2 + \frac{1}{2} \|w - w^0\|_2^2 \\ \text{s.t.} \quad & v_i \geq v_{[k-\bar{k}_0]}, \quad \forall i \in \{1, \dots, |\bar{\beta}_1|\}, \\ & v_i = v_{[k-\bar{k}_0]}, \quad \forall i \in \{|\bar{\beta}_1| + 1, \dots, |\bar{\beta}_1| + |\bar{\beta}_2|\}, \\ & v_i \leq v_{[k-\bar{k}_0]}, \quad \forall i \in \{|\bar{\beta}_1| + |\bar{\beta}_2| + 1, \dots, |\bar{\beta}_1| + |\bar{\beta}_2| + |\bar{\beta}_3|\}, \\ & \mathbf{1}_{\bar{\alpha}}^\top u + \bar{\mu}_{\bar{\beta}}^\top v = 0 \end{aligned} \tag{33}$$

where $\bar{\beta}_1 = \{i \in \bar{\beta} : (\bar{\mu}_{\bar{\beta}})_i = 1\}$, $\bar{\beta}_2 = \{i \in \bar{\beta} : (\bar{\mu}_{\bar{\beta}})_i \in (0, 1)\}$, and $\bar{\beta}_3 = \{i \in \bar{\beta} : (\bar{\mu}_{\bar{\beta}})_i = 0\}$. Problem (33) is obtained by using Lemma 2.5 of Wu et al. (2014), where σ is a permutation of $\bar{\beta}$ that sorts the $\bar{\beta}_1$ components, preserves the $\bar{\beta}_2$ components, and sorts the $\bar{\beta}_3$ components so that $(v_\sigma)_{\bar{\beta}_1} = (\bar{v}_{\bar{\beta}_1})$, $(v_\sigma)_{\bar{\beta}_2} = (\bar{v}_{\bar{\beta}_2})$, and $(v_\sigma)_{\bar{\beta}_3} = (\bar{v}_{\bar{\beta}_3})$. The condition $\bar{y}_{k_0+1}^0 - \bar{\lambda} < \bar{y}_k < \bar{y}_{k_0+1}^0$ implies that $0 < \bar{\mu}_{\bar{\beta}} = (\bar{y}_{\bar{\beta}}^0 - \bar{y}_{\bar{\beta}})/\bar{\lambda} < \mathbf{1}_{\bar{\beta}}$ since

$$\bar{y}_{k_1}^0 - \bar{\theta} > 0 \implies \dots \implies \bar{y}_{k_0+1}^0 - \bar{\theta} > 0, \quad \bar{y}_{k_0+1}^0 - \bar{\theta} < \bar{\lambda} \implies \dots \implies \bar{y}_{k_1}^0 - \bar{\theta} < \bar{\lambda}.$$

Thus $\bar{\beta}_1 \cup \bar{\beta}_3 = \emptyset$. Since $\bar{\beta}_1 \cup \bar{\beta}_3 = \emptyset$, projection onto the critical cone simplifies in (33), and the permutation σ may be disregarded. The linear subspace is then easily verified as that defined in (12) by noting that the equality constraints can be arranged successively.

Case 3 follows the same argument as Case 2, where the constraints on v_i for $i \in \{1, \dots, |\bar{\beta}_1|\}$ in (33) are automatically satisfied by the sorting. \square

Proof of Lemma 3. Let $h(y) := \frac{1}{2} \|\max(\bar{g} - y, 0)\|_2^2$. Then h is continuously differentiable with gradient $\nabla h(y) = -\max(\bar{g} - y, 0)$. Since both (20) and (21) are convex, in order to prove this lemma, it suffices to show that $\bar{y} - \bar{g} \in (\mathcal{T}_{\mathcal{B}_k}(\bar{y}))^* \implies -\max(\bar{g} - \bar{y}, 0) \in (\mathcal{T}_{\mathcal{B}_k}(\bar{y}))^*$ is an element in the dual of the tangent cone of \mathcal{B}_k at \bar{y} . Using the form of the KKT conditions which ensures that $\bar{y} \leq \bar{g}$, the desired conclusion holds trivially since

$$\bar{y} - \bar{g} = \max(\bar{y} - \bar{g}, 0) + \min(\bar{y} - \bar{g}, 0) = \max(\bar{y} - \bar{g}, 0) - \max(\bar{g} - \bar{y}, 0) = -\max(\bar{g} - \bar{y}, 0). \quad \square$$

Appendix B. Constraint non-degeneracy

Here we describe the constraint non-degeneracy conditions of (1) with a single superquantile term in the objective function. Such conditions are used in Section 3.1 when discussing convergence guarantees for the ALM. The constraint non-degeneracy follows the approach outlined in Ding (2016).

Since the top- k -sum operator T_k is positive homogeneous, we can consider the conic problem

$$\begin{aligned} \min_{x \in X} \quad & f(x) + t \\ \text{s.t.} \quad & (G(x), t) \in \mathcal{K}_{(k)} := \text{epi}(\mathsf{T}_k(\cdot)), \end{aligned} \quad (34)$$

where epi denotes the epigraph. Let $\mathcal{G}(x, t) := (G(x), t)$. Then at any $(\bar{x}, \bar{t}) \in X \times \mathbb{R}$, Robinson's constraint qualification

$$J\mathcal{G}(\bar{x}, \bar{t})(X \times \mathbb{R}) + \mathcal{T}_{\mathcal{K}_{(k)}}(\mathcal{G}(\bar{x}, \bar{t})) = \mathbb{R}^m \times \mathbb{R}$$

holds since there exists an interior to the feasible region by taking t large enough. In particular, this means that at an optimal solution $x^* \in X$, the corresponding set of multipliers are nonempty, convex, and compact (Zowe and Kurcyusz, 1979). By Robinson (1984, 1987), the constraint non-degeneracy holds at $(\bar{x}, \mathsf{T}_k(G(\bar{x})))$ if and only if

$$JG(\bar{x})X + \mathcal{T}^{\text{lin}}(G(\bar{x})) = \mathbb{R}^m, \quad (35)$$

where $\mathcal{T}^{\text{lin}}(y)$ is the linear subspace of \mathbb{R}^m defined by

$$\mathcal{T}^{\text{lin}}(y) := \{d \in \mathbb{R}^m : \mathsf{T}'_k(y; d) = -\mathsf{T}'_k(y; -d)\}.$$

Using (Rockafellar, 1970, Theorem 23.4) and the form of $\partial \mathsf{T}_k(y)$ in Wu et al. (2014, Lemma 2.2), the directional derivative of the top- k -sum operator at a point $y \in \mathbb{R}^m$ along a direction $d \in \mathbb{R}^m$ is given by

$$\mathsf{T}'_k(y; d) = \sup_{\mu} \{\langle \mu, d \rangle : \mu \in \partial \mathsf{T}_k(y)\} = \mathbf{1}_{\alpha}^{\top}(d_{\pi})_{\alpha} + \mathsf{T}_{k-k_0}((d_{\pi})_{\beta}), \quad (36)$$

where α, β, γ are the index-sets associated with $y_{\pi} = \bar{y}$, and k_0, k_1 are the corresponding sorting indices. That is, the directional derivative is linear in the α components. Letting $\bar{y} := G(\bar{x})$, we obtain

$$\mathcal{T}^{\text{lin}}(\bar{y}) = \{d \in \mathbb{R}^m : \mathsf{T}_k((d_{\pi})_{\beta}) = -\mathsf{T}_k(-(d_{\pi})_{\beta})\} = \{d \in \mathbb{R}^m : (d_{\pi})_{\beta} = s\mathbf{1}_{\beta}, s \in \mathbb{R}\}. \quad (37)$$

In the case when $X = \mathbb{R}^n$, and we assume that $\bar{y} = G(\bar{x})$ is sorted, then system (35) is equivalent to the existence of $v \in \mathbb{R}^n$ for any $w \in \mathbb{R}^m$ such that

$$\begin{aligned} JG(\bar{x})v + \mathcal{T}^{\text{lin}}(\bar{y}) = w & \iff \begin{bmatrix} JG(\bar{x})_{\alpha,:} \\ JG(\bar{x})_{\beta,:} \\ JG(\bar{x})_{\gamma,:} \end{bmatrix} v + \begin{bmatrix} u_{\alpha} \\ u_{\beta} \\ u_{\gamma} \end{bmatrix} = \begin{bmatrix} w_{\alpha} \\ w_{\beta} \\ w_{\gamma} \end{bmatrix}, \quad u \in \mathcal{T}^{\text{lin}}(\bar{y}) \\ & \iff JG(\bar{x})_{\beta,:}v + s\mathbf{1}_{\beta} = w_{\beta}, \quad s \in \mathbb{R}, \end{aligned}$$

where the second equivalence holds due to the form of $\mathcal{T}^{\text{lin}}(\bar{y})$ because u_α and u_γ are free. Therefore, the constraint non-degeneracy condition (35) holds if and only if $[G'(\bar{x})_{\beta,:} \quad \mathbf{1}_\beta]$ has full row rank, which requires $|\beta| \leq n + 1$.

For comparison, consider the linear independence constraint qualification of the nonlinear reformulated problem in (3). The following conditions can be verified to hold at a solution $(x, t, z) \in \mathbb{R}^{n+1+m}$ with an active top- k -sum constraint (left) leading to the following active constraint matrix (right):

$$0 = t + \frac{1}{k} \mathbf{1}^\top z \quad (38.1)$$

$$z_\alpha = G(x)_\alpha - t \mathbf{1}_\alpha \quad z_\alpha > 0 \quad (38.2)$$

$$z_\beta = G(x)_\beta - t \mathbf{1}_\beta \quad z_\beta = 0 \quad (38.3)$$

$$z_\gamma > G(x)_\gamma - t \mathbf{1}_\gamma \quad z_\gamma = 0 \quad (38.4)$$

$$C := \begin{bmatrix} 0 & 1 & \frac{1}{k} \mathbf{1}^\top \\ -G'(x)_{\alpha,:} & \mathbf{1}_\alpha & I_{\alpha,:} \\ -G'(x)_{\beta,:} & \mathbf{1}_\beta & I_{\beta,:} \\ 0 & 0 & I_{\beta,:} \\ 0 & 0 & I_{\gamma,:} \end{bmatrix}. \quad (39)$$

It is clear that C has full row rank if and only if two conditions hold: (i) the submatrix $[JG(x)_{\beta,:} \quad \mathbf{1}_\beta]$ has full row rank, i.e., $\text{rank}[JG(x)_{\beta,:} \quad \mathbf{1}_\beta] = |\beta|$ and $|\beta| \leq n + 1$; and (ii) $|\beta| \leq n$. Condition (i) arises from observing that the second and fifth block rows of C are always linearly independent from each other and from the remaining block rows because the index sets α , β , and γ are distinct. Therefore, the only possibility of linear dependence is due to blocks three and four. Finally, given a full row rank matrix E , any block matrix $\begin{bmatrix} D & E \\ & E \end{bmatrix}$ has full row rank if and only if D has full row rank (Horn and Johnson, 2013). Thus we recover condition (i). Condition (ii) is obtained from the requirement that $1 + m + |\beta| \leq n + m + 1$, i.e., the number of rows of C must not exceed the number of columns. Therefore, necessary and sufficient conditions for satisfying the constraint non-degeneracy condition under the standard reformulation are slightly stronger than the top- k -sum operator formulation.

Appendix C. Reformulation of the quantile regression

Given input-output data $\{(a^i, b_i)\}_{i=1}^m \subseteq \mathbb{R}^n \times \mathbb{R}$ and some $\tau \in (0, 1)$ with integer $(1 - \tau)m$, consider the empirical quantile regression problem given by

$$\underset{x \in \mathbb{R}^n, x_0 \in \mathbb{R}}{\text{minimize}} \quad \frac{1}{m} \sum_{i=1}^m \rho_\tau(b_i - x^\top a^i - x_0), \quad (40)$$

where $z = z^+ - z^-$ with $z^+ = \max(z, 0)$ and $z^- = \max(-z, 0)$, and $\rho_\tau(z) = (1 - \tau) \cdot z^- + \tau \cdot z^+$ is the so called “check loss” of Koenker and Bassett, Jr. (1978). By simple manipulation, we see that

$$\rho_\tau(z) = (1 - \tau) \cdot z^- + \tau \cdot z^+ = (1 - \tau)z^- + \tau z^+ + (z^+ - z^+) = (1 - \tau) \cdot (-z) + z^+.$$

It is easy to see that one can scale the objective in (40) by $(1 - \tau)^{-1}$ and rewrite the problem as

$$\begin{aligned} & \underset{x \in \mathbb{R}^n, x_0 \in \mathbb{R}}{\text{minimize}} \quad \frac{1}{(1-\tau) \cdot m} \sum_{i=1}^m \max(0, b_i - x^\top a^i - x_0) - \frac{1}{m} \sum_{i=1}^m (b_i - x^\top a^i - x_0) \\ &= \underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \underset{x_0 \in \mathbb{R}}{\text{minimize}} \quad x_0 + \frac{1}{(1-\tau) \cdot m} \sum_{i=1}^m \max(0, b_i - x^\top a^i - x_0) - \frac{1}{m} \sum_{i=1}^m (b_i - x^\top a^i) \\ &= \underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \text{CVaR}_\tau(b_i - x^\top a^i) - \frac{1}{m} \sum_{i=1}^m (b_i - x^\top a^i) \\ &= \underset{x, t}{\text{minimize}} \quad t - \frac{1}{m} \sum_{i=1}^m (b_i - x^\top a^i) \quad \text{subject to} \quad \text{CVaR}_\tau(\{b_i - x^\top a^i\}_{i=1}^m) \leq t. \end{aligned}$$

By the translation equivariance property of the superquantile operator, we obtain formulation (31).