TRANSPORTATION SCIENCE

Target-Oriented Robust Inventory Management in Electric Vehicle Battery Swapping Networks

Journal:	*Transportation Science*
Manuscript ID	Draft
Manuscript Type:	Original Manuscript
Keywords/Area of Expertise:	Battery swapping, Distributionally robust optimization, Lateral transshipment

https://pubsonline.informs.org/journal/trsc

TRANSPORTATION SCIENCE

Vol. 00, No. 0, Xxxxx 0000, pp. 000–000

Submitted to *Transportation Science*

ISSN 0041-1655, EISSN 1526-5447

Target-Oriented Robust Inventory Management in Electric Vehicle Battery Swapping Networks

Abstract. Battery swapping for refueling has been embraced by major electric vehicle (EV) manufacturers. To facilitate the operations of battery swapping, this paper studies the management of a battery swapping network that operates in a "swap-locally, charge-centrally" way: batteries are swapped at local swapping stations and transported to a central charging station for recharging. It remains unclear how to replenish batteries from central charging stations, fulfill uncertain demand, and transship batteries among swapping stations to guarantee a high service level. In this paper, we exploit the network as a repairable inventory management problem with lateral transshipment. Combining the ideas from process flexibility, we first develop a stochastic dynamic program to characterize coupling network operations. Then we introduce a novel robust satisficing model from a target-oriented perspective to address the curse of dimensionality and distributional ambiguity. Specifically, we first employ a utility-based Service Level Measure by constraining the severity of inventory violations caused by uncertain demand. We solve this problem by leveraging an event-wise ambiguity set and a pooling-group-based enhanced linear decision rule. Out-of-sample tests demonstrate that our model outperforms several benchmarks in terms of robustness and target attainment. The major insights are twofold: (i) Incorporating lateral transshipment in the service network reduces substantial costs compared to a system without transshipment. Furthermore, a chaining network captures most of the benefits of a fully flexible network, and (ii) increasing the cost budget can enhance the robustness, whereas increasing the penalty cost reduces the probability of inventory violation by building more battery inventory.

Key words: Battery swapping, Adjustable robust optimization, Robust satisficing, Distributionally robust optimization, Lateral transshipment

1. Introduction

This research aims to address the battery inventory management problem with uncertain demand that arises in an EV battery swapping-charging network. In this problem, decision-makers (DMs) determine initial battery stocks and periodic network operations, including battery replenishment, fulfillment, and lateral transshipment. The problem is motivated by the wave of vehicle electrification, which holds great promise to reduce greenhouse gas emissions in the transport sector (accounting for around 30% globally). Over 95% of greenhouse gas emissions by the transport sector are attributable to oil consumption of internal combustion engines (IEA 2024). Correspondingly, major economies concentrate on EVs powered by batteries and electric motors to alleviate the dependence on oil: China, the forerunner in this field, incentivizes companies along the EV supply chain with subsidies to ramp up domestic production. The U.S. Inflation Reduction Act allocates USD $369 billion for climate investments to promote EV adoption (U.S. Government 2023).

However, the widespread adoption of EVs faces several barriers, primarily due to limitations in batteries: (i) *range anxiety* and (ii) *high costs*. The first upfront issue, *range anxiety*, refers to the fear of running out of battery before reaching the destination. This is primarily caused by the significantly longer recharging time of EVs compared to refueling internal combustion engine counterparts. Fully recharging an EV battery typically takes an entire night with a home charger, which is impractical for long-distance trips. Another key issue is the *high costs* of EV models. Currently, the total cost of owning EVs remains 10%-50% higher than that of ICE cars. For a light-duty vehicle, a 100 kWh lithium-ion battery pack costs USD$13,900, which curbs many potential adopters despite the low running cost of EVs. As a result, many companies are transitioning to an innovative business model of *battery swapping* to address these issues, i.e., establishing a network of swapping stations to refuel EVs with fully charged batteries.

Battery swapping provides an alternative to plug-in charging by replacing the depleted EV battery with a fully charged one. This fully automated process can be completed within three minutes, which largely reduces the waiting time compared to traditional plug-in charging. The unique benefits of battery swapping tackle the aforementioned issues: (i) *Fast service*: Range anxiety can be significantly reduced through completing a swap in under three minutes, far quicker than the 30 minutes required by Tesla's Supercharger. (ii) *High affordability*: Battery swapping reduces costs through bundling with a battery lease contract, under which the EV company retains ownership and management of the batteries. Consumers can subscribe for just $120 per month instead of paying $10,000 upfront (Asadi and Nurre Pinkley 2022).

Nevertheless, EV companies are now facing challenges in promoting battery swapping service mainly because: (i) *Inefficient power capacity*. Simultaneous fast charging at swapping stations creates high-peak power demands, which risks local grid congestion. (ii) *Heavy upfront costs*. Battery swapping requires considerable investment in pervasive associated infrastructures. Previous attempts by Tesla and Better Place to deploy battery charging-swapping networks failed due to heavy upfront investment costs (Mak, Rong, and Shen 2013).

To overcome these challenges, we are inspired by Qi, Zhang, and Zhang (2023) to consider a *swap-locally, charge-centrally* network where batteries are swapped at distributed swapping stations and transported to a concentrated charging station for recharging. Unlike traditional battery swapping systems that equip each swapping station with charging bays, the ”swap-locally, charge-centrally” network leverages the sufficient power capacity of a substation to reduce the pressure to upgrade local power grids. Moreover, swapping stations can be adapted from the existing gas stations without the need for expensive charging bays. Therefore, an accessible and affordable battery swapping network could be achieved. However, how to operate such a network to guarantee a high service level? How to replenish batteries from central charging stations, fulfill uncertain demand and transship batteries among swapping stations? These problems remain unresolved.

Our work provides models, solution approaches, and managerial implications for such a novel business model under demand uncertainty. We summarize our key contributions as follows:

- We present a novel problem relevant to the emerging EV industry by exploiting the “swap-locally, charge-centrally” network as a two-echelon repairable inventory model with lateral transshipment. We jointly determine coupling network decisions to facilitate battery operations: the initial battery stocks, replenishment, fulfillment, and lateral transshipment of batteries.
- We develop a stochastic dynamic programming model incorporating ideas from process flexibility. To tackle the curse of dimensionality and distributional ambiguity, we propose a target-oriented multi-stage *robust satisficing* (RS) model that allows an inherent consideration in risk mitigation across the entire probability space. We further define a utility-based Service Level Measure to assess inventory violations caused by uncertain demand. We also provide theoretical proofs for its salient properties.
- We derive a tractable linear program leveraging the lifted event-wise ambiguity set and a pooling-group-based enhanced linear decision rule. The result formulation under event-wise distributional information can be effectively solved by the state-of-the-art solvers.
- We perform numerical studies on synthetic data to demonstrate the robustness and efficiency of our proposed RS model. When evaluated on out-of-sample distributions, our model achieves target inventory levels with greater robustness than traditional distributional robust optimization (DRO) models. The major managerial insights are derived: (i) Applying lateral transshipment in a regular battery swapping-charging network ensures a minimum saving of USD$240,000 (5.26%) in total costs compared to a network without transshipment. (ii) Our proposed chaining structure network can preserve the advantages of a fully flexible network. These benefits are dimmed by

the increasing demand correlations among swapping stations. (iii) The overall risk of violating inventory targets decreases as the budget for operational costs increases, but escalates sharply as back-order penalty costs go down.

The rest of this paper is organized as follows: Section 2 reviews the related literature. Section 3 formally defines the problem as a stochastic dynamic program. In Section 4, we extend the model to an RS setting and incorporate our proposed utility-based SLM. Section 5 presents the proposed solution approach. Section 6 reports results from extensive numerical studies and discusses the derived managerial insights. Section 7 concludes with a summary. The appendix contains the proofs of statements, extensions, as well as additional results for the numerical studies.

2. Literature Review

This section presents a review of related literature pertinent to four main streams: (i) EV battery swapping operations, (ii) closed-loop repairable inventory management, (iii) lateral transshipment and process flexibility, and (iv) target-oriented DRO. We also discuss the research gap between the extant studies and our work.

2.1. EV Battery Swapping Operations

As battery swapping plays a crucial role in scaling up EV adoption, a growing research interest has arisen in operations management. Mak, Rong, and Shen (2013) were the first to investigate the swapping station model. They proposed a two-stage robust optimization model to support the deployment of battery-swapping infrastructure networks. Following this direction, Avci, Girotra, and Netessine (2015) presented a rigorous analytical comparison between battery swapping and plug-in charging models. They formulated the problem as a sequential game combining a repairable inventory model and a behavioral model of motorists. Both papers resolved strategy-level problems in the early planning stage of battery swapping adoption. Another direction is the operation-level problem in swapping station management. Schneider, Thonemann, and Klabjan (2018) considered an infinite time horizon of operation of battery swapping stations. They formulated a Markov Decision Process (MDP) and proposed an approximate dynamic programming (ADP) algorithm to solve it. In the same spirit, Asadi and Nurre Pinkley (2022) applied the ADP method to model the operations at a battery swapping station to decide stochastic scheduling, allocation, and inventory replenishment problems. Xie, Dai, and Pei (2024) proposed a stationary queuing network model to characterize the operations of the e-bike battery swapping system.

The aforementioned research overlooked the ”swap-locally, charge-centrally” service network, a setting that has not been thoroughly explored in the existing literature. To the best of our knowledge, Qi, Zhang, and Zhang (2023) was among the first to rigorously analyze the deployment and management of the ”swap-locally, charge-centrally” service network. They proposed a joint location and repairable inventory model to capture non-Poisson demand arrivals and stochastic charging time. However, Qi, Zhang, and Zhang (2023) did not consider lateral transshipments among charging or swapping stations. Our work extended their work to incorporate lateral transshipment and provided computationally efficient solutions under the RS framework.

2.2. Closed-loop Repairable Inventory Management

The battery swapping-charging network operation is closely relevant to closed-loop repairable inventory management, where depleted batteries function as failed parts and recharging is considered the repairing process (Qi, Zhang, and Zhang 2023). The key distinction is that the central charging station serves as both a repair factory and a base in our context. The cornerstone of repairable inventory management, both in theory and practice, is the Multi-Echelon Technique for Recoverable Item Control (METRIC) model developed by Sherbrooke (1968). The METRIC system consists of one depot and multiple bases, where bases replace the failed items with on-hand stocks. Then failed items are returned to the depot for repair. Simultaneously, the base orders stock replenishment from the depot. The METRIC model has been extensively applied in the military, e.g., aircraft components, and in the commercial area (Guide Jr and Srivastava 1997). Mak, Rong, and Shen (2013) and Avci, Girotra, and Netessine (2015) first introduced this classic METRIC model into the novel context of battery swapping.

Traditional repairable inventory theory focuses on providing exact analytical solutions, but these solutions depend on restrictive assumptions such as the demand arrival process follows a Poisson distribution with a fixed mean (Lee 1987). Although they offer valuable insights, these assumptions are too strict and incompatible with real-world uncertain demands. Qi, Zhang, and Zhang (2023) extended the batched order repairable inventory model by relaxing the above assumptions. Additionally, they provided nuanced structural properties of models through analytical analysis. However, Qi, Zhang, and Zhang (2023) still failed to consider demand ambiguity and flexible ordering quantities. Therefore, our work is distinct from this previous literature by employing a multi-stage robust model to consider demand ambiguity.

2.3. Lateral Transshipment and Process Flexibility

Inspired by Schneider, Thonemann, and Klabjan (2018), we introduce lateral transshipment into the ”swap-locally, charge-centrally” network. This process of reallocating batteries between stations is rooted in Lee (1987), which derived the expected backorder level and the number of lateral transshipments in a multi-echelon repairable inventory system. Extensive settings such as periodic and continuous review have already been resolved (Olsson 2015). Recently, there has been growing interest in investigating lateral transshipment (inventory repositioning) in other settings, including shared mobility and container reallocating in freight logistics; as seen in He, Hu, and Zhang (2020), Benjaafar et al. (2022), and the references therein. Unlike shared mobility and container reallocation, our work focuses on a setting where resources are repairable and reusable.

Process flexibility refers to the flexibility for manufacturers to promptly switch between different products using the same production line (Jordan and Graves 1995). At a broader level, both process flexibility and lateral transshipment fall under resource pooling (Cui et al. 2023a). Recently, there have been several efforts to incorporate process flexibility into the design of last-mile delivery operations, vehicle routing problems, and online retailing (Lyu et al. 2019, Ledvina et al. 2022, DeValve et al. 2023). However, these studies did not consider the multi-period setting, nor did they incorporate distributional robust approaches. Our work sheds light on the consideration of process flexibility within the battery swap process.

2.4. Target-Oriented DRO Approach

The last mainstream study related to our study is the target-oriented DRO. According to Simon (1955), major DMs often tend to achieve a specified target rather than maximize utility, a philosophy known as target *satisficing*, which combines "satisfy" and "suffice". Since Simon (1955)'s work, target-oriented optimization has seen considerable advancements (Föllmer and Schied 2002, Brown and Sim 2009). Brown and Sim (2009) axiomatized the decision criterion *satisficing measure*, which rewards diversification and aligns with risk-aversion behavior. The satisficing measure is prevalent in optimization involving probability ambiguity, such as surgery allocation (Chow, Cui, and Long 2022), humanitarian operations (Avishan et al. 2023), and inventory routing problems (Cui et al. 2023b). Given the difficulty in estimating demand distribution before the network deployment, employing satisficing measures is a natural approach for operational problems in battery swapping networks.

All the aforementioned works employ the DRO framework, which has gained popularity in the past decade (Mohajerin Esfahani and Kuhn 2018, Bertsimas, Sim, and Zhang 2019). DRO addresses uncertainty using an ambiguity set to restrict the distribution to within the vicinity of the empirical distribution. Instead, our work uses a novel RS framework proposed by Long, Sim, and Zhou (2023). Unlike DRO, which specifies an ambiguity set of probability distributions, the

RS framework puts no restriction on the distributions and systematically mitigates the risks of not achieving targets. Under the RS framework, Cui et al. (2023a) formulated a two-stage resource pooling problem using DRO from a service perspective. In this paper, we first introduce the RS framework to decide coupling adaptive operations in the "swap-locally, charge-centrally" network.

Notation. We use brackets $[\cdot]$ to represent running indices, for example, $[N]$ denotes $\{1, 2, \ldots, N\}$ ($[0] = \emptyset$). $|\cdot|$ denotes the number of elements in the set. The boldface lowercase letters represent column vectors, such as $\boldsymbol{x} = (x_1, x_2, \ldots, x_N)$. Operator $(\cdot)^+$ is defined as $(x)^+ = \max\{x, 0\}$. $\mathbf{1}(\mathbf{0})$ denotes a vector with all elements as one (zero), and $\boldsymbol{e}_i$ denotes i-th basic vector. The uncertain quantities are denoted with the symbol "~", e.g., $\tilde{\boldsymbol{x}}$ is an uncertain vector, with $\boldsymbol{x}$ being its realization. For a random variable $\tilde{x}$ with $\tilde{x} \sim \mathbb{P}, \mathbb{P} \in \mathcal{P}_0$, $\mathcal{P}_0$ represent the set of all possible distributions on $\mathbb{R}$. For a multi-variant random variable $\tilde{\boldsymbol{x}}$ with support $\mathcal{Z}$, we use $\mathcal{P}_0(\mathcal{Z})$ to denote its space of distributions. Finally, we denote $\mathcal{R}^{N,M}$ the space of all measurable functions from $\mathbb{R}^N$ to $\mathbb{R}^M$ that are bounded on compact sets.

3. Battery Swapping Network Operations

In Section 3.1, we model the operations of a "swap-locally, charge-centrally" network as an MDP and formulate the associated stochastic dynamic programming model. Section 3.2 introduces process flexibility to design pooling groups for lateral transshipment a priori. We then incorporate these pooling groups into a multi-stage DRO model to address demand ambiguity.

3.1. Problem Statement

We consider a two-echelon "swap-locally, charge-centrally" network consisting of one central charging station and N local swapping stations, as illustrated in Figure 1. Local swapping stations handle uncertain EV arrivals and exchange depleted batteries from customers for fully charged batteries. Batteries are centrally charged at the charging station 0, and once fully charged, they are distributed to each swapping station $i \in [N]$. The sequence of events in each period is as follows:

Before the planning horizon, we should consider the initial battery investment at each station, i.e., x_0 for charging station 0 and x_i^0 for swapping station $i \in [N]$. Each battery requires a unit investment of K. *At the beginning* of each period $t \in [T]$, the swapping station $i \in [N]$ holds an on-hand stock of x_i^t fully charged batteries to meet customers' demand. Let x_0^t denote the fully charged batteries reserved in the central charging station 0. We further denote $\boldsymbol{x}^t = (x_0^t, x_1^t, \ldots, x_N^t)$ the vector of inventory levels at all stations. Prior to the EV arrivals, the central planner decides on the replenishment quantity $\boldsymbol{q}^t = (q_1^t, \ldots, q_N^t)$ for swapping stations from the charging station 0,

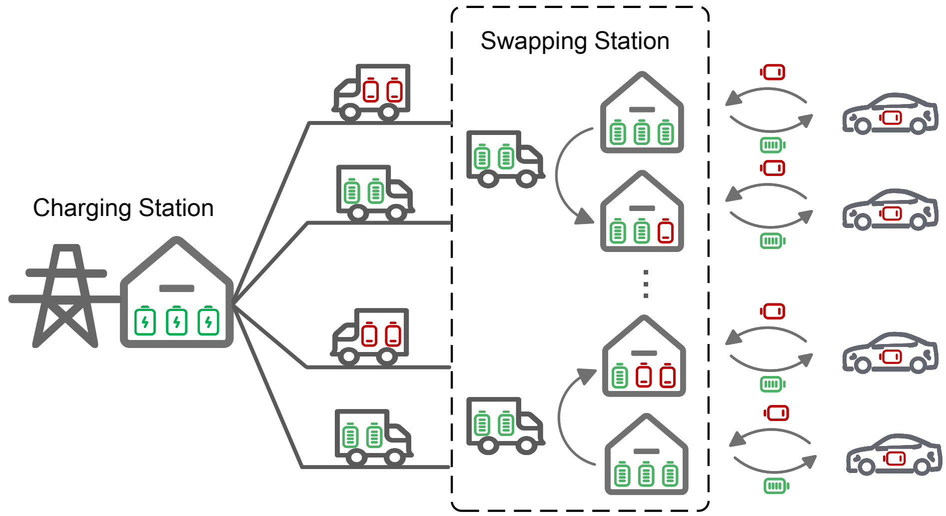

Figure 1 **"Swap-Locally, Charge-Centrally" Network**

with a unit ordering cost $c_t > 0$ and the transportation time $L_{0i} > 0$. Then the order $\boldsymbol{q}^t$ in period t will arrive at $t + L_{0i}$. Besides order quantity, the planner determines the number of transshipments $\boldsymbol{y}^t = (y_{ij}^t)$ for all $i, j \in [N]$, where y_{ij}^t fully charged batteries are transshipped from the swapping station i to j in period t. Let $L_{ij} \geq 0$ denote the transportation time from i to j and $L_{ij} = L_{ji}$. The transshipment order $\boldsymbol{y}^t$ will arrive at the beginning of period $t + L_{ij}$. Battery transshipment incurs a unit cost of p_{ij}^t, and the costs depend on the transportation distance $L_{ij} \geq 0$ ($L_{ii} = 0$). When all swapping stations are out of stock, the swapping station will wait for regular replenishment from the central charging station.

During the period t, there are a total of $\tilde{d}_i^t$ uncertain EV arrivals at the station i. Let the random vector $\tilde{\boldsymbol{d}}^t = (\tilde{d}_1^t, \tilde{d}_2^t, \ldots, \tilde{d}_n^t)$ denote demands of all swapping stations in period t; correspondingly, $\tilde{\boldsymbol{d}}^{[t]} = (\tilde{\boldsymbol{d}}^1, \tilde{\boldsymbol{d}}^2, \ldots, \tilde{\boldsymbol{d}}^t)$ represents demands from period 1 to period t ($\tilde{\boldsymbol{d}}^{[0]} = \emptyset$). The demand process $(\boldsymbol{d}^t)$ is assumed to be independent over horizons, and it follows a joint probability distribution $\mathbb{P}$. When an EV arrives at the swapping station i, the station prioritizes using its stock on hand to fulfill the demand. The fulfillment quantity is denoted by $\boldsymbol{u}^t = (u_1^t, \ldots, u_N^t)$. Any excess demand $(d_i^t - u_i^t)^+$ is assumed to be lost sales with a unit penalty cost $b > 0$. We exclude holding costs for batteries because capital costs are negligible compared to lost-sales penalties (Schneider, Thonemann, and Klabjan 2018).

At the end of period t (i.e., *at the beginning of* period $t + 1$), the swapping station i will dispatch a vehicle to transfer u_i^t depleted batteries to charging station 0 for recharging. As defined in Qi, Zhang, and Zhang (2023), we use $\Delta_i > 0$ to represent the lead time of batteries, including charging and

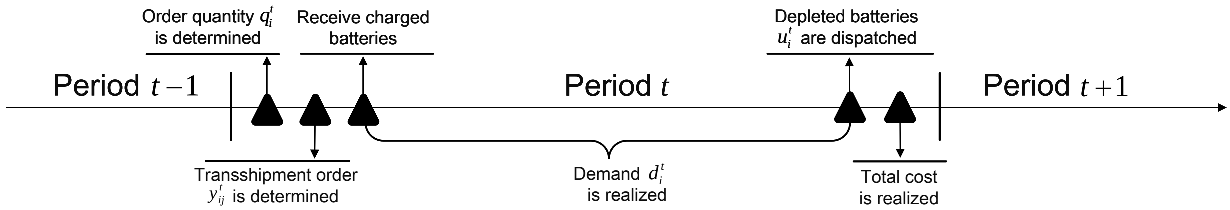

Figure 2 **Sequence of Events at A Swapping Station**

transportation time. Therefore, the depleted batteries sent to the charging station 0 in period $t - \Delta_i$ will be available *at the beginning* of period t. We assume all the transportation can be accomplished while satisfying time windows and capacity constraints. The sequence of events is illustrated in Figure 2.

This problem is formulated as a stochastic dynamic program with a planning horizon of T periods. Given a period $t \in [T]$, the system state includes two parts. The first part is on-hand inventory $\boldsymbol{x}^t$. The second part is historical information: (i) demand realizations $\boldsymbol{d}^{[t-1]}$; (ii) historical replenishment orders $\boldsymbol{q}^{[t-1]}$; (iii) historical transshipment orders $\boldsymbol{y}^{[t-1]}$ and (iv) historical charging batteries: $\boldsymbol{u}^{[t-1]}$. After observing the system state $\boldsymbol{s}_t = \left(\boldsymbol{x}^t, \boldsymbol{d}^{[t-1]}, \boldsymbol{q}^{[t-1]}, \boldsymbol{y}^{[t-1]}, \boldsymbol{u}^{[t-1]}\right)$, the planner makes an action $\boldsymbol{a}_t = (\boldsymbol{q}^t, \boldsymbol{y}^t, \boldsymbol{u}^t)$ in response. The action space satisfies the following constraints:

- The total allocated charged batteries $\sum_{i\in[N]} q_i^t$ cannot exceed available batteries at the central charging station x_0^t. The second constraint enforces that the replenishment quantity is always nonnegative:

$$\begin{aligned} &\sum_{i\in[N]} q_i^t \le x_0^t + \sum_{i\in[N]} u_i^{t-\Delta_i} \quad && \forall t \in [T], i \in [N] \\ &q_i^t \ge 0 && \forall t \in [T], i \in [N] \end{aligned} \tag{1}$$

- The total transshipment $\sum_{j\in[N]} y_{ij}^t$ sent from the swapping station i cannot surpass the on-hand stock x_i^t. The transshipment amount is always nonnegative:

$$\begin{aligned} &\sum_{j\in[N]} y_{ij}^t \le x_i^t \quad && \forall t \in [T], i \in [N] \\ &y_{ij}^t \ge 0 && \forall t \in [T], i, j \in [N] \end{aligned} \tag{2}$$

- Fulfillment decision u_i^t is either the current demand d_i^t or the updated inventory level, depending on which one is smaller. To facilitate the discussion, we introduce an auxiliary variable $\bar{x}_i^t$ to denote the updated inventory level:

$$\bar{x}_i^t = x_i^t + q_i^{t-L_{0i}} + \sum_{j\in[N]} y_{ji}^{t-L_{ij}} - \sum_{j\in[N]} y_{ij}^t$$

Moreover, the fulfillment quantity is always nonnegative. Therefore, fulfillment decision u_i^t satisfies:

$$u_i^t = \min\{d_i^t, \bar{x}_i^t\} \quad \forall t \in [T], i \in [N]$$
$$u_i^t \geq 0 \quad \forall t \in [T], i \in [N]$$

The above fulfillment constraint is nonlinear and induces a computational burden. We need to reformulate it into linear constraints that are easier to handle. For this purpose, the following assumptions are proposed for the cost parameters c_t, p_{ij}^t, b:

ASSUMPTION 1. *The cost parameters satisfy $b > c_{\max} + p_{\max}$, where $c_{\max} = \max_{t\in[T]} c_t$ and $p_{\max} = \max_{i,j,t} p_{ji}^t$.*

This assumption implies that the profit of satisfying demand in the current period is greater than the costs of ordering and transshipping the battery from any other station $i \in [N]$. Similar assumptions on cost parameters have been made in the transshipment literature to ensure the convexity of the cost function; see He, Hu, and Zhang (2020) and Benjaafar et al. (2022). Benjaafar et al. (2022) stated that "the assumption prevents the unpleasant situation where one might want to 'hide' the inventory due to the difference in the repositioning cost." Under Assumption 1, we can linearize the fulfillment constraints by the following:

$$\begin{aligned} u_i^t &\leq \bar{x}_i^t \quad \forall t \in [T], i \in [N] \\ u_i^t &\leq d_i^t \quad \forall t \in [T], i \in [N] \end{aligned} \tag{3}$$

For brevity, we denote the action space by $\mathcal{X}(\boldsymbol{s}_t) = \{\boldsymbol{a}_t \mid (1), (2), (3)\}$. Given state $\boldsymbol{s}_t$, the action space $\mathcal{X}(\boldsymbol{s}_t)$ is a polyhedron. The transition probabilities are induced by the system state update function for any $i \in [N]$ and $t \in [T]$:

$$x_i^{t+1} = x_i^t + \underbrace{q_i^{t-L_{0i}} + \sum_{j\in[N]} y_{ji}^{t-L_{ij}}}_{\text{Battery income}} - \underbrace{\sum_{j\in[N]} y_{ij}^t}_{\text{Battery expenditure}} - u_i^t \tag{4}$$

Observe that $x_i^{t+1} = \bar{x}_i^t - u_i^t$ following the definition. Together with the fulfillment constraint (3), we can replace the system update function (4) equivalently with $x_i^{t+1} \geq 0$. In the same spirit, the inventory dynamics for charging station 0 in period t is:

$$x_0^{t+1} = x_0^t + \underbrace{\sum_{i\in[N]} u_i^{t-\Delta_i}}_{\text{Battery income}} - \underbrace{\sum_{i\in[N]} q_i^t}_{\text{Battery expenditure}} \tag{5}$$

Together with the replenishment constraint (1), we equivalently have $x_0^{t+1} \geq 0$. Historical information $(\boldsymbol{d}^{[t]}, \boldsymbol{q}^{[t]}, \boldsymbol{y}^{[t]}, \boldsymbol{u}^{[t]})$ is updated by simple concatenation. In resolving the initial boundary conditions $\boldsymbol{s}_0$, we enforce that there is no outstanding order in transit or fulfillment order before the planning horizon, e.g.,

$$\boldsymbol{q}^t = \mathbf{0},\ \boldsymbol{y}^t = \mathbf{0},\ \boldsymbol{u}^t = \mathbf{0} \quad \forall t \leq 0$$

where all zero vectors are in the right dimensions.

Given an action $\boldsymbol{a}^t$ in period t, the single-period operational cost $C_t(\boldsymbol{a}^t)$ is the sum of order cost, transportation cost, and backorder penalty:

$$C_t(\boldsymbol{a}^t) = \sum_{i \in [N]} \left(c_t q_i^t + \sum_{j \in [N], j \neq i} p_{ij}^t y_{ij}^t + b(d_i^t - u_i^t) \right)$$

Additionally, since each battery incurs a unit cost of K, the total initial battery investment is given by $K\mathbf{1}^\top \boldsymbol{x}$. Our goal is to minimize the initial investment and the expected operational costs:

$$Z_{\text{STOC}} = \min_{\boldsymbol{x}^{\mathbf{0}} \geq \mathbf{0}} \left\{ K\mathbf{1}^\top \boldsymbol{x} + \min_{(\boldsymbol{a}^t)_{t \in [T]}} \mathbb{E}_{\mathbb{P}}^{\boldsymbol{x}^{\mathbf{0}}} \left[\sum_{t \in [T]} C_t(\boldsymbol{a}^t) \right] \right\} \tag{$\mathcal{STOC}$}$$

where $\mathbb{P}$ represents the joint probability distribution of the demand process, and the transfer probability is induced by inventory dynamics (4) and (5).

We can recast the Problem ($\mathcal{STOC}$) to a stochastic dynamic program. In the case of initial inventory $\boldsymbol{x}^{\mathbf{0}}$, the optimality equation for each period $t \in [T]$ is defined as follows:

$$\begin{aligned} V_t(\boldsymbol{s}_t, \boldsymbol{x}^0) = \min_{\boldsymbol{a}^t \in \mathcal{X}(\boldsymbol{s}_t)} & \left\{ \sum_{i \in [N]} \left(c_t q_i^t + \sum_{j \in [N], j \neq i} p_{ij}^t y_{ij}^t \right) \right. \\ & \left. + \mathbb{E}_{\boldsymbol{d}^t} \left[\sum_{i \in [N]} b(d_i^t - u_i^t) + V_{t+1}(\boldsymbol{s}_{t+1}, \boldsymbol{x}^0) \mid \boldsymbol{s}_t \right] \right\} \end{aligned} \tag{6}$$

where $V_t(\boldsymbol{s}_t, \boldsymbol{x}^0)$ stands for the optimal total expected costs from period t until the last period T with the initial battery investment $\boldsymbol{x}^0$. The terminal cost is given by $V_{T+1}(\boldsymbol{s}_{T+1}) = 0$. The expectation in (6) is taken over the conditional probability of $\boldsymbol{d}^t$ given $\boldsymbol{s}_t$. Hence, we can use $V_0(\boldsymbol{s}_0, \boldsymbol{x}^0)$ to denote the optimal cost over T periods. Problem ($\mathcal{STOC}$) can be rewritten as:

$$Z_{\text{STOC}} = \min_{\boldsymbol{x}^0 \geq \mathbf{0}} \left\{ K\mathbf{1}^\top \boldsymbol{x} + V_0(\boldsymbol{s}_0, \boldsymbol{x}^0) \right\} \tag{$\mathcal{DP}$}$$

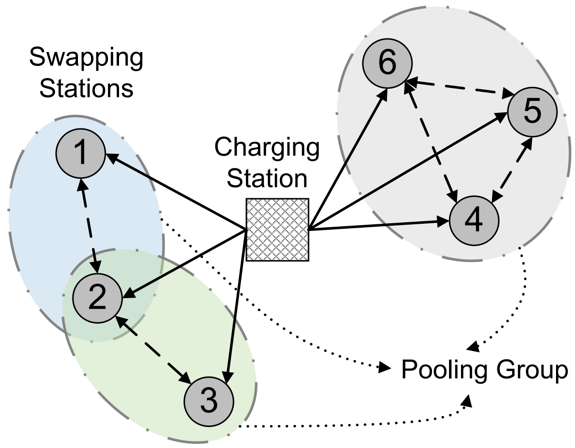

Figure 3 Three Pooling Groups with 1 Charging Station and 6 Swapping Stations.

Note. The solid arrow shows regular replenishment from the central charging station; the dashed arrow indicates potential lateral transshipment. Colored shadow regions represent different pooling groups.

The stochastic dynamic programming model ($\mathcal{DP}$) is confronted with the following challenges. First, it is computationally prohibitive due to high-dimensional state and action variables. In particular, the state variable $\boldsymbol{s}^t$ in a period is of $N + (N+1)^2(t-1)$ dimensions. Maintaining the value function $V_0(\boldsymbol{s}_0, \boldsymbol{x}^0)$ induces the notorious "curse of dimensionality". The second challenge is the inaccessibility of the true joint distribution $\mathbb{P}$ of uncertain demands $\tilde{\boldsymbol{d}}$. Although DMs obtain distributional knowledge from past demand samples (e.g., means and covariances), the real-world demand distribution is generally unobservable. The estimated distribution from a data set may perform poorly in out-of-sample data, which is called "optimizer's curse" (Smith and Winkler 2006). In reality, high-dimensional distribution poses difficulty in obtaining sufficient data and accurate estimation (Perakis et al. 2023). Subsequently, we first introduce process flexibility to simplify lateral transshipment, and then present the multi-stage DRO model to overcome computational challenges and distributional ambiguity.

3.2. Multi-stage DRO Model with Pooling Groups

Followed by process flexibility, we define flexibility as the ability to satisfy the demands of different locations by a single swapping station (Jordan and Graves 1995). Inspired by Lee (1987), we divide N swapping stations into *pooling groups* $\mathcal{N}_h$ for $h \in [H]$ to pool the demand within each group. Each swapping station is assigned to at least one pooling group. Figure 3 gives an illustration example of pooling groups. The ideal transshipment strategy allows any two swapping stations to transship inventory. Such a strategy features full flexibility to match battery supply and uncertain

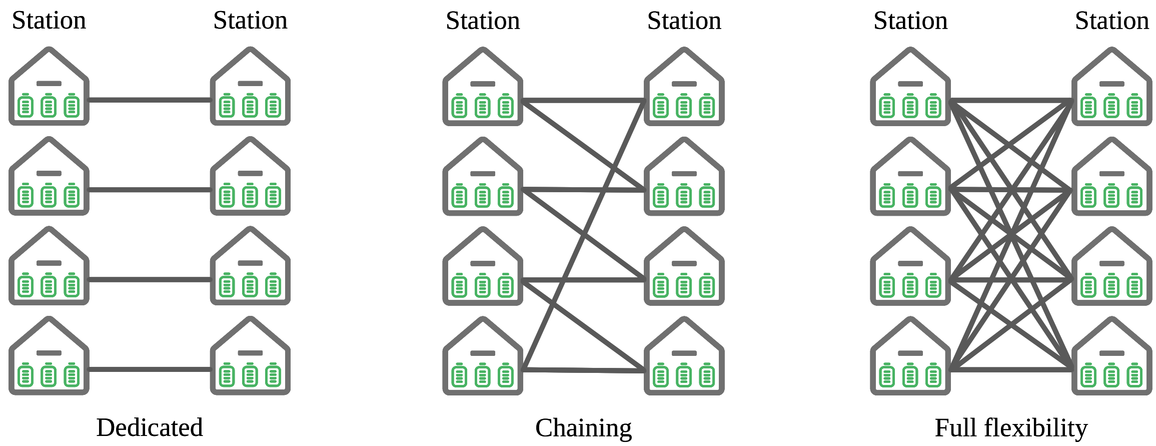

Figure 4 Flexibility Structures

Note. (i) Dedicated structure $\mathcal{D}$: Each swapping station i can only fulfill its own demand. (ii) Chaining (2-chain) structure $\mathcal{D}$: swapping station i is permitted to share inventory with the neighboring station $i+1$. (iii) Full flexibility structure $\mathcal{F}$: any two swapping stations that can share inventory through lateral transshipment.

demand. Despite its captivating strengths, full flexibility involves large dimensions of decision variables in each period (e.g., N^2 variables for N stations). Moreover, operating a fully flexible system can be prohibitively expensive. Therefore, a natural problem arises: *How can we implement partial flexibility while maintaining the most performance of full flexibility*?

From the process flexibility literature, we introduce three important flexibility networks $\mathcal{G} \in \{\mathcal{D}, \mathcal{C}, \mathcal{F}\}$, see Figure 4. Each bipartite flexibility network corresponds to a partition of pooling groups, where the resource nodes are swapping stations $i \in [N]$ and demand nodes are customers $j \in [N]$. An arc (i, j) connecting node i to node j implies that station i can use its extra inventory x_i to fulfill the demand d_j that has not been satisfied by station j.

It is useful to specify the neighborhood set of a swapping station i. Given a flexibility structure $\mathcal{G} \in \{\mathcal{D}, \mathcal{C}, \mathcal{F}\}$, we define the neighborhood set $\Gamma(i)$ as $\Gamma(i) = \{j \mid (i, j) \in \mathcal{G}\}$. In Figure 3, $\Gamma(1) = \{2\}$ and $\Gamma(2) = \{1, 3\}$; the corresponding pooling groups are $\mathcal{N}_1 = \{1, 2\}$ and $\mathcal{N}_2 = \{2, 3\}$. Under a defined flexibility $\mathcal{G}$ and neighborhood sets, the range of transshipment (2) can be restricted as follows:

$$\sum_{j \in \Gamma(i)} y_{ij}^t \leq x_i^t \quad \forall t \in [T], i \in [N]$$

$$y_{ij}^t \geq 0 \quad \forall t \in [T], (i, j) \in \mathcal{G}$$

Similarly, the above constraints (1)-(5) involving transshipment decisions y_{ij}^t for $j \in [N]$ can be adjusted to y_{ij}^t for $j \in \Gamma(i)$.

The main justification for introducing flexibility structures is to lower the decision dimensions. Under the chaining structure, we have $O(N)$ transshipment variables in a single period compared to $O(N^2)$ transshipment variables in the full flexibility network. However, we do admit there is a loss of optimality. Due to the restrictions on transshipment, there might exist a situation where the shortest path between two stations does not permit a transfer. The next result states the relations between optimal values of flexibility structures.

PROPOSITION 1 (**Performance Comparison of Flexibility Structures**). *Let $Z^*(\mathcal{G})$ denote the optimal objective value of Problem $(\mathcal{STOC})$ that can be achieved by a flexibility structure $\mathcal{G} \in \{\mathcal{D}, \mathcal{C}, \mathcal{F}\}$, we have: $Z^*(\mathcal{D}) \geq Z^*(\mathcal{C}) \geq Z^*(\mathcal{F})$.*

Proof See Online Appendix A.1. □

The result is straightforward since more flexibility indicates a larger solution space. Indeed, it is noted that as we add more flexibility, we can obtain more benefits (Simchi-Levi and Wei 2012). As we will later expound in Section 6, the chaining structure captures the most advantage of a full flexibility network, which underscores the effectiveness of pooling groups.

To address the computational burdensome and distributional ambiguity of Problem $(\mathcal{STOC})$, we next develop a novel multi-stage DRO model. Instead of assuming true demand distribution $\mathbb{P}$ is known, the DRO model assumes that the empirical distribution $\hat{\mathbb{P}}$ belongs to a certain ambiguity set $\hat{\mathcal{P}} \subset \mathcal{P}_0(\mathbb{R}^{NT})$. $\mathcal{P}_0$ is the entire distribution space of NT-dimensional random demand $\tilde{\boldsymbol{d}}$. The ambiguity set $\hat{\mathcal{P}}$ is characterized by partial distribution information estimated from historical data; see details in Bertsimas, Sim, and Zhang (2019). As a result, we seek to minimize the expected cost under the worst-case distribution within the ambiguity set, i.e., $\hat{\mathbb{P}} \in \hat{\mathcal{P}}$:

$$
\begin{aligned}
Z_{\mathrm{DRO}} = \min_{\boldsymbol{x},\boldsymbol{q},\boldsymbol{y},\boldsymbol{u} \geq \boldsymbol{0}} \quad & K\mathbf{1}^\top \boldsymbol{x} + \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}} \left[\sum_{t \in [T]} C_t(\boldsymbol{a}^t) \right] \\
\text{s.t.} \quad & \text{Constraints } (1) - (5)
\end{aligned}
\tag{7}
$$

Note that the objective function Z_{DRO} depends on the underlying flexibility structure $\mathcal{G}$. When the context is clear, we omit this dependence for simplicity. Unlike Problem $(\mathcal{STOC})$, which requires solving the problem sequentially, Problem (7) provides all the solutions concurrently against distributional ambiguity. However, there exist several limitations of the DRO model (7): First, the objective function of minimizing the expected cost does not account for the DM's attitude towards risk, nor does it consider service level. Second, DRO does not have a performance guarantee when the true distribution deviates from the ambiguity set, namely, $\mathbb{P} \in \mathcal{P}_0 \backslash \hat{\mathcal{P}}$. To obtain a fully robust program mitigating risk on the whole probability space, we introduce the framework of RS proposed by Long, Sim, and Zhou (2023) in the next section.

4. Target-oriented Robust Satisficing Model

This section first extends our multi-stage DRO model to the RS model, and then proposes a utility-based Service Level Measure to control the degree of inventory violation over the entire probability space. Finally, we demonstrate how to incorporate it into the RS model to obtain a tractable model.

4.1. Robust Satisficing

It is of vital importance to maintain on-hand inventory within a required window to provide high-level accessibility. Motivated by Cui et al. (2023b), we anticipate the period t inventory level x_i^t at each swapping station falling into a pre-specified interval $[\underline{\tau}_i^t, \overline{\tau}_i^t]$. For instance, $[\underline{\tau}_i^t, \overline{\tau}_i^t] = [0, 5]$ indicates that DMs are strongly averse to backorder as well as limiting the excess inventory to less than 5 units. Let $\tilde{\boldsymbol{x}} \in \mathcal{V}^N$ denote an N-dimensional random vector. The double-sided target violation of $\tilde{\boldsymbol{x}}$ associated with the required inventory interval $[\underline{\boldsymbol{\tau}}, \overline{\boldsymbol{\tau}}]$ is defined by the function:

$$
\begin{aligned}
v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) &= \max\{\tilde{\boldsymbol{x}} - \overline{\boldsymbol{\tau}}, \underline{\boldsymbol{\tau}} - \tilde{\boldsymbol{x}}\} \\
&= \max(\{\tilde{x}_i^t - \overline{\tau}_i, \underline{\tau}_i^t - \tilde{x}_i^t\})_{t\in[T], i\in[N]}
\end{aligned}
$$

Note that the negative violation indicates $\tilde{\boldsymbol{x}}$ falls within the interval, which indicates no target violation. Instead of limiting the distribution in our reference ambiguity set $\hat{\mathbb{P}} \in \hat{\mathcal{P}}_0$, we measure the system's worst-case threshold violation over all possible distributions $\mathbb{P} \in \mathcal{P}_0$. For this purpose, we adopt the RS framework by Long, Sim, and Zhou (2023) as follows:

$$
\begin{aligned}
\min_{k\geq 0} \quad & k \\
\text{s.t.} \quad & \text{Constraints (1)} - \text{(5)} \\
& \mathbb{E}_{\mathbb{P}}\left[v_{\underline{\tau}_i,\overline{\tau}_i}(x_i^t)\right] \leq k\Delta(\mathbb{P}, \hat{\mathbb{P}}) \qquad\qquad (\mathcal{RS}) \\
& \qquad \forall t \in [T], i \in [N], \mathbb{P} \in \mathcal{P}_0, \hat{\mathbb{P}} \in \hat{\mathcal{P}} \\
& K\mathbf{1}^\top \boldsymbol{x} + \sup_{\hat{\mathbb{P}}\in\hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[\sum_{t\in[T]} C_t(\boldsymbol{a}^t)\right] \leq C
\end{aligned}
$$

Constraints (1)-(5) are retained from the DRO model (7). The second constraint imposes that the worst-case expected loss $\sup_{\hat{\mathbb{P}}\in\hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}[\tilde{v}]$ is less than a scaled probability distance $k\Delta(\mathbb{P}, \hat{\mathbb{P}})$. The third constraint enforces a predetermined budget C on the worst-case expected total costs.

The magnitude of the violation $k \geq 0$ is measured by the statistical distance between a reference distribution $\hat{\mathbb{P}}$ in the ambiguity set $\hat{\mathcal{P}}$ and all possible distributions $\mathbb{P} \in \mathcal{P}_0$. A larger k will incur a larger expected target violation when evaluated on an unobservable distribution. Hence, intuitively, k represents the model's ability to avoid excessive deviation from the target interval $[\underline{\boldsymbol{\tau}}, \overline{\boldsymbol{\tau}}]$. Long, Sim, and Zhou (2023) used the empirical distribution as the reference distribution $\hat{\mathbb{P}}$. We extend a single empirical distribution to a reference ambiguity set $\hat{\mathcal{P}}$ to enhance distributional robustness. We configure the budget by solving the DRO counterpart of Problem $(\mathcal{RS})$ (He, Hu, and Zhang 2020, Long, Sim, and Zhou 2023):

$$
\begin{aligned}
C_{\text{DRO}} = \min_{\boldsymbol{x},\boldsymbol{q},\boldsymbol{y},\boldsymbol{u} \geq \boldsymbol{0}} \quad & K\mathbf{1}^\top \boldsymbol{x} + \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}} \left[\sum_{t \in [T]} C_t(\boldsymbol{a}^t) \right] \\
\text{s.t.} \quad & \text{Constraints (1)} - \text{(5)} \qquad\qquad (\mathcal{DRO}) \\
& \mathbb{E}_{\hat{\mathbb{P}}} \left[v_{\underline{\tau}_i, \overline{\tau}_i}(x_i^t) \right] \leq 0 \quad \forall t \in [T], i \in [N], \hat{\mathbb{P}} \in \hat{\mathcal{P}}
\end{aligned}
$$

Compared to our proposed $(\mathcal{RS})$ that minimizes k, $(\mathcal{DRO})$ minimizes the worst-case expected total costs. Compared to Problem (7), $(\mathcal{DRO})$ additionally ensures that the worst-case expected inventory violation does not exceed 0 within the reference ambiguity set $\mathcal{P}$, namely $\mathbb{E}_{\hat{\mathbb{P}}} \left[v_{\underline{\tau}_i, \overline{\tau}_i}(x_i^t) \right] \leq 0$ for any $\hat{\mathbb{P}} \in \hat{\mathcal{P}}$. Model $(\mathcal{DRO})$ recovers the empirical stochastic model $(\mathcal{STOC})$ as a special case when $\hat{\mathcal{P}} = \{\hat{\mathbb{P}}\}$, that is, the ambiguity set is a singleton. We present the relationship between $(\mathcal{RS})$ and $(\mathcal{DRO})$ in Proposition 2.

PROPOSITION 2. *Suppose the model $(\mathcal{DRO})$ is feasible, then the corresponding RS model $(\mathcal{RS})$ is also feasible for all $C \geq C_{DRO}$. In particular, if $C = C_{DRO}$, the optimal solution to the Problem $(\mathcal{RS})$ is also optimal to the Problem $(\mathcal{DRO})$. Moreover, if the Problem $(\mathcal{RS})$ is rewritten as a function of budget C, i.e., $\rho(C)$, then we have $\rho(C_1) \leq \rho(C_2)$ for any $C_{DRO} \leq C_1 \leq C_2$.*

Proof See Online Appendix A.2. □

Proposition 2 implies that $(\mathcal{RS})$ encapsulates $(\mathcal{DRO})$ as a special case when $C = C_{\text{DRO}}$. Observe that the RS constraint reduces to a robust constraint in $(\mathcal{DRO})$ when we evaluate the expected violation $\mathbb{E}_{\mathbb{P}} \left[v_{\underline{\tau}_i, \overline{\tau}_i}(x_i^t) \right]$ concerning $\mathbb{P} = \hat{\mathbb{P}}$ and set $\Delta(\mathbb{P}, \hat{\mathbb{P}}) = 0$ by definition. The objective value C_{DRO} indicates the minimum worst-case costs given that the expected inventory violation is less than 0 in the ambiguity set. Hence, the budget C is set at least as C_{DRO} of $(\mathcal{DRO})$. Otherwise, $(\mathcal{RS})$ would be infeasible even when $k = 0$. The surplus cost $C - C_{\text{DRO}}$ justifies the DM's acceptable target violation to withstand greater ambiguity.

From a computational perspective, our ($\mathcal{RS}$) model and its DRO counterpart ($\mathcal{DRO}$) have a similar computational complexity since they share most constraints; see details in the reformulation part in Section 6. From a modeling perspective, the model ($\mathcal{DRO}$) solely immunizes against uncertainty within a restricted ambiguity set $\hat{\mathcal{P}}$, but does not safeguard losses when $\mathbb{P} \in \mathcal{P}_0 \backslash \hat{\mathcal{P}}$ (Long, Sim, and Zhou 2023). Given the budget C_{DRO}, the DRO model ($\mathcal{DRO}$) implies that:

$$\mathbb{E}_{\hat{\mathbb{P}}}\left[\tilde{v}\right] \leq 0 \qquad \forall \hat{\mathbb{P}} \in \hat{\mathcal{P}}$$
$$\mathbb{E}_{\hat{\mathbb{P}}}\left[\tilde{v}\right] \leq +\infty \qquad \forall \hat{\mathbb{P}} \notin \hat{\mathcal{P}}$$

By contrast, our model ($\mathcal{RS}$) suggests a fundamental consideration in risk mitigation over the whole probability space $\mathcal{P}_0$. The RS constraint imposes upper bounds on all distributions:

$$\mathbb{E}_{\mathbb{P}}[\tilde{v}] \leq k\Delta(\mathbb{P}, \hat{\mathbb{P}}) \quad \forall \mathbb{P} \in \mathcal{P}_0, \hat{\mathbb{P}} \in \hat{\mathcal{P}}$$

These are justified by our experiments in Section 6. From a managerial standpoint, compared to tuning the hyperparameter in the DRO literature, the budget C in our model ($\mathcal{RS}$) can be attained by solving ($\mathcal{DRO}$), which is easier for managers to configure and interpret.

4.2. Utility-based Service Level Measure

This section gives details on how to obtain a tractable model ($\mathcal{RS}$). The key point is to define the explicit form of the probability distance function $\Delta(\hat{\mathbb{P}}, \mathbb{P})$. The widely-used probability distances include the Kullback-Leibler divergence (Chen et al. 2018) and the Wasserstein metric (Mohajerin Esfahani and Kuhn 2018). However, these classic probability distances place a heavy computational burden, especially in multi-stage DRO models. To ease the computational burden, we are motivated by Cui et al. (2023a) to define Service Level Measure (SLM) as follows:

Definition 1 (Service Level Measure). Given a random inventory violation $v_{\underline{\tau},\overline{\tau}}(\boldsymbol{x})$ and the piece-wise linear utility function $u(v)$. We only know that the reference distribution $\hat{\mathbb{P}}$ belongs to an ambiguity set $\hat{\mathcal{P}}$. The SLM $\rho_{\underline{\tau},\overline{\tau}}(\cdot) : \mathcal{V}^T \mapsto \mathbb{R}_+$ is defined as :

$$\begin{aligned}
\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) = \inf \Bigg\{ k > 0 \,\Bigg|\, & \forall t \in [T], \\
& \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k} \right) \right] \leq 0 \Bigg\} \\
= \inf \Bigg\{ k > 0 \,\Bigg|\, & \forall t \in [T], m \in [M], \\
& \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) + b_m k \right] \leq 0 \Bigg\}
\end{aligned}$$

Furthermore, its underlying shortfall risk measure is:

$$\mu^u_{\hat{\mathbb{P}}}(\tilde{v}) = \inf \left\{ a \,\middle|\, \mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m\in[M]} \{a_m(\tilde{v}-a)+b_m\} \right] \le 0 \right\}$$

with the associated probability distance Δ_u represented as:

$$\Delta_u(\mathbb{P},\hat{\mathbb{P}}) = \sup_{\tilde{v}\in\mathcal{V}} \left\{ \mathbb{E}_{\mathbb{P}}[\tilde{v}] \,\middle|\, \forall m\in[M], \ \mathbb{E}_{\hat{\mathbb{P}}}[a_m\tilde{v}+b_m] \le 0 \right\}$$

We define the SLM based on utility-based probability distance, which is a dual form of the shortfall risk measure (Föllmer and Schied 2002). The underlying utility function $u(v)$ has many options. When we choose $u(x) = \exp(x) - 1$, then $\Delta_u(\mathbb{P},\hat{\mathbb{P}})$ is equivalent to Kullback–Leibler divergence $\Delta_{\mathrm{KL}}(\mathbb{P},\hat{\mathbb{P}})$; see more instances in (Cui et al. 2023a)). For tractability, we choose the utility function $u(\cdot)$ as a piece-wise linear function:

$$u(v) = \max_{m\in[M]} \{a_m v + b_m\}$$

where $a_m \ge 0, b_m$ for each segment $m \in [M]$ are given. This function preserves the linear structure of the model to alleviate the computational burden. Furthermore, the piece-wise linear function can approximate any non-decreasing convex utility function with sufficient pieces (Boyd and Vandenberghe 2004, Cui et al. 2023b).

We emphasize that SLM inherently considers multiple periods and selects the most robust solution over all periods. In the case of a single period, SLM falls under the umbrella of *Satisficing Measure* framework proposed by Brown and Sim (2009). SLM is also an instance of *Coherent Risk Enveloping Measure* that sets the enveloping bounds at all probabilistic levels (Chow, Cui, and Long 2022). Furthermore, SLM is an instance of the *Fragility Measure* with the utility-based probability distance (Long, Sim, and Zhou 2023). Consequently, SLM satisfies a series of salient properties in the following.

PROPOSITION 3 **(Properties of SLM).** *For any $\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{y}} \in \mathcal{V}^T$, the utility-based SLM $\rho_{\underline{\tau},\overline{\tau}}(\cdot) : \mathcal{V}' \to \mathbb{R}_+$ satisfies the following properties:*

1. *Monotonicity. If $v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) \ge v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}})$, then $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) \ge \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}})$.*
2. *Positive homogeneity. For any $\lambda \ge 0$, we have $\rho_{\lambda\underline{\tau},\lambda\overline{\tau}}(\lambda\tilde{\boldsymbol{x}}) = \lambda\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}})$.*
3. *Convexity. For any $\lambda \in [0,1]$, we have $\rho_{\underline{\tau},\overline{\tau}}(\lambda\tilde{\boldsymbol{x}} + (1-\lambda)\tilde{\boldsymbol{y}}) \le \lambda\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) + (1-\lambda)\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}})$.*
4. *Fragility*

(a) *Pro-Robustness. If* $v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) \le 0$*, then* $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) = 0$.

(b) *Antifragility. If there exists an* $t \in [T]$ *such that* $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) > 0$*, then* $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) = \infty$.

5. *Enveloping bound of violation probability. If utility function is lower bounded by* $\underline{u} \le 0$*, i.e.,* $u(v) \ge \underline{u}$*, and we obtain the optimal SLM* $k^* = \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}})$*, then for each* $\hat{\mathbb{P}} \in \hat{\mathcal{P}}$*,* $\theta > 0$*, we have*

$$\hat{\mathbb{P}}\left(v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) \ge \theta\right) \le \frac{-\underline{u}}{u\left(\frac{\theta}{k^*}\right) - \underline{u}}$$

6. *Right continuity.* $\lim_{a\downarrow 0} \rho_{-\boldsymbol{\infty},\boldsymbol{0}}(\tilde{\boldsymbol{x}} + a\boldsymbol{1}) = \rho_{-\boldsymbol{\infty},\boldsymbol{0}}(\tilde{\boldsymbol{x}})$*, where* $-\boldsymbol{\infty}$ *denotes the vector with all elements as* $-\infty$.

Proof See Online Appendix A.3. □

Monotonicity captures the DMs' aversion against higher inventory violation, i.e., if one violation almost surely dominates another, it would be less preferred under the SLM. The properties of monotonicity and positive homogeneity are inherited from coherent risk measures proposed by Artzner et al. (1999). Convexity property indicates a favor of risk diversification, which is justified by Föllmer and Schied (2002). Besides, convexity is of vital importance to obtain a tractable model. Pro-robustness dictates that if the random violation is less than 0 almost surely, the corresponding SLM attains the lowest value 0. On the contrary, antifragility alleges that the violation that always exceeds the inventory window has an infeasible solution, that is $\rho_{\underline{\tau},\overline{\tau}}(\cdot) = +\infty$. Enveloping bound property provides a probability bound on the inventory violation exceeding any level. We illustrate the enveloping bound with an example in Figure 5. Thus, using SLM as the objective function provides a performance guarantee on the inventory violation in the face of uncertain demands. Finally, we ensure that SLM remains unchanged to an infinitely small decrement with the technical condition of continuity.

Equipped with SLM, we focus on the reformulation of a tractable model. According to Theorem 1 by Cui et al. (2023a) and SLM defined by Definition 1, the Problem ($\mathcal{RS}$) is rewritten as follows.

THEOREM 1. *Under utility-based SLM, the model* ($\mathcal{RS}$) *can be equivalently reformulated as:*

$$\begin{aligned} \kappa^* = \inf \quad & \max_{i\in[N]} \{\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i)\} \\ s.t. \quad & (1)-(5) \qquad\qquad (\mathcal{RS}-\mathcal{SLM}) \\ & K\boldsymbol{1}^\top \boldsymbol{x} + \sup_{\hat{\mathbb{P}}\in\hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[\sum_{t\in[T]} C_t(\boldsymbol{a}^t)\right] \le C \end{aligned}$$

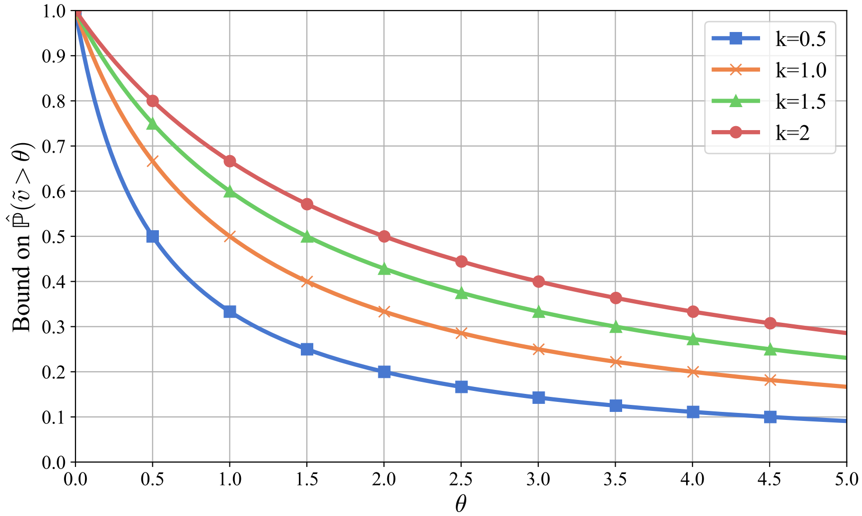

Figure 5 **Enveloping Probability Bounds on the Inventory Violation Exceeding** θ

Note. The underlying utility function is $u(v) = \max\{-1, v\}$, and the enveloping bound is $1/(1+\theta/k^*)$ for all $\hat{\mathbb{P}} \in \hat{\mathcal{P}}$. A lower k^* is preferable since it imposes a tighter bound on the probability of inventory violation exceeding all levels of θ.

where:

$$\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i) = \inf\left\{ k_i > 0 \,\middle|\, \forall t \in [T], m \in [M], \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(\tilde{x}_i^t) + b_m k_i\right] \leq 0 \right\}$$

Proof See Online Appendix A.4. □

The relations among our proposed models are depicted in Figure 6. In the model ($\mathcal{RS}-\mathcal{SLM}$), we consider minimizing the largest fragility $\max_{i\in[N]}\{\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i)\}$ over N swapping stations to ensure fair network service accessibility. The min-max form is used for two reasons: First, it provides individual SLM $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i)$ for each swapping station. We can identify the infeasible subset of constraints and then develop a new algorithm to improve the solution. Second, it affords extra flexibility to attach different weights w_i to each swapping station to reflect their relative importance:

$$\inf \max\{w_1\rho(\tilde{\boldsymbol{x}}_1), w_2\rho(\tilde{\boldsymbol{x}}_2), \ldots, w_N\rho(\tilde{\boldsymbol{x}}_N)\}$$

It may yield multiple optimal solutions, but some may not be Pareto optimal. The non-Pareto optimality issue calls for a new model. Hence, a Pareto optimization procedure can be conducted following the details provided in Algorithm 1 in Online Appendix B.1, which extends the original RS model. Our proposed RS model is also related to a recently proposed budget-driven DRO model by Hu, Chen, and Wang (2024). The relationship between our model and theirs is discussed in Appendix B.2.

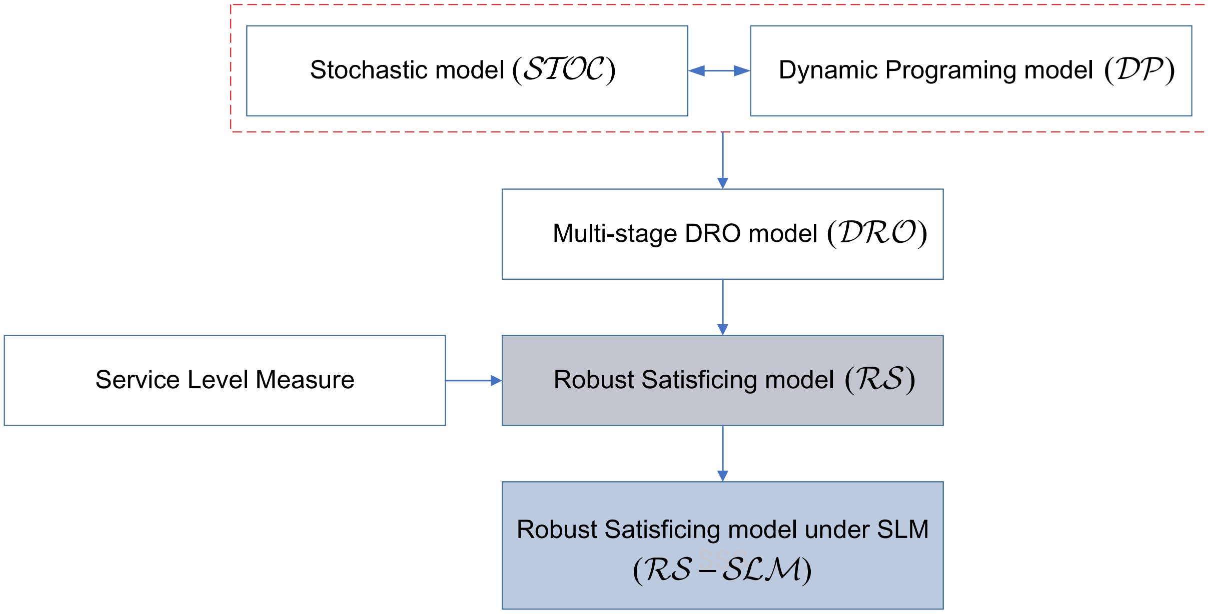

Figure 6 The Relations Among Models

5. Solution Approach

This section presents a solution approach to transform Problem ($\mathcal{RS}-\mathcal{SLM}$) into a computationally tractable model that can be solved by off-the-shelf solvers. It addresses two key issues: how to (i) construct the ambiguity set $\hat{\mathcal{P}}$ and (ii) propose approximation solutions based on a pooling-group-based linear decision rule.

5.1. Event-wise ambiguity set

To specify the reference ambiguity set $\hat{\mathcal{P}}$, we deploy the *Robust Stochastic Optimization* (RSO) framework proposed by Chen, Sim, and Xiong (2020). Under the RSO framework, the uncertainty is described by an event-wise ambiguity set, where an event can be a sample path or a surrogate scenario specified by covariate information. In the context of battery swapping, weather conditions and holidays can be seen as covariates to characterize uncertain demands (Perakis et al. 2023). Suppose there exist L historical samples coupled with covariate information, one can partition the space of covariates Ω into S disjoint subsets: $\Omega_s, s \in [S]$ with $\Omega_s \cap \Omega_{s'} = \emptyset$ for all $s, s' \in [S], s \neq t$ and $\cup_{s\in[S]}\Omega_s = \Omega$. Each subset Ω_s contains the samples with the same covariate information. Formally, we designate each Ω_s as an event $\tilde{s}$, which is a discrete random variable taking values in a finite set $[S]$. The probability p_s of the event s satisfies $\sum_{s\in[S]} p_s = 1$.

Conditioning on the realization of the event $\tilde{s} = s$, the support set of the uncertain demand $\tilde{\boldsymbol{d}}$ can be different and is denoted by $\mathcal{Z}_s$. The joint distribution of $(\tilde{\boldsymbol{d}}, \tilde{s})$ is denoted by $\hat{\mathbb{P}} \in \hat{\mathcal{P}}$. Therefore, we construct the event-wise ambiguity set as follows:

$$
\hat{\mathcal{P}} = \left\{ \begin{array}{l} \hat{\mathbb{P}} \in \mathcal{P}_0\left(\mathbb{R}^{NT} \times [S]\right) \\ \left| \begin{array}{l} (\tilde{\boldsymbol{d}}, \tilde{s}) \sim \hat{\mathbb{P}} \\ \hat{\mathbb{P}}(\underline{\boldsymbol{d}}^s \leq \tilde{\boldsymbol{d}} \leq \overline{\boldsymbol{d}}^s \mid \tilde{s} = s) = 1 \\ \qquad\qquad\qquad\qquad \forall s \in [S] \\ \mathbb{E}_{\hat{\mathbb{P}}}[\tilde{\boldsymbol{d}} \mid \tilde{s} = s] = \boldsymbol{\mu}_s \qquad \forall s \in [S] \\ \mathbb{E}_{\hat{\mathbb{P}}}[|\tilde{\boldsymbol{d}} - \boldsymbol{\mu}_s| \mid \tilde{s} = s] \leq \boldsymbol{\sigma}_s \\ \qquad\qquad\qquad\qquad \forall s \in [S] \\ \mathbb{E}_{\hat{\mathbb{P}}}\left[\left|\sum_{i \in \mathcal{N}_h} \frac{\tilde{d}_i^t - \mu_i^{ts}}{\sigma_i^s}\right| \middle| \tilde{s} = s\right] \leq \epsilon_h^s \\ \qquad \forall s \in [S], t \in [T], h \in [H] \\ \hat{\mathbb{P}}(\tilde{s} = s) = p_s, \qquad \forall s \in [S] \end{array} \right. \end{array} \right\}
$$

The event-wise ambiguity set includes the descriptive statistics: the support $[\underline{\boldsymbol{d}}^s, \overline{\boldsymbol{d}}^s]$, means $\boldsymbol{\mu}_s = (\mu_i^t)_{i \in [N], t \in [T]}$, and the upper bound for mean absolute deviations $|\tilde{\boldsymbol{d}} - \boldsymbol{\mu}_s|$. In particular, the fourth constraint is introduced to restrict the correlation between the random demand within a pooling group $\mathcal{N}_h$ for all $h \in [H]$. It preserves the linearity instead of using second-order moment information (e.g., covariance). Given S events, we can estimate the parameters in the ambiguity set by standard techniques in statistics. For example, the empirical means $\boldsymbol{\mu}_s$ are determined by $\boldsymbol{\mu}_s = \sum_{l \in \Omega_s} \boldsymbol{d}_s^l / |\Omega_s|$, where l index the sample in subset Ω_s.

The event-wise ambiguity set $\hat{\mathcal{F}}$ contains non-linear constraints, therefore, we introduce auxiliary variables $\tilde{\boldsymbol{w}} \in \mathbb{R}^{NT}$ and $\tilde{\boldsymbol{v}} \in \mathbb{R}^{HT}$ in the following. From the lifting and projection theorem of Bertsimas, Sim, and Zhang (2019), we can define a lifted ambiguity set $\hat{\mathcal{F}}$ which ensures $\hat{\mathcal{P}} = \prod_{(\tilde{\boldsymbol{d}}, \tilde{s})} \hat{\mathcal{F}}$.

$$
\hat{\mathcal{F}} = \left\{ \begin{array}{l} \hat{\mathbb{Q}} \in \mathcal{Q}_0\left(\mathbb{R}^{NT} \times \mathbb{R}^{NT} \times \mathbb{R}^{HT} \times [S]\right) \\ \left| \begin{array}{l} (\tilde{\boldsymbol{d}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{v}}, \tilde{s}) \sim \hat{\mathbb{Q}} \\ \mathbb{E}_{\hat{\mathbb{Q}}}[\tilde{\boldsymbol{d}} \mid \tilde{s} = s] = \boldsymbol{\mu}_s \quad \forall s \in [S] \\ \mathbb{E}_{\hat{\mathbb{Q}}}[\tilde{\boldsymbol{w}} \mid \tilde{s} = s] \leq \boldsymbol{\sigma}_s \quad \forall s \in [S] \\ \mathbb{E}_{\hat{\mathbb{Q}}}[\tilde{\boldsymbol{v}} | \tilde{s} = s] \leq \boldsymbol{\epsilon}_s \quad \forall s \in [S] \\ \hat{\mathbb{Q}}\left[(\tilde{\boldsymbol{d}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{v}}) \in \bar{\mathcal{Z}}_s \middle| \tilde{s} = s\right] = 1 \\ \qquad\qquad\qquad\qquad \forall s \in [S] \\ \hat{\mathbb{Q}}(\tilde{s} = s) = p_s \qquad \forall s \in [S] \end{array} \right. \end{array} \right\} \tag{8}
$$

where for each $\tilde{s} = s$, $\boldsymbol{\epsilon}_s = (\epsilon_h^s)_{h \in [H], t \in [T]}$, and the lifted support set $\bar{\mathcal{Z}}_s$ is defined as:

$$
\bar{\mathcal{Z}}_s = \left\{ \begin{array}{l} (\tilde{\boldsymbol{d}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{v}}) \in \mathbb{R}^{NT} \times \mathbb{R}^{NT} \times \mathbb{R}^{HT} \\ \left| \begin{array}{l} \underline{d}_i^{ts} \leq \tilde{d}_i^t \leq \overline{d}_i^{ts} \quad \forall i \in [N],\ t \in [T] \\ \tilde{d}_i^t - \mu_i^{ts} \leq \tilde{w}_i^t \quad \forall i \in [N],\ t \in [T] \\ \mu_i^{ts} - \tilde{d}_i^t \leq \tilde{w}_i^t \quad \forall i \in [N],\ t \in [T] \\ \sum_{i \in \mathcal{N}_h} \frac{\tilde{d}_i^t - \mu_i^{ts}}{\sigma_i^{ts}} \leq \tilde{v}_h^t \\ \qquad\qquad \forall h \in [H],\ t \in [T] \\ \sum_{i \in \mathcal{N}_h} \frac{\mu_i^{ts} - \tilde{d}_i^t}{\sigma_i^{ts}} \leq \tilde{v}_h^t \\ \qquad\qquad \forall h \in [H],\ t \in [T] \end{array} \right. \end{array} \right\} \tag{9}
$$

By introducing auxiliary variables, the lifted ambiguity set $\hat{\mathcal{F}}$ governs the distributions of the random variable tuple $(\tilde{\boldsymbol{d}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{v}}, \tilde{s})$. $\hat{\mathcal{F}}$ contains only linear expectation constraints, and the lifted support set $\bar{\mathcal{Z}}_s$ is a polyhedron. Note that any distribution $\hat{\mathbb{Q}} \in \hat{\mathcal{F}}$ can be rewritten as $\hat{\mathbb{Q}} = \sum_{s\in[S]} p_s \hat{\mathbb{Q}}_s$, where each marginal distribution $\hat{\mathbb{Q}}_s \in \hat{\mathcal{F}}_s$ corresponds to an event $s \in [S]$ with support set $\bar{\mathcal{Z}}_s$ (Chen, Sim, and Xiong 2020). Hence, the ambiguity set contains multi-modal distributions that are a mixture of S distributions.

Our focus is to solve the Problem ($\mathcal{RS}-\mathcal{SLM}$), where uncertainty is revealed progressively. In particular, the adaptive decisions $x_0^t, x_i^t, q_i^t, y_{ij}^t$ are dependent on $(\boldsymbol{d}^{[t-1]}, \boldsymbol{w}^{[t-1]}, \boldsymbol{v}^{[t-1]}, s)$ and u_i^t are dependent on $(\boldsymbol{d}^{[t]}, \boldsymbol{w}^{[t]}, \boldsymbol{v}^{[t]}, s)$. For notational, we denote this dependency by $x_0^{ts}(\cdot), x_i^{ts}(\cdot), q_i^{ts}(\cdot), u_i^{ts}(\cdot)$, and $y_{ij}^{ts}(\cdot)$. Under the lifted ambiguity set, we can transform Problem ($\mathcal{RS}-\mathcal{SLM}$) to an equivalent adjustable robust optimization (ARO) optimization problem in Theorem 2.

THEOREM 2. *With the lifted ambiguity set defined by* (8) *and* (9)*, Problem* ($\mathcal{RS}-\mathcal{SLM}$) *can be formulated as an ARO problem.*

Proof See Online Appendix A.4.

The equivalent ARO problem is presented in the online appendix when proving Theorem 2. The ARO problem is the robust counterpart of Problem ($\mathcal{RS}-\mathcal{SLM}$) under the lifted support set $\bar{\mathcal{Z}}_s$ for each $s \in [S]$. It is still unsolvable by commercial solvers since it involves infinite constraints and adaptive variables. Hence, we elaborate on obtaining an implementable model using the pooling-group-based linear decision rule proposed by Bertsimas, Sim, and Zhang (2019).

5.2. Pooling-group-based Enhanced Linear Decision Rule

The framework of *enhanced linear decision rule* (ELDR) is employed to endow the model with affine adaptability (Bertsimas, Sim, and Zhang 2019, Chen, Sim, and Xiong 2020). In particular, we restrict adaptive decisions to be affinely dependent on the revealed information within one pooling group. The adaptive decisions are restricted to an affine function set $\mathcal{A}$. In particular, we propose the following affine functions mapping from $\mathbb{R}^{2Nt+Ht}$ to $\mathbb{R}^M$ that are linearly dependent on the available information within pooling groups:

$$\mathcal{A}^M\left(\boldsymbol{d}^{[t]}, \boldsymbol{w}^{[t]}, \boldsymbol{v}^{[t]}\right) = \left\{ \boldsymbol{f} \in \mathcal{R}^{2Nt+Ht,M} \;\middle|\; \begin{array}{l} \exists \boldsymbol{f}^0, \boldsymbol{f}^1_{il}, \boldsymbol{f}^2_{il}, \boldsymbol{f}^3_{hl} \in \mathbb{R}^M : \\ f_i\left(\boldsymbol{d}^{[t]}, \boldsymbol{w}^{[t]}, \boldsymbol{v}^{[t]}\right) = f_i^0 \\ \quad + \sum_{j\in\Gamma(i), l\in[t]} f_{ij}^{l1} d_i^l \\ \quad + \sum_{j\in\Gamma(i), l\in[t]} f_{ij}^{l2} w_i^l \\ \quad + \sum_{h\in\mathcal{H}(i), l\in[t]} f_{ih}^{l3} v_h^l \end{array} \right\}$$

where $\Gamma(i) = \{j \mid (i, j) \in \mathcal{G}\}$ denotes the neighborhood of the station i in $\mathcal{G}$ and $\mathcal{H}(i) = \{h \mid i \in \mathcal{N}_h\}$ denotes the index set of pooling groups that include the station i. For the charging station i, we set $\Gamma(0) = [N]$ and $\mathcal{H}(0) = [H]$ since it connects every swapping station i. For any station $i \in [N] \cup 0$, we make adaptive decisions based on the demand information $\boldsymbol{d}^{[t]}, \boldsymbol{w}^{[t]}$ from $\Gamma(i)$ swapping stations and auxiliary information $\boldsymbol{v}^{[t]}$ from the pooling groups set $\mathcal{H}(i)$. This is different from the classic LDR, which considers only demand information (He, Hu, and Zhang 2020). As proved in Theorem 2 of Bertsimas, Sim, and Zhang (2019), incorporating auxiliary variables in the LDR would yield a better approximation than the classic LDR. Also, it differs from ELDR by He, Hu, and Zhang (2020), we emphasize the information based on the flexible structure $\mathcal{G}$.

Having introduced the pooling-group-based ELDR for adaptive decisions into the ARO model, we then have a classical affinely adjustable robust counterpart (AARC). Recall that the SLM and lifted ambiguity set contain only linear constraints; standard techniques from the robust literature can be leveraged to transform AARC into a linear program. The following proposition demonstrates the reformulation procedure.

PROPOSITION 4. *For any $t \in [T], i \in [N], s \in [S]$ and $m \in [M]$, the constraint*

$$\begin{aligned} &o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \\ &\quad \geq p_s\left[a_m(x_i^{ts}(\cdot) - \overline{\tau}_i) + b_m k\right] \;\forall\, (\boldsymbol{d}, \boldsymbol{w}, \boldsymbol{v}) \in \bar{\mathcal{Z}}_s \end{aligned}$$

can be formulated into the following linear constraints based on the proposed ELDR.

$$\begin{aligned}
&\boldsymbol{\pi}_{i1}^{tsm}\underline{\boldsymbol{d}}^s - \boldsymbol{\pi}_{i2}^{tsm}\overline{\boldsymbol{d}}^s + (\boldsymbol{\pi}_{i3}^{tsm} - \boldsymbol{\pi}_{i4}^{tsm})\boldsymbol{\mu}^s + \\
&\sum_{h\in[H]}\sum_{l\in[T]}\sum_{i'\in\mathcal{N}_h}\frac{\mu_{i'}^{ls}}{\sigma_{i'}^{ls}}(\eta_{ihl1}^{tsm} - \phi_{ihl1}^{tsm}) \\
&\qquad \geq p_s a_m (x_i^{ts0} - \overline{\tau}_i) + p_s b_m k_i - o_i^{ts} \\
&\pi_{ii'l1}^{tsm} - \pi_{ii'l2}^{tsm} + \pi_{ii'l3}^{tsm} - \pi_{ii'l4}^{tsm} \\
&+ \frac{1}{\sigma_{i'}^{ls}}\sum_{h'\in\mathcal{H}(i')}(\eta_{ihl1}^{tsm} - \phi_{ihl1}^{tsm}) \leq z_{ii'l}^{ts} - p_s a_m x_{ii'l}^{ts1} \\
&\qquad \forall i' \in \Gamma(i), l \in [t-1] \\
&\pi_{ii'l1}^{tsm} - \pi_{ii'l2}^{tsm} + \pi_{ii'l3}^{tsm} - \pi_{ii'l4}^{tsm} \\
&+ \frac{1}{\sigma_{i'}^{ls}}\sum_{h'\in\mathcal{H}(i')}(\eta_{ihl1}^{tsm} - \phi_{ihl1}^{tsm}) \leq z_{ii'l}^{ts} \\
&\qquad \forall (i', l) \in ([N]\times[T]) \setminus (\Gamma(i)\times[t-1]) \\
&\pi_{ii'l3}^{tsm} + \pi_{ii'l4}^{tsm} \leq \beta_{ii'l}^{ts} - p_s a_m x_{ii'l}^{ts2} \\
&\qquad \forall i' \in \Gamma(i), l \in [t-1] \\
&\pi_{ii'l3}^{tsm} + \pi_{ii'l4}^{tsm} \leq \beta_{ii'l}^{ts} \\
&\qquad \forall (i', l) \in ([N]\times[T]) \setminus (\Gamma(i)\times[t-1]) \\
&\eta_{ihl}^{tsm} + \phi_{ihl}^{tsm} \leq \gamma_{ihl}^{ts} - p_s a_m x_{ihl}^{ts3} \\
&\qquad \forall h \in \mathcal{H}(i), l \in [t-1] \\
&\eta_{ihl}^{tsm} + \phi_{ihl}^{tsm} \leq \gamma_{ihl}^{ts} \\
&\qquad \forall (h, l) \in ([H]\times[T]) \setminus (\mathcal{H}(i)\times[t-1]) \\
&\boldsymbol{\pi}_{i1}^{tsm}, \boldsymbol{\pi}_{i2}^{tsm}, \boldsymbol{\pi}_{i3}^{tsm}, \boldsymbol{\pi}_{i4}^{tsm} \in \mathbb{R}_+^{NT}, \boldsymbol{\eta}_i^{tsm}, \boldsymbol{\phi}_i^{tsm} \in \mathbb{R}_+^{HT}
\end{aligned}$$

Proof See Online Appendix A.6.

With Proposition 4, Problem ($\mathcal{RS}-\mathcal{SLM}$) can be reformulated into a linear program. The result reformulation preserves linearity and facilitates practical implementation. The linear problem includes $O(N^2T^2SM)$ continuous decision variables and $O(N^2T^2SM)$ constraints in total. Given the number of events S, both the number of variables and the number of constraints of the resulting linear program grow quadratically with the number of stations/periods.

6. Numerical Studies

This section conducts a series of numerical studies on synthetic data to evaluate the effectiveness of our proposed model and explore managerial insights. The locations of swapping and charging stations are sampled from NIO's 45 swapping stations in Shanghai using Amap 2024 (Amap 2024). For presentation brevity, most details of input data and parameters are relegated to Appendix C.1. In what follows, we conduct experiments in Section 6.1 to compare our RS model with benchmark models, showcasing its effectiveness and robustness. The model's out-of-sample performance under perturbed distributions is also investigated. Section 6.2 demonstrates the values of transshipment by comparing models on different flexibility structures. Finally, the impact of budget and penalty cost are analyzed. The program is coded in Python and executed on a PC with an Intel Core i5-11300H 3.10GHz and 16 GB RAM using Gurobi 11.0.3. We set $\epsilon = 10^{-4}$ and time limit of 2 hours throughout the experiments. All the problems are solved using their respective AARC.

6.1. Comparison with benchmark models

We benchmark our RS model against the other four models, including variants of our proposed RS model and the related DRO counterparts. In all variants of the RS model, the embedding utility function of SLM is set as $u(x) = \max\{-1, x, ex - 1\}$, which can approximate the exponential utility function $e^x - 1$. Unless stated explicitly, each model uses the ambiguity set in (8) and the ELDR proposed in Section 5.1. We introduce these five models briefly:

1. DRO model with event-wise information (DRO-S). The first model minimizes the worst-case expected total cost based on Problem ($\mathcal{DRO}$).

2. DRO model without event-wise information (DRO-NS). This model adapts from DRO-S by excluding the event-wise ambiguity set. That is to say, we only consider one event $\tilde{s} = 1$ and $p = 1$ as in He, Hu, and Zhang (2020).

3. RS model with event-wise information (RS-S). The third model is our proposed robust satisficing model in Problem ($\mathcal{RS} - \mathcal{SLM}$).

4. RS model without event-wise information (RS-NS). This model is similar to RS-S except that we do not incorporate scenario information in the ambiguity set.

5. RS model based on mean value (RS-PE). The last benchmark is a special case of RS-S where we set $\underline{d}_i^{ts} = \overline{d}_i^{ts} = \mu_i^{ts}$ and there is only one demand realization in each event $s \in [S]$.

For a fair comparison, all three RS model variants adopt the same budget based on the maximum total cost of both DRO benchmarks. We calculate the statistics of individual violation $v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(\tilde{x}_i^t)$,

Table 1 Comparison of the Inventory Violation Between Solutions from the Five Benchmarks

N	Model	Prob %	Mean	Std.	VaR95%	CVaR95%
3	RS-S	**5.30**	**0.40**	2.53	**0.30**	**0.52**
	DRO-S	17.22	0.74	**1.98**	6.24	7.62
	DRO-NS	17.22	0.74	**1.98**	6.24	7.62
	RS-NS	16.98	0.76	2.43	5.52	9.70
	RS-PE	60.47	5.08	7.16	21.19	3.34
4	RS-S	**3.34**	**0.25**	2.49	**0.00**	**0.25**
	DRO-S	7.80	0.40	**1.55**	4.24	5.92
	DRO-NS	7.80	0.40	**1.55**	4.24	5.92
	RS-NS	13.83	1.14	4.01	9.13	17.76
	RS-PE	63.44	5.01	8.20	12.29	30.80
5	RS-S	**6.34**	**0.39**	**1.64**	**3.98**	**7.44**
	DRO-S	17.65	0.77	2.04	7.48	5.92
	DRO-NS	17.65	0.77	2.04	7.48	5.92
	RS-NS	21.88	1.95	6.77	9.17	28.19
	RS-PE	52.41	3.96	6.87	7.97	27.53
7	RS-S	**3.30**	**0.23**	**1.53**	**0.00**	**0.23**
	DRO-S	17.31	0.76	2.07	7.48	7.65
	DRO-NS	17.31	0.76	2.07	7.48	7.65
	RS-NS	–	–	–	–	–
	RS-PE	62.60	3.53	3.95	7.48	10.95

Note. Values in bold indicate the best performance.

including the probability of being positive (Prob), the expected value (Mean), the standard deviation (Std), the 95% CVaR, and value at risk at 95% (VaR95%).

Table 1 demonstrates the performance metrics of 5 models for $N \in \{3,4,5,7\}$. Overall, the RS-S generally dominates other methods in all five metrics because lower metric values indicate better performance. Second, we observe that our proposed RS-S manifests superior performance over other models in controlling the violation in average and extreme cases (see Mean and CVaR95%). This confirms that RS framework mitigates risk over the whole probability space. It can be seen that RS-PE performs poorly in all metrics due to the overfitting in-sample, since it only considers deterministic mean values. Moreover, in case $N = 7$, the RS-NS is infeasible with the given moderate budget. This underscores the importance of incorporating event information inthe RS model. We also show the box plot of the inventory violation for five models in Figure 7. It is evident that our proposed ($\mathcal{RS}-\mathcal{SLM}$) outperforms the other four models in terms of both the number and interval of violation samples.

To delve into the optimal decisions of the five models, we investigate the mean values of inventory dynamics in Figure 8. Almost all the inventory of RS-S falls within the predefined intervals and exhibits high stability, while RS-PE fluctuates with inconsistent variations. An interesting

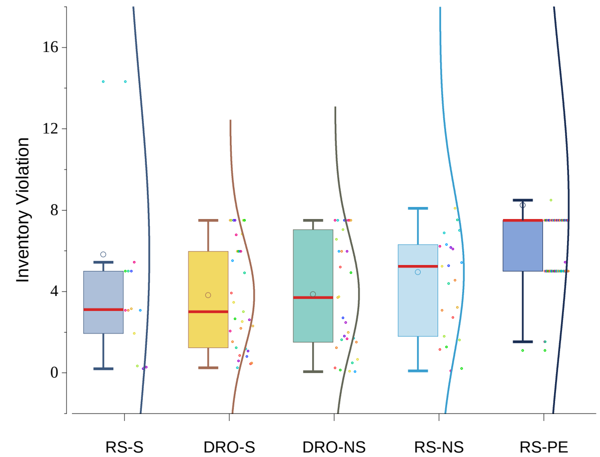

Figure 7 Box plot of inventory violation of five models ($N = 3$)

Note. The dots denote violation samples and the curve represents the empirical normal distribution of violations.

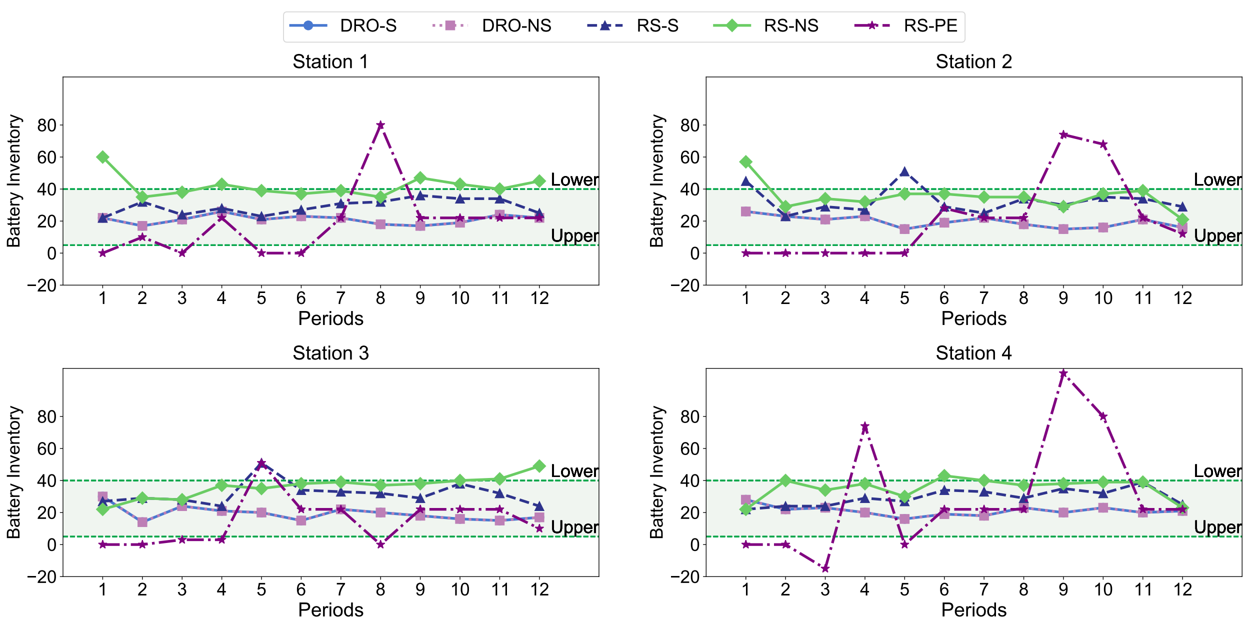

Figure 8 Mean values of battery inventory dynamics at swapping stations for five Models ($\tilde{s} = 1$)

Note. The shadowed area implies the imposed inventory window.

observation is that the RS-S maintains the on-hand inventory around the mean demand steadily ($\mu_1 = 15, \mu_2 = 30$). Hence, this accords with the classic base-stock policy that we replenish the inventory to reach a pre-defined base stock.

To solidify the anti-fragility and robustness of our proposed model, we conduct another out-of-sample test that uses the network of 1 charging station and $N = 4$ swapping stations. We add a perturbation Δ to the mean of the out-of-sample demand as in Li et al. (2024). See details in Appendix C.2. Additionally, the performance of RS-PE is far inferior to the other four models, so

Table 2 **Out-of-Sample Statistics for Four Models**

Distribution	Model	Prob %	Mean	Std.	VaR95%	CVaR95%
Truncated Normal ($\Delta = -0.1$)	RS-S	**6.90**	**0.30**	**1.40**	**2.21**	**5.58**
	DRO-S	11.13	0.57	1.97	4.32	7.34
	DRO-NS	11.13	0.57	1.97	4.32	7.34
	RS-NS	13.02	0.97	3.70	6.29	15.30
Truncated Normal ($\Delta = -0.05$)	RS-S	**6.07**	**0.25**	**1.24**	**1.28**	**4.89**
	DRO-S	9.29	0.46	1.71	4.32	6.26
	DRO-NS	9.29	0.46	1.71	4.32	6.26
	RS-NS	12.75	0.96	3.70	6.26	15.30
Truncated Normal ($\Delta = 0.05$)	RS-S	**4.31**	**0.16**	**0.93**	**0.00**	**0.16**
	DRO-S	9.38	0.40	1.46	4.32	5.73
	DRO-NS	9.38	0.40	1.46	4.32	5.73
	RS-NS	12.75	0.96	3.70	6.26	15.30
Truncated Normal ($\Delta = 0.1$)	RS-S	**3.80**	**0.13**	**0.80**	**0.00**	**0.13**
	DRO-S	20.50	0.78	1.93	4.78	7.26
	DRO-NS	20.50	0.78	1.93	4.78	7.26
	RS-NS	16.60	1.05	3.73	6.50	15.36
Uniform	RS-S	**10.55**	**0.60**	**2.44**	**5.00**	**9.21**
	DRO-S	16.31	1.05	3.36	7.48	12.73
	DRO-NS	16.31	1.05	3.36	7.48	12.73
	RS-NS	17.60	1.19	3.87	7.89	15.36
Poisson	RS-S	**5.02**	**0.21**	**1.11**	**0.03**	**4.15**
	DRO-S	7.89	0.38	1.45	4.32	5.77
	DRO-NS	7.89	0.38	1.45	4.32	5.77
	RS-NS	11.97	0.88	3.58	5.36	14.58

Note. Bold values highlight the best performance in each distribution scenario.

we do not include it in the following comparison. Table 2 summarizes the statistics of out-of-sample inventory violation from four models under several distributions. Additionally, we normalize the performance of the four models in Figure 9. The previous observation remains solid: the RS-S reduces inventory violation and withstands greater demand uncertainty. Hence, we can conclude that our RS-S model provides highly robust inventory management decisions in battery swapping.

6.2. The Value of Transshipment

We further investigate the value of transshipment by comparing the total costs of ($\mathcal{DRO}$). The parameter setting is consistent with the former section, except that we assume demands are i.i.d. among swapping stations. As shown in Figure 10, we depict the average performance indicators of each flexibility structure. As we expected, the lateral transshipment strategy reduces the total costs, we can observe the following fact, which supports Proposition 1: Cost(Dedicated) ≥ Cost(2-chain) ≥ Cost(Full Flexibility). It is demonstrated that adding a little flexibility in the network, i.e., transforming a dedicated system to the 2-chain structure, evidently drives down the violation in five metrics. Our results add new evidence to the well-known fact that "the chaining structure is almost as good as a full flexibility structure" (Chou et al. 2010).

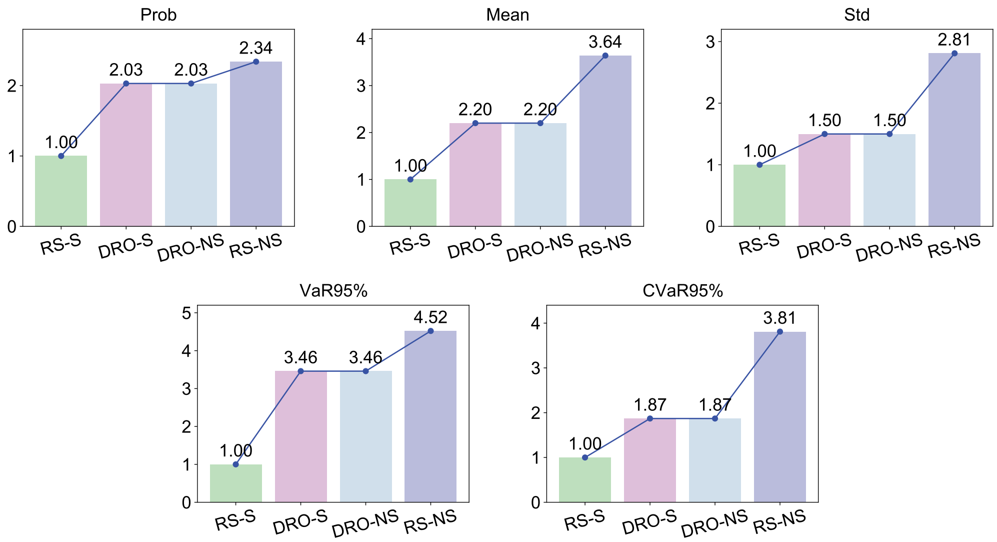

Figure 9 Average Performance of RS-S and Benchmark Models

Note. A lower value is preferable, and 1.00 indicates the best performance.

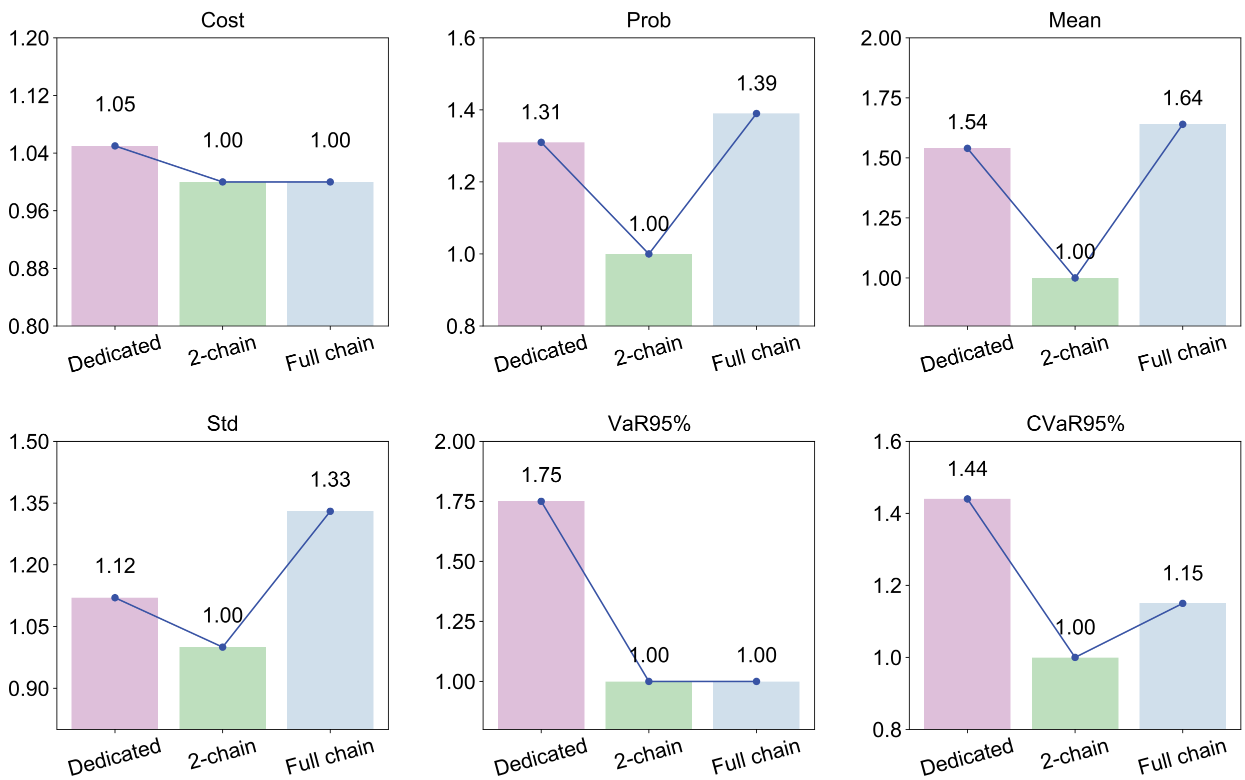

Figure 10 Performance of DRO-S Model under Different Flexibility Structures

We further investigate how the demand correlation affects inventory pooling. We conduct the experiment using varying levels of demand correlation $\alpha \in \{0, 0.25, 0.50, 0.75, 1.00\}$. We also consider a general correlation where the correlation can be positive and negative. Table 3 suggests

Table 3 **Comparison of Total Costs $(\times 10^4)$ of DRO-S Model Under Different Demand Correlation Levels**

Type	α	Dedicated		2-chain		Full Flexibility
		Cost	Gap	Cost	Gap	
Positive	0	496.30	25.44 (5.40%)	472.15	1.28 (0.27%)	470.86
	0.25	496.30	25.19 (5.35%)	472.48	1.37 (0.29%)	471.11
	0.5	496.30	24.98 (5.30%)	472.87	1.55 (0.33%)	471.32
	0.75	496.30	24.77 (5.25%)	473.39	1.86 (0.39%)	471.53
	1	496.30	24.67 (5.23%)	474.38	2.75 (0.58%)	471.63
General	0	496.30	25.36 (5.38%)	472.15	1.20 (0.26%)	470.94
	0.25	496.30	25.15 (5.34%)	472.48	1.32 (0.28%)	471.16
	0.5	496.30	25.12 (5.33%)	472.87	1.69 (0.36%)	471.18
	0.75	496.30	24.86 (5.27%)	473.39	1.95 (0.41%)	471.44
	1	496.30	24.61 (5.22%)	474.38	2.69 (0.57%)	471.69

Note. Gap = (Cost(Dedicated or k-Chain) − Cost(Full Flexibility))/Cost(Full Flexibility)

a minimum saving of USD$240,000 (5.26%) in total costs when transshipment is allowed. The performance gap between 2-chain and full flexibility is very close, ranging from 0.26% to 0.57%. Hence, the 2-chain achieves the most benefits of inventory sharing, and also incurs a lower transshipment cost than full flexibility does. Furthermore, the cost gap for 2-chain and full flexibility models slightly increases, from 1.28×10^4 to 2.75×10^4, while the gap for the dedicated structure decreases with respect to the correlation α. It is in line with the observation that the value of inventory pooling is shrinking with the increase of demand correlation (Cui et al. 2023b).

6.3. Impact of Cost Budget

In Figure 11, we check the budget's impact on the performance of RS-S and RS-NS models. Notably, as the budget increases, both models' average violation probability and means exhibit an overall decreasing trend, except for the cases where $\delta = 2, 3$. Our RS-S have a smaller probability and means of inventory violation than RS-NS. This underscores the effectiveness of event-wise ambiguity set, which captures more accurate covariate information and minimizes fragility caused by uncertain demand. We provide a tool for the DMs to specify the scaling factor δ to accommodate their risk preference. In practice, adjusting the scaling factor δ manually is more interpretable and straightforward than tuning hyperparameters in the DRO models.

6.4. Impact of Penalty Cost

Figure 12 reports the total costs for three models as the penalty cost scales up from 100 to 5100. There is an intuitive fact that higher penalty costs for stock-out will result in higher battery inventory and total costs. Of special interest are the relative tendencies of total costs and battery costs. The total cost shows a steady growth trend, whereas the battery cost exhibits a highly unstable fluctuation

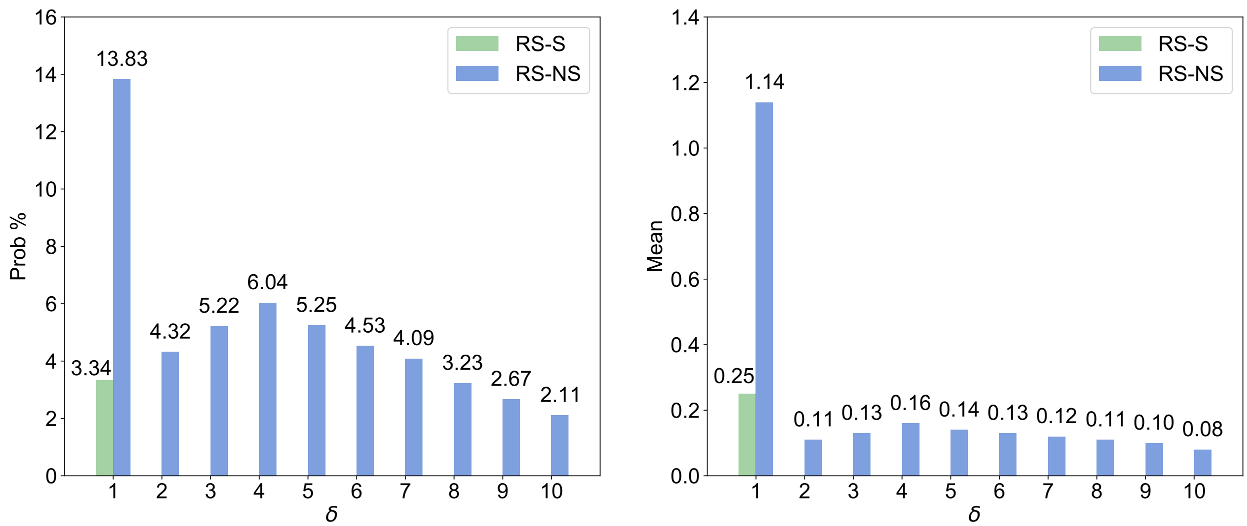

Figure 11 Comparison of (a) Violation Probability and (b) Mean Values of Inventory Violation under Different Budget Levels C.

Note. Here we vary the budget $C \in \{1.2C_{\text{DRO}}, 1.2C_{\text{DRO}} + \delta * 10000\}$ in which scaling factor δ ranges from 0 to 9. The RS-S achieves a 0 violation with high probability when the scaling factor δ exceeds 1.

trend. In particular, the battery costs remain relatively low when penalty costs are below \$3600 but rise sharply as penalty costs increase. The trends exhibit that most costs are incurred by stock-out penalties rather than investments in batteries. We also observe that the probability of inventory violation stays relatively high (from 30% to 70%) with low penalty costs (< \$3600) but plummet from 40% to 3% as penalty costs surpass \$4100. The sharp change is mainly due to the fact that, when penalty costs are low relative to battery costs, paying stock-out penalties is more cost-effective than investing in large quantities of batteries. This counterintuitive result highlights the limitation of using cost minimization as the sole system objective. Therefore, we recommend excluding penalty costs from operational costs and instead adopting our proposed SLM as decision criterion.

7. Conclusion

This paper studies the battery inventory management problem with uncertain demand in a "swap-locally, charge-centrally" network, where distributed swapping stations replenish from the charging station and transship to neighboring stations. To leverage the benefits of lateral transshipment, we first introduce ideas from process flexibility to develop a stochastic dynamic programming model. To deal with the "curse of dimensionality" and distributionally ambiguity, we introduce a novel target-oriented robust satisficing model ($\mathcal{RS}-\mathcal{SLM}$) under the robust satisficing framework. To evaluate the uncertain violation in achieving two-sided targets, we construct a utility-based Service Level Measure based on the utility function-based probability distance. For tractability, we

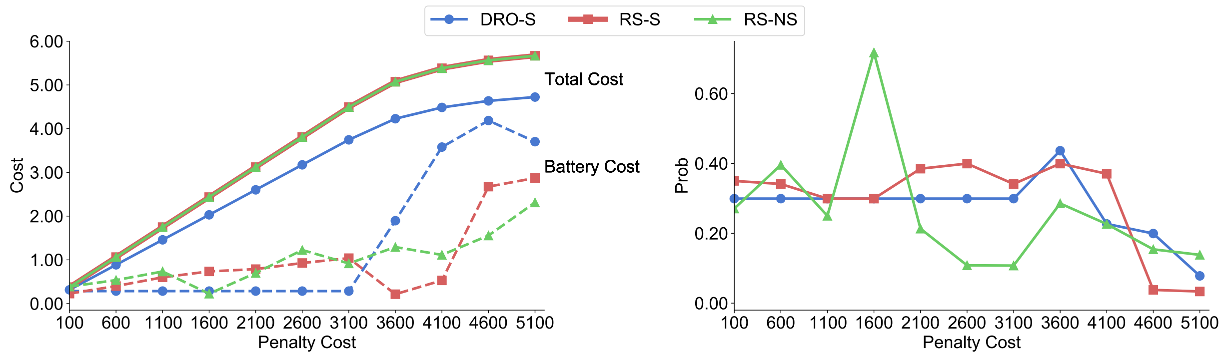

Figure 12 Comparison of (a) Violation probability and (b) Mean Values of Inventory violation under Different Penalty Costs.

Note. The solid line represents total costs and the dashed line represents battery investment, respectively.

employ an event-wise ambiguity set and a pooling-group-based enhanced linear decision rule. The numerical study validates our proposed model is highly robust to distributional ambiguity in out-of-sample tests. Our framework provides a practical tool for operators to operate a "swap-locally, charge-centrally" network and balance their cost objectives with service level requirements.

Further research questions pertaining to this problem setting remain to be explored in the future. From a model perspective, while this study adopts a Service Level Measure constructed by utility-based probability distance, it would be beneficial to compare the effects of different fragility measures constructed by other distributional distances. Limited by computational capacity, we do not investigate the marginal profits of flexibility structure in larger networks. It would be valuable to quantify the benefits of various flexibility structures (3-chain, 4-chain, etc.) on a large scale in our problem setting. It would be beneficial to consider designing an efficient sparse structure for the battery swapping-charging network by optimization techniques. Our model is also flexible to extend to operational issues, for example, backlogged demand and on-site charging at swapping stations. Furthermore, integrating charging decisions in battery swap stations with the electricity market would also be a worthwhile direction.

References

Amap, 2024 *A mapping and navigation service*. URL `https://lbs.amap.com/`, Accessed August 11, 2024.

Artzner P, Delbaen F, Eber JM, Heath D, 1999 *Coherent measures of risk. Math. Finance* 9(3):203–228.

Asadi A, Nurre Pinkley S, 2022 *A monotone approximate dynamic programming approach for the stochastic scheduling, allocation, and inventory replenishment problem: Applications to drone and electric vehicle battery swap stations. Transportation Sci.* 56(4):1085–1110.

Avci B, Girotra K, Netessine S, 2015 *Electric vehicles with a battery switching station: Adoption and environmental impact*. *Management Sci.* 61(4):772–794.

Avishan F, Elyasi M, Yanıkoğlu İ, Ekici A, Özener OÖ, 2023 *Humanitarian relief distribution problem: An adjustable robust optimization approach*. *Transportation Sci.* 57(4):1096–1114.

Benjaafar S, Jiang D, Li X, Li X, 2022 *Dynamic inventory repositioning in on-demand rental networks*. *Management Sci.* 68(11):7861–7878.

Bertsimas D, Sim M, Zhang M, 2019 *Adaptive distributionally robust optimization*. *Management Sci.* 65(2):604–618.

Boyd SP, Vandenberghe L, 2004 *Convex Optimization* (Cambridge University Press, Cambridge).

Brown DB, Sim M, 2009 *Satisficing measures for analysis of risky positions*. *Management Sci.* 55(1):71–84.

Chen Y, Guo Q, Sun H, Li Z, Wu W, Li Z, 2018 *A distributionally robust optimization model for unit commitment based on Kullback – Leibler divergence*. *IEEE Trans. Power System* 33(5):5147–5160.

Chen Z, Sim M, Xiong P, 2020 *Robust stochastic optimization made easy with RSOME*. *Management Sci.* 66(8):3329–3339.

Chou MC, Chua GA, Teo CP, Zheng H, 2010 *Design for process flexibility: Efficiency of the long chain and sparse structure*. *Oper. Res.* 58(1):43–58.

Chow VTF, Cui Z, Long DZ, 2022 *Target-oriented distributionally robust optimization and its applications to surgery allocation*. *INFORMS J. Comput.* 34(4):2058–2072.

Cui Z, Ding J, Long DZ, Zhang L, 2023a *Target-based resource pooling problem*. *Production Oper. Management* 32(4):1187–1204.

Cui Z, Long DZ, Qi J, Zhang L, 2023b *The inventory routing problem under uncertainty*. *Oper. Res.* 71(1):378–395.

DeValve L, Wei Y, Wu D, Yuan R, 2023 *Understanding the value of fulfillment flexibility in an online retailing environment*. *Manufacturing Service Oper. Management* 25(2):391–408.

Föllmer H, Schied A, 2002 *Convex measures of risk and trading constraints*. *Finance Stochast.* 6:429–447.

Guide Jr VDR, Srivastava R, 1997 *Repairable inventory theory: Models and applications*. *Eur. J. Oper. Res.* 102(1):1–20.

He L, Hu Z, Zhang M, 2020 *Robust repositioning for vehicle sharing*. *Manufacturing Service Oper. Management* 22(2):241–256.

Hu J, Chen Z, Wang S, 2024 *Budget-driven multiperiod hub location: A robust time-series approach*. *Oper. Res.* 73(2):613–631.

IEA, 2024 *Global EV Outlook 2024*. URL `https://iea.blob.core.windows.net/assets/a9e3544b-0b12-4e15-b407-65f5c8ce1b5f/GlobalEVOutlook2024.pdf`, Accessed: September 13, 2024.

Jordan WC, Graves SC, 1995 *Principles on the benefits of manufacturing process flexibility*. *Management Sci.* 41(4):577–594.

Ledvina K, Qin H, Simchi-Levi D, Wei Y, 2022 *A new approach for vehicle routing with stochastic demand: Combining route assignment with process flexibility*. *Oper. Res.* 70(5):2655–2673.

Lee HL, 1987 *A multi-echelon inventory model for repairable items with emergency lateral transshipments*. *Management Sci.* 33(10):1302–1316.

Li H, Delage E, Zhu N, Pinedo M, Ma S, 2024 *Distributional robustness and inequity mitigation in disaster preparedness of humanitarian operations*. *Manufacturing Service Oper. Management* 26(1):197–214.

Long DZ, Sim M, Zhou M, 2023 *Robust satisficing*. *Oper. Res.* 71(1):61–82.

Lyu G, Cheung WC, Chou MC, Teo CP, Zheng Z, Zhong Y, 2019 *Capacity allocation in flexible production networks: Theory and applications*. *Management Sci.* 65(11):5091–5109.

Mak HY, Rong Y, Shen ZJM, 2013 *Infrastructure planning for electric vehicles with battery swapping*. *Management Sci.* 59(7):1557–1575.

Mohajerin Esfahani P, Kuhn D, 2018 *Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations*. *Math. Programming* 171(1-2):115–166.

Olsson F, 2015 *Emergency lateral transshipments in a two-location inventory system with positive transshipment leadtimes*. *Eur. J. Oper. Res.* 242(2):424–433.

Perakis G, Sim M, Tang Q, Xiong P, 2023 *Robust pricing and production with information partitioning and adaptation*. *Management Sci.* 69(3):1398–1419.

Qi W, Zhang Y, Zhang N, 2023 *Scaling up electric-vehicle battery swapping services in cities: A joint location and repairable-inventory model*. *Management Sci.* 69(11):6855–6875.

Schneider F, Thonemann UW, Klabjan D, 2018 *Optimization of battery charging and purchasing at electric vehicle battery swap stations*. *Transportation Sci.* 52(5):1211–1234.

Sherbrooke CC, 1968 *METRIC: A multi-echelon technique for recoverable item control*. *Oper. Res.* 16(1):122–141.

Simchi-Levi D, Wei Y, 2012 *Understanding the performance of the long chain and sparse designs in process flexibility*. *Oper. Res.* 60(5):1125–1141.

Simon HA, 1955 *A behavioral model of rational choice*. *The Quarterly Journal of Economics* 99–118.

Smith JE, Winkler RL, 2006 *The optimizer's curse: Skepticism and postdecision surprise in decision analysis*. *Management Sci.* 52(3):311–322.

US Government, 2023 *Inflation Reduction Act*. URL `https://www.whitehouse.gov/wp-content/uploads/2022/12/Inflation-Reduction-Act-Guidebook.pdf`, Accessed: September 13, 2024.

Xie X, Dai X, Pei Z, 2024 *Empowering the capillary of the urban daily commute: Battery deployment analysis for the locker-based e-bike battery swapping*. *Transportation Sci.* 58(1):176–197.

Online Appendix for "Target-Oriented Robust Inventory Management in Electric Vehicle Battery Swapping Networks"

The online appendix is organized as follows. Appendix A provide the proofs for all statements. Appendix B shows the extensions of our proposed RS model. Appendix C presents all supplementary materials for numerical studies.

A Proofs of Statements

A.1 Proof of Proposition 1

Proof. Let $(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger})$ be an optimal solution of model $(\mathcal{DP})$ for dedicated structure $\mathcal{D}$ and define the objective value $C\left(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger} \mid \mathcal{D}\right)$ such that:

$$Z^*(\mathcal{D}) = C\left(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger} \mid \mathcal{D}\right)$$

where $Z^*(\mathcal{D})$ denote the optimal objective value under dedicated structure $\mathcal{D}$. Observe that $(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger})$ is feasible in $(\mathcal{DP})$ for chaining structure $\mathcal{C}$. Furthermore, $(\mathcal{DP})$ is a minimization problem, we have:

$$\begin{aligned} Z^*(\mathcal{D}) &= C\left(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger} \mid \mathcal{D}\right) \\ &= C\left(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger} \mid \mathcal{C}\right) \\ &\geq Z^*(\mathcal{C}) \end{aligned}$$

The first equality holds since the dedicated structure $\mathcal{D}$ is a subset of a chaining structure $\mathcal{C}$, so $(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger})$ yields the same objective value in $\mathcal{D}$ and $\mathcal{C}$. The inequality holds true because $(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger})$ is a feasible solution, and $\mathcal{DP}$ is a minimization problem. A similar reasoning applies to the chaining structure $\mathcal{C}$ and the full flexibility structure $\mathcal{F}$, so we omit the details for brevity. We now have $Z^*(\mathcal{C}) \geq Z^*(\mathcal{F})$. Combining two inequalities, therefore,

$$Z^*(\mathcal{D}) \geq Z^*(\mathcal{C}) \geq Z^*(\mathcal{F})$$

□

A.2 Proof of Proposition 2

Proof. Let $(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger})$ be an optimal solution of the model $(\mathcal{DRO})$. Note that $(\boldsymbol{x}^{\dagger}, \boldsymbol{u}^{\dagger}, \boldsymbol{q}^{\dagger}, \boldsymbol{y}^{\dagger})$ satisfy all the constraints in the RS model $(\mathcal{RS})$ since $C \geq C_{DRO}$, therefore it is feasible for model $(\mathcal{RS})$. Therefore, any optimal solution to the model $(\mathcal{DRO})$ must be feasible in the RS model $(\mathcal{RS})$.

We state that, if $C = C_{DRO}$, then the optimal solution to the model $(\mathcal{RS})$ is also optimal to the model $(\mathcal{DRO})$. That is to say, the robust satisficing constraint $\sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[C\left(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y} \mid \mathcal{G}\right)\right] \leq C$ is tight at the optimality of the RS model $(\mathcal{RS})$. To see this, we first assume the robust satisficing constraint is not tight, then we have a budget lower than C_{DRO}. This implies a strictly lower objective in the model $(\mathcal{DRO})$, which leads to the contradiction of the optimality assumption. Therefore, if $C = C_{DRO}$, then the optimal solution to the model $(\mathcal{RS})$ is also optimal to the model $(\mathcal{DRO})$.

If we write the RS model as a parameterized function $\rho(C)$ dependent on a budget C, then we have $\rho(C_1) \leq \rho(C_2)$ for any $C_{DRO} \leq C_1 \leq C_2$. This is straightforward since the feasible region enlarges when we set a higher budget. Together with the feasibility of the model $(\mathcal{DRO})$, we complete the proof. □

A.3 Proof of Proposition 3

Proof. The proof follows partly from the existing research, refer to Hall et al. (2015), Cui et al. (2023), and Long et al. (2023). Given a utility-based SLM defined by Definition 1, we define the set

$$\mathcal{K}(\tilde{\boldsymbol{x}}) = \left\{ k > 0 \,\middle|\, \mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m \in [M]} \left\{a_m v_{\mathcal{I}_t, \overline{\tau}_t}(\tilde{x}_t) + b_m k\right\}\right] \leq 0, \forall t \in [T], \hat{\mathbb{P}} \in \hat{\mathcal{P}} \right\} \tag{S.1}$$

1. *Monoticity.* If $v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) \geq v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}})$, then $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) \geq v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{y}_t), \forall t \in [T]$. For brevity, we would omit indices when the context is clear. By the monotonicity of $\mathbb{E}_{\hat{\mathbb{P}}}[u(\cdot)]$ for any $\hat{\mathbb{P}} \in \hat{\mathcal{P}}$, we have

$$\mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m\in[M]} \left\{a_m v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) + b_m k\right\}\right] \geq \mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m\in[M]} \left\{a_m v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{y}_t) + b_m k\right\}\right]$$

Hence, for any $k \in \mathcal{K}(\tilde{\boldsymbol{x}})$, the condition $k \in \mathcal{K}(\tilde{\boldsymbol{y}})$ must holds. In other words, $\mathcal{K}(\tilde{\boldsymbol{x}})$ is a subset of $\mathcal{K}(\tilde{\boldsymbol{y}})$. Finally, applying the infimum proves $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) \geq \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}})$.

2. *Convexity.* To prove the convexity, we first consider any $k^x \in \mathcal{K}(\tilde{\boldsymbol{x}})$ and $k^y \in \mathcal{K}(\tilde{\boldsymbol{y}})$, we define $k^\lambda = \lambda k^x + (1-\lambda)k^y, \lambda \in [0,1]$, where k^x and k^y satisfy the constraints in $\mathcal{K}(\tilde{\boldsymbol{x}})$ and $\mathcal{K}(\tilde{\boldsymbol{y}})$, respectively. Recall that $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) = \max\{x_t - \overline{\tau}_t, \underline{\tau}_t - x_t\}$, hence $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)$ is a convex function of $\tilde{x}_t$. Without loss of generality, the utility function is represented by $u(\cdot)$, therefore

$$\begin{aligned}
&\mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\overline{\tau}_t,\underline{\tau}_t}(\lambda\tilde{x}_t + (1-\lambda)\tilde{y}_t)}{k^\lambda}\right)\right] \\
&\leq \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{\lambda v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) + (1-\lambda)v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{y}_t)}{k^\lambda}\right)\right] \\
&= \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{\lambda k^x}{k^\lambda}\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k^x} + \frac{(1-\lambda)k^y}{k^\lambda}\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{y}_t)}{k^y}\right)\right] \\
&\leq \mathbb{E}_{\hat{\mathbb{P}}}\left[\mu u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k^x}\right) + (1-\mu)u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{y}_t)}{k^y}\right)\right] \\
&= \mu\mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k^x}\right)\right] + (1-\mu)\mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{y}_t)}{k^y}\right)\right] \\
&\leq 0
\end{aligned}$$

where $\mu = \lambda k^x/k^\lambda, \lambda \in [0,1]$. The first inequality holds due to the convexity of $v_{\underline{\tau}_t,\overline{\tau}_t}(\cdot)$. Then the second equation breaks the whole term into two separate parts. The second inequality holds since the utility function $u(\cdot)$ is convex. The last inequality is true for the definition of $\mathcal{K}(\tilde{\boldsymbol{x}})$ and $\mathcal{K}(\tilde{\boldsymbol{y}})$. Hence, we can conclude that for any $\lambda \in [0,1], k^x \in \mathcal{K}(\tilde{\boldsymbol{x}}), k^y \in \mathcal{K}(\tilde{\boldsymbol{y}})$, we have $k^\lambda \in \mathcal{K}(\lambda\tilde{\boldsymbol{x}} + (1-\lambda)\tilde{\boldsymbol{y}})$.

Therefore, when we take the convex combination of $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}})$ and $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}})$, we have

$$\begin{aligned}
&\lambda\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) + (1-\lambda)\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{y}}) \\
&= \lambda \inf\left\{k^x, k^x \in \mathcal{K}(\tilde{\boldsymbol{x}})\right\} + (1-\lambda)\inf\left\{k^y, k^y \in \mathcal{K}(\tilde{\boldsymbol{y}})\right\} \\
&= \inf\left\{\lambda k^x + (1-\lambda)k^y, k^x \in \mathcal{K}(\tilde{\boldsymbol{x}}), k^y \in \mathcal{K}(\tilde{\boldsymbol{y}})\right\} \\
&\geq \inf\left\{k^\lambda, k^\lambda \in \mathcal{K}(\lambda\tilde{\boldsymbol{x}} + (1-\lambda)\tilde{\boldsymbol{y}})\right\} \\
&= \rho_{\underline{\tau},\overline{\tau}}(\lambda\tilde{\boldsymbol{x}} + (1-\lambda)\tilde{\boldsymbol{y}})
\end{aligned}$$

The inequality is true due to $\lambda\mathcal{K}(\tilde{\boldsymbol{x}})$ is a convex function of $\tilde{\boldsymbol{x}}$, i.e., $\lambda\mathcal{K}(\tilde{\boldsymbol{x}})+(1-\lambda)\mathcal{K}(\tilde{\boldsymbol{y}}) \subseteq \mathcal{K}(\lambda\tilde{\boldsymbol{x}}+(1-\lambda)\tilde{\boldsymbol{y}})$.

3. *Positive homogeneity.* From the definition of $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) = \max\{x_t - \overline{\tau}_t, \underline{\tau}_t - x_t\}$, for all $\lambda > 0$, we have

$$v_{\lambda\underline{\tau}_t,\lambda\overline{\tau}_t}(\lambda\tilde{x}_t) = \max\{\lambda(x_t - \overline{\tau}_t), \lambda(\underline{\tau}_t - x_t)\} = \lambda v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)$$

Hence,

$$\begin{aligned}\rho_{\lambda\underline{\tau},\lambda\overline{\tau}}(\lambda\tilde{\boldsymbol{x}}) &= \inf\left\{k>0 \middle| \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\lambda\underline{\tau}_t,\lambda\overline{\tau}_t}(\lambda\tilde{x}_t)}{k}\right)\right] \le 0, \forall t\in[T], \hat{\mathbb{P}}\in\hat{\mathcal{P}}\right\}\\ &= \inf\left\{k>0 \middle| \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{\lambda v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k}\right)\right] \le 0, \forall t\in[T], \hat{\mathbb{P}}\in\hat{\mathcal{P}}\right\}\\ &= \inf\left\{k>0 \middle| \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k/\lambda}\right)\right] \le 0, \forall t\in[T], \hat{\mathbb{P}}\in\hat{\mathcal{P}}\right\}\\ &= \inf\left\{\lambda\beta>0 \middle| \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{\beta}\right)\right] \le 0, \forall t\in[T], \hat{\mathbb{P}}\in\hat{\mathcal{P}}\right\}\\ &= \lambda\inf\left\{\beta>0 \middle| \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{\beta}\right)\right] \le 0, \forall t\in[T], \hat{\mathbb{P}}\in\hat{\mathcal{P}}\right\}\\ &= \lambda\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}})\end{aligned}$$

where $\beta = k/\lambda$. The equation holds true after the notation substitution, as the notation itself does not affect the function.

4. *Pro-Robustness.* If $v_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) \le 0$, then $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) \le 0, \forall t\in[T]$. Recall that utility function $u(\cdot)$ is non-decreasing and normalized by $u(0)=0$, for any $k>0$ and $\hat{\mathbb{P}}\in\hat{\mathcal{P}}$, we observe that

$$\begin{aligned}\mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)\right)\right] &= \mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m\in[M]}\left\{a_m v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) + b_m k\right\}\right]\\ &\le \mathbb{E}_{\hat{\mathbb{P}}}\left[u(0)\right] = 0\end{aligned}$$

The RS constraint holds true almost surely. Hence, we have $\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}) = \inf\{k>0\} = 0$.

5. *Antifragility.* If there exists an $t\in[T]$ such that $v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) > 0$, then for any $k>0$ and $\hat{\mathbb{P}}\in\hat{\mathcal{P}}$, there must be

$$\mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m\in[M]}\left\{a_m v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) + b_m k\right\}\right] \ge \mathbb{E}_{\hat{\mathbb{P}}}\left[u(0)\right] = 0$$

which means the feasible set $\mathcal{K}(\tilde{\boldsymbol{x}})$ is an empty set since the t-th robust constraint is always violated. It is straightforward to see $\inf\mathcal{K}(\tilde{\boldsymbol{x}}) = \inf\emptyset = \infty$.

6. *Enveloping bound of violation probability.* First, we define the optimal SLM $k^* = \inf_{k\in\mathcal{K}(\boldsymbol{x})} k$. Then we define the set for any $t\in[T]$,

$$\mathcal{K}(\tilde{x}_t) = \left\{k>0 \;\middle|\; \mathbb{E}_{\hat{\mathbb{P}}}\left[\max_{m\in[M]}\left\{a_m v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) + b_m k\right\}\right] \le 0, \hat{\mathbb{P}}\in\hat{\mathcal{P}}\right\} \tag{S.2}$$

It is straightforward to see that $\mathcal{K}(\tilde{\boldsymbol{x}}) \subseteq \mathcal{K}(\tilde{x}_t)$ because any feasible solution in $\mathcal{K}(\tilde{\boldsymbol{x}})$ is always feasible in $\mathcal{K}(\tilde{x}_t)$. Let k_t^* be the optimal solution on $\mathcal{K}(\tilde{x}_t)$, we have

$$k^* \ge k_t^*, \quad t\in[T]$$

Now we construct the probability enveloping bound for inventory violation. Recall that the utility function is lower bounded by $\underline{u} \le 0$, thus we have $u(v) - \underline{u} \ge 0$. Then for $\forall\hat{\mathbb{P}}\in\hat{\mathbb{P}}, \forall\theta>0$,

$$\begin{aligned}\hat{\mathbb{P}}\left(v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t) \ge \theta\right) &= \hat{\mathbb{P}}\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k_t^*} \ge \frac{\theta}{k_t^*}\right)\\ &= \hat{\mathbb{P}}\left(u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k_t^*}\right) - \underline{u} \ge u\left(\frac{\theta}{k_t^*}\right) - \underline{u}\right)\\ &\le \frac{\mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_t,\overline{\tau}_t}(\tilde{x}_t)}{k_t^*}\right)\right] - \underline{u}}{u\left(\frac{\theta}{k_t^*}\right) - \underline{u}}\\ &\le \frac{-\underline{u}}{u\left(\frac{\theta}{k_t^*}\right) - \underline{u}}\\ &\le \frac{-\underline{u}}{u\left(\frac{\theta}{k^*}\right) - \underline{u}}\end{aligned}$$

The second equation holds since the utility function is non-decreasing. Then the first inequality holds based on Markov inequality and the fact that $u(v) - \underline{u} \geq 0$. The second inequality is true because of the definition of $\mathcal{K}(\tilde{x}_t)$, and the last inequality holds because of $k^* \geq k_t^*$ and $\underline{u} \leq 0$.

7. *Right continuity.* The proof of *Right continuity* is similar to that of Theorem 1 in Cui et al. (2023).

□

A.4 Proof of Theorem 1

Proof. For convenience, we define the feasible set of decisions $(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y})$ as $\mathcal{X}$ including all the constraints in $(\mathcal{RS})$. By the definition of utility-based probability distance Δ_u:

$$\Delta_u(\mathbb{P}, \hat{\mathbb{P}}) = \sup_{\tilde{v}\in\mathcal{V}} \left\{ \mathbb{E}_{\mathbb{P}}[\tilde{v}] \;\middle|\; \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m \tilde{v} + b_m\right] \leq 0, m \in [M] \right\}$$

We first derive the equivalence between the robust satisficing constraint using utility-based probability distance and the Service Level Measure.

$$\begin{aligned}
\kappa^* &= \min\left\{ k \geq 0 \;\middle|\; \mathbb{E}_{\mathbb{P}}\left[v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)\right] \leq k\Delta_u(\mathbb{P}, \hat{\mathbb{P}}), \forall t \in [T], \right.\\
&\qquad \left. i \in [N], \mathbb{P} \in \mathcal{P}_0, \hat{\mathbb{P}} \in \hat{\mathcal{P}}, (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}\\
&= \inf\left\{ k > 0 \;\middle|\; \mathbb{E}_{\mathbb{P}}\left[\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}}{k}\right] - \Delta_u(\mathbb{P}, \hat{\mathbb{P}}) \leq 0, \forall t \in [T], \right.\\
&\qquad \left. i \in [N], \mathbb{P} \in \mathcal{P}_0, \hat{\mathbb{P}} \in \hat{\mathcal{P}}, (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}\\
&= \inf\left\{ k > 0 \;\middle|\; \sup_{\mathbb{P}\in\mathcal{P}_0} \left\{ \mathbb{E}_{\mathbb{P}}\left[\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)}{k}\right] - \Delta_u(\mathbb{P}, \hat{\mathbb{P}}) \right\} \leq 0, \right.\\
&\qquad \left. \forall t \in [T], i \in [N], \hat{\mathbb{P}} \in \hat{\mathcal{P}}, (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}
\end{aligned} \tag{S.3}$$

According to Föllmer and Schied (2002), the utility-based probability distance is the dual form of the shortfall risk measure:

$$\mu^u_{\hat{\mathbb{P}}}(\tilde{v}) = \sup_{\mathbb{P}\in\mathcal{P}_0} \{\mathbb{E}_{\mathbb{P}}[\tilde{v}] - \Delta_u(\mathbb{P}, \hat{\mathbb{P}})\}$$

Hence, we can rewrite (S.3) by definition of the shortfall risk measure as:

$$\begin{aligned}
\kappa^* &= \inf\left\{ k > 0 \;\middle|\; \mu^u_{\hat{\mathbb{P}}}\left(\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)}{k}\right) \leq 0, \forall t \in [T], i \in [N], \right.\\
&\qquad \left. \hat{\mathbb{P}} \in \hat{\mathcal{P}}, (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}\\
&= \inf\left\{ k > 0 \;\middle|\; \inf\left\{ a \;\middle|\; \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)}{k} - a\right)\right] \leq 0 \right\} \leq 0, \right.\\
&\qquad \left. \forall t \in [T], i \in [N], \hat{\mathbb{P}} \in \hat{\mathcal{P}}, (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}\\
&= \inf\left\{ k > 0 \;\middle|\; \exists a \leq 0, \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)}{k} - a\right)\right] \leq 0 \right.\\
&\qquad \left. \forall t \in [T], i \in [N], \hat{\mathbb{P}} \in \hat{\mathcal{P}}, (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}\\
&= \inf\left\{ k > 0 \;\middle|\; \sup_{\hat{\mathbb{P}}\in\hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)}{k}\right)\right] \leq 0, \right.\\
&\qquad \left. \forall t \in [T], i \in [N], (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \right\}\\
&= \inf\left\{ k > 0 \;\middle|\; \sup_{\hat{\mathbb{P}}\in\hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(\tilde{x}_i^t) + b_m k\right] \leq 0, \right.
\end{aligned}$$

$$\forall t \in [T], i \in [N], m \in [M], (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \bigg\}$$

The first equality and the second equality hold, through Föllmer and Schied (2002)'s dual representation. Then, the third equality holds since the utility function $u(\cdot)$ is non-decreasing: There exists $a \geq 0$ such that $\mathbb{E}_{\hat{\mathbb{P}}}\left[u\left(\frac{v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(x_i^t)}{k} - a\right)\right] \leq 0$. The reverse is trivial as long as $a = 0$. The last equality holds due to the definition of the utility function.

Next, we show the equivalence between Problem ($\mathcal{RS}$) and Problem ($\mathcal{RS} - \mathcal{SLM}$). The next lemma states that the LHS in the robust satisficing constraint is non-decreasing in k, which implies that when $k' > k$, the constraint still holds.

Lemma S1. *Let $k > 0$ be feasible to the constraint*

$$\sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m v_{\underline{\tau}_t, \overline{\tau}_t}(\tilde{x}_t) + b_m k\right] \leq 0, \forall t \in [T], m \in [M]$$

For any $k' > k$, we must have

$$\sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m v_{\underline{\tau}_t, \overline{\tau}_t}(\tilde{x}_t) + b_m k'\right] \leq 0, \forall t \in [T], m \in [M]$$

Since $k > 0$, the constraint can be rewritten as:

$$\sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m \frac{v_{\underline{\tau}_t, \overline{\tau}_t}(\tilde{x}_t)}{k} + b_m\right] \leq 0, \forall t \in [T], m \in [M]$$

Given $k' > k$ and $a_m \geq 0$, we must have $0 \leq a_m \frac{v_{\underline{\tau}_t, \overline{\tau}_t}(\tilde{x}_t)}{k'} \leq a_m \frac{v_{\underline{\tau}_t, \overline{\tau}_t}(\tilde{x}_t)}{k}$. Therefore, for any $k' > k$, the following constraint holds because taking the expectation does not affect the inequality:

$$\sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m \frac{v_{\underline{\tau}_t, \overline{\tau}_t}(\tilde{x}_t)}{k'} + b_m\right] \leq 0, \forall t \in [T], m \in [M]$$

This completes the proof of Lemma S1.

By Lemma S1, we can equivalently minimize the maximum of individual k_i as follows:

$$\begin{aligned} \inf \quad & \max\{k_1, k_2, \ldots, k_N\} \\ \text{s.t.} \quad & \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}}\left[a_m v_{\underline{\tau}_i^t, \overline{\tau}_i^t}(\tilde{x}_i^t) + b_m k_i\right] \leq 0 \\ & \qquad \forall t \in [T], i \in [N], m \in [M] \\ & (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \end{aligned} \tag{S.4}$$

Let κ^* and $\{k_i^*\}_{i \in [N]}$ denote the optimal values for Problem ($\mathcal{RS}$) and Problem (S.4). The direction (S.4) $\Rightarrow$ ($\mathcal{RS}$) is straightforward by adding auxiliary $k \geq k_i, \forall i \in [N]$. By the fact that the LHS in robust satisficing constraint is non-decreasing in k, k satisfies the constraints in ($\mathcal{RS}$) and hence is a feasible solution to Problem ($\mathcal{RS}$). Thus we have $\max_{i \in [N]}\{k_i^*\} = k \geq \kappa^*$. To see the direction ($\mathcal{RS}$) $\Rightarrow$ (S.4), we consider any feasible solution k in ($\mathcal{RS}$), we can construct a feasible solution of (S.4) by setting $k_i = k, \forall i \in [N]$ such that $\kappa^* \geq \max_{i \in [N]}\{k_i^*\}$. Combining two inequalities, we have $\kappa^* = \max_{i \in [N]}\{k_i^*\}$. Finally, with the definition of SLM, we can reduce the notations to:

$$\begin{aligned} \inf \quad & \max\{\rho(\tilde{\boldsymbol{x}}_1), \rho(\tilde{\boldsymbol{x}}_2), \ldots, \rho(\tilde{\boldsymbol{x}}_N)\} \\ \text{s.t.} \quad & (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \end{aligned}$$

□

A.5 The proof of Theorem 2

Proof. With the definition of the lifted ambiguity set, we rewrite the explicit form of Problem ($\mathcal{RS} - \mathcal{SLM}$) as follows:

$$\begin{aligned} \kappa^* = \inf_{k_i > 0} \quad & \max_{i \in [N]}\{k_i\} \\ \text{s.t.} \quad & x_i^{ts}(\cdot) = x_i^{1s} + \sum_{l=1}^{t} q_i^{(l-L_{0i})s}(\cdot) + \sum_{j \in \Gamma(i)} \sum_{l=1}^{t} y_{ji}^{(l-L_{ij})s}(\cdot) \end{aligned}$$

$$
\begin{aligned}
& - \sum_{j\in\Gamma(i)} \sum_{l=1}^{t} y_{ij}^{ls}(\cdot) - \sum_{l=1}^{t} u_i^{ls}(\cdot) && \forall t\in[T], i\in[N], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] && \text{(S.5a)}\\
& x_0^{ts}(\cdot) = x_0^{1s} + \sum_{i\in[N]}\sum_{l=1}^{t} u_i^{(l-\Delta_i)s}(\cdot) - \sum_{i\in[N]}\sum_{l=1}^{t} q_i^{ls}(\cdot) && \forall t\in[T], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] && \text{(S.5b)}\\
& \sum_{j\in\Gamma(i)} y_{ij}^{ts}(\cdot) \le x_i^{ts}(\cdot) && \forall t\in[T], i\in[N], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] && \text{(S.5c)}\\
& u_i^{ts}(\cdot) \le d_i^{ts} && \forall t\in[T], i\in[N], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] && \text{(S.5d)}\\
& \sup_{\hat{\mathbb{Q}}\in\hat{\mathcal{F}}} \mathbb{E}_{\hat{\mathbb{Q}}}\left[a_m v_{\underline{\tau}_i^t,\overline{\tau}_i^t}(x_i^{t\tilde{s}}(\cdot)) + b_m k_i\right] \le 0 && \forall t\in[T], i\in[N], m\in[M] && \text{(S.5e)}\\
& \sup_{\hat{\mathbb{Q}}\in\hat{\mathcal{F}}} \mathbb{E}_{\hat{\mathbb{Q}}}\Bigg[K(x_0^{1\tilde{s}} + \sum_{i\in[N]} x_i^{1\tilde{s}}) + \sum_{t\in[T]}\sum_{i\in[N]} c_t q_i^{t\tilde{s}}(\cdot) \\
& \qquad \sum_{t\in[T]}\sum_{i\in[N]}\Big(\sum_{j\in\Gamma(i)} p_{ij}^t y_{ij}^{t\tilde{s}}(\cdot) + b(d_i^{t\tilde{s}} - u_i^{t\tilde{s}}(\cdot))\Big)\Bigg] \le C && && \text{(S.5f)}\\
& \boldsymbol{x}^s(\cdot), \boldsymbol{q}^s(\cdot), \boldsymbol{y}^s(\cdot), \boldsymbol{u}^s(\cdot) \ge \boldsymbol{0} && \forall (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] && \text{(S.5g)}\\
& \boldsymbol{x}^s(\cdot), \boldsymbol{q}^s(\cdot), \boldsymbol{y}^s(\cdot), \boldsymbol{u}^s(\cdot) \in \mathcal{A} && \forall s\in[S] && \text{(S.5h)}
\end{aligned}
$$

We consider the robust satisficing constraint (S.5e) in Problem (S.5). By the law of total probability, we rewrite it by $\hat{\mathbb{Q}} = \sum_{s\in[S]} p_s \hat{\mathbb{Q}}_s$:

$$
\sum_{s\in[S]} p_s \sup_{\hat{\mathbb{Q}}_s\in\hat{\mathcal{F}}_s} \mathbb{E}_{\hat{\mathbb{Q}}_s}\left[a_m v_{\underline{\tau}_i^t,\overline{\tau}_i^t}(x_i^{t\tilde{s}}(\cdot)) + b_m k_i\right] \le 0 \quad \forall t\in[T], i\in[N], m\in[M] \tag{S.6}
$$

The left term is a semi-infinite maximization problem. Given a specific $t\in[T]$ and $i\in[N]$, we start with the equivalent moment problem:

$$
\begin{aligned}
\sup \sum_{s\in[S]} \quad & p_s \mathbb{E}_{\hat{\mathbb{Q}}_s}\left[\max_{m\in[M]}\left\{a_m v_{\underline{\tau}_i^t,\overline{\tau}_i^t}(x_i^t(\cdot)) + b_m k_i\right\}\right] \\
\text{s.t.} \quad & \hat{\mathbb{Q}}_s\left[(\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s\right] = 1, \quad \forall s\in[S], \quad \Rightarrow \text{dual variable } o_i^{ts}\in\mathbb{R} \\
& \mathbb{E}_{\hat{\mathbb{Q}}_s}[\boldsymbol{d}] = \boldsymbol{\mu}^s, \quad \forall s\in[S], \quad \Rightarrow \text{dual variable } \boldsymbol{z}_i^{ts}\in\mathbb{R}^{NT} \\
& \mathbb{E}_{\hat{\mathbb{Q}}_s}[\boldsymbol{w}] \le \boldsymbol{\sigma}^s, \quad \forall s\in[S], \quad \Rightarrow \text{dual variable } \boldsymbol{\beta}_i^{ts}\in\mathbb{R}_+^{NT} \\
& \mathbb{E}_{\hat{\mathbb{Q}}_s}[\boldsymbol{v}] \le \boldsymbol{\epsilon}^s, \quad \forall s\in[S], \quad \Rightarrow \text{dual variable } \boldsymbol{\gamma}_i^{ts}\in\mathbb{R}_+^{HT}
\end{aligned} \tag{S.7}
$$

The strong duality holds according to Wiesemann et al. (2014) and Bertsimas et al. (2019), and we associate constraints of the problem with dual variables $o_i^{ts}\in\mathbb{R}$, $\boldsymbol{z}_i^{ts}\in\mathbb{R}^{NT}$, $\boldsymbol{\beta}_i^{ts}\in\mathbb{R}_+^{NT}$ and $\mathbb{R}_+^{HT}$. Applying the moment duality techniques, we can obtain the dual problem:

$$
\begin{aligned}
\inf \quad & \sum_{s\in[S]} o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{\mu}^s + \boldsymbol{\beta}_i^{ts}\boldsymbol{\sigma}^s + \boldsymbol{\gamma}_i^{ts}\boldsymbol{\epsilon}^s \\
\text{s.t.} \quad & o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge p_s \max_{m\in[M]}\left\{a_m v_{\underline{\tau}_i^t,\overline{\tau}_i^t}(x_i^{ts}(\cdot)) + b_m k_i\right\} \quad \forall (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] \\
& o_i^{ts}\in\mathbb{R},\ \boldsymbol{z}_i^{ts}\in\mathbb{R}^{NT}, \boldsymbol{\beta}_i^{ts}\in\mathbb{R}_+^{NT}, \boldsymbol{\gamma}_i^{ts}\in\mathbb{R}_+^{HT}
\end{aligned} \tag{S.8}
$$

Observe that the first constraint involves a piece-wise linear function, hence, we have equivalently:

$$
o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge p_s\left[a_m v_{\underline{\tau}_i^t,\overline{\tau}_i^t}(x_i^{ts}(\cdot)) + b_m k_i\right] \quad \forall m\in[M], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] \tag{S.9}
$$

By the definition of $v_{\underline{\tau}_i^t,\overline{\tau}_i^t} = \max\{x_i^t - \overline{\tau}_i^t, \underline{\tau}_i^t - x_i^t\}$, the constraint (S.9) is reformulated to two sets of constraints:

$$
\begin{aligned}
o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge \quad & p_s\left[a_m(x_i^{ts}(\cdot) - \overline{\tau}_i) + b_m k_i\right] \\
& \forall m\in[M], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S] \\
o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge \quad & p_s\left[a_m(\underline{\tau}_i - x_i^{ts}(\cdot)) + b_m k_i\right] \\
& \forall m\in[M], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S]
\end{aligned} \tag{S.10}
$$

Inserting the above constraints (S.10) into the dual problem (S.8), and combining with the primal constraint (S.6), we have:

$$\begin{aligned}
&\sum_{s\in[S]} o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{\mu}^s + \boldsymbol{\beta}_i^{ts}\boldsymbol{\sigma}^s + \boldsymbol{\gamma}_i^{ts}\boldsymbol{\epsilon}^s \le 0 \qquad \forall t\in[T], i\in[N]\\
&o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge p_s\left[a_m(x_i^{ts}(\cdot)-\overline{\tau}_i)+b_m k_i\right]\\
&\qquad\qquad \forall t\in[T], i\in[N], m\in[M], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S]\\
&o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge p_s\left[a_m(\underline{\tau}_i - x_i^{ts}(\cdot))+b_m k_i\right]\\
&\qquad\qquad \forall t\in[T], i\in[N], m\in[M], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S]\\
&o_i^{ts}\in\mathbb{R},\ \boldsymbol{z}_i^{ts}\in\mathbb{R}^{NT}, \boldsymbol{\beta}_i^{ts}\in\mathbb{R}_+^{NT}, \boldsymbol{\gamma}_i^{ts}\in\mathbb{R}_+^{HT} \quad \forall t\in[T], i\in[N], s\in[S]
\end{aligned} \tag{S.11}$$

By the similar reformulation, we rewrite the constraint (S.5f) as:

$$\begin{aligned}
&\sum_{s\in[S]} o_s + \boldsymbol{z}_s\boldsymbol{\mu}^s + \boldsymbol{\beta}_s\boldsymbol{\sigma}^s + \boldsymbol{\gamma}_s\boldsymbol{\epsilon}^s \le C\\
&o_s + \boldsymbol{z}_s\boldsymbol{d} + \boldsymbol{\beta}_s\boldsymbol{w} + \boldsymbol{\gamma}_s\boldsymbol{v} \ge p_s\left[\sum_{i\in[N]\cup 0} K x_i^{1s} + \sum_{i=1}^{N}\sum_{t=1}^{T}\left(c_t q_i^{ts}(\cdot) + b_t(d_i^{ts} - u_i^{ts}(\cdot)) + \sum_{j\in\Gamma(i)} p_{ij}^t y_{ji}^{ts}(\cdot)\right)\right]\\
&\qquad\qquad \forall\,(\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S]\\
&o_s\in\mathbb{R},\ \boldsymbol{z}_s\in\mathbb{R}^{NT}, \boldsymbol{\beta}_s\in\mathbb{R}_+^{NT}, \boldsymbol{\gamma}_s\in\mathbb{R}_+^{HT}
\end{aligned} \tag{S.12}$$

Inserting the result constraints into Problem (S.5), we obtain the result ARO problem. □

A.6 Proof of Proposition 4

Proof. For any $t\in[T], i\in[N], s\in[S]$ and $m\in[M]$, the reformulation is independent. We focus on the following constraint:

$$o_i^{ts} + \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} \ge p_s\left[a_m(x_i^{ts}(\cdot)-\overline{\tau}_i)+b_m k_i\right] \quad \forall\,(\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S]$$

For adaptive decisions, we replace adaptive decision $x_i^{ts}(\cdot)$ by explicit ELDR:

$$x_i^{ts}(\cdot) = x_i^{ts0} + \sum_{\substack{i'\in\Gamma(i),\\ l\in[t-1]}} x_{ii'l}^{ts1} d_{i'}^l + \sum_{\substack{i'\in\Gamma(i),\\ l\in[t-1]}} x_{ii'l}^{ts2} w_{i'}^l + \sum_{\substack{h\in\mathcal{H}(i),\\ l\in[t-1]}} x_{ihl}^{ts3} v_h^l$$

By rearranging the terms in the equation, we have

$$\begin{aligned}
&\boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} - p_s a_m\left[\sum_{\substack{i'\in\Gamma(i),\\ l\in[t-1]}} x_{ii'l}^{ts} d_{i'}^l + \sum_{\substack{i'\in\Gamma(i),\\ l\in[t-1]}} x_{ii'l}^{ts2} w_{i'}^l + \sum_{\substack{h\in\mathcal{H}(i),\\ l\in[t-1]}} x_{ihl}^{ts3} v_h^l\right]\\
&\qquad \ge p_s a_m(x_i^{ts0}-\overline{\tau}_i) + p_s b_m k_i - o_i^{ts} \qquad \forall t\in[T], i\in[N], m\in[M], (\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})\in\bar{\mathcal{Z}}_s, s\in[S]
\end{aligned}$$

It remains to solve the minimization subproblem on the left-hand side. Inserting the lifted support set, we now have:

$$\begin{aligned}
\inf_{(\boldsymbol{d},\boldsymbol{w},\boldsymbol{v})}\ & \boldsymbol{z}_i^{ts}\boldsymbol{d} + \boldsymbol{\beta}_i^{ts}\boldsymbol{w} + \boldsymbol{\gamma}_i^{ts}\boldsymbol{v} - p_s a_m\left[\sum_{\substack{i'\in\Gamma(i),\\ l\in[t-1]}} x_{ii'l}^{ts1} d_{i'}^l + \sum_{\substack{i'\in\Gamma(i),\\ l\in[t-1]}} x_{ii'l}^{ts2} w_{i'}^l + \sum_{\substack{h\in\mathcal{H}(i),\\ l\in[t-1]}} x_{ihl}^{ts3} v_h^l\right]\\
\text{s.t.}\quad & \boldsymbol{d} \ge \underline{\boldsymbol{d}}^s && \Rightarrow \text{dual variable } \boldsymbol{\pi}_{i1}^{tsm}\in\mathbb{R}_+^{NT}\\
& -\boldsymbol{d} \ge -\overline{\boldsymbol{d}}^s && \Rightarrow \text{dual variable } \boldsymbol{\pi}_{i2}^{tsm}\in\mathbb{R}_+^{NT}\\
& \boldsymbol{w} + \boldsymbol{d} \ge \boldsymbol{\mu}^s && \Rightarrow \text{dual variable } \boldsymbol{\pi}_{i3}^{tsm}\in\mathbb{R}_+^{NT}\\
& \boldsymbol{w} - \boldsymbol{d} \ge -\boldsymbol{\mu}^s && \Rightarrow \text{dual variable } \boldsymbol{\pi}_{i4}^{tsm}\in\mathbb{R}_+^{NT}\\
& v_h^l + \sum_{i'\in\mathcal{N}_h}\frac{d_{i'}^l}{\sigma_{i'}^{ls}} \ge \sum_{i'\in\mathcal{N}_h}\frac{\mu_{i'}^{ls}}{\sigma_{i'}^{ls}} \quad \forall h\in[H], l\in[T] && \Rightarrow \text{dual variable } \eta_{ihl}^{tsm}\in\mathbb{R}_+\\
& v_h^l - \sum_{i'\in\mathcal{N}_h}\frac{d_{i'}^l}{\sigma_{i'}^{ls}} \ge -\sum_{i'\in\mathcal{N}_h}\frac{\mu_{i'}^{ls}}{\sigma_{i'}^{ls}} \quad \forall h\in[H], l\in[T] && \Rightarrow \text{dual variable } \phi_{ihl}^{tsm}\in\mathbb{R}_+
\end{aligned}$$

By strong duality (Bertsimas et al., 2019), we have:

$$\begin{aligned}
\sup_{\substack{(\boldsymbol{\pi}_{i1}^{tsm},\boldsymbol{\pi}_{i2}^{tsm},\boldsymbol{\pi}_{i3}^{tsm},\\ \boldsymbol{\pi}_{i4}^{tsm},\boldsymbol{\eta}_i^{tsm},\boldsymbol{\phi}_i^{tsm})\geq \mathbf{0}}} \quad & \boldsymbol{\pi}_{i1}^{tsm}\underline{\boldsymbol{d}}^s - \boldsymbol{\pi}_{i2}^{tsm}\overline{\boldsymbol{d}}^s + (\boldsymbol{\pi}_{i3}^{tsm} - \boldsymbol{\pi}_{i4}^{tsm})\boldsymbol{\mu}^s + \sum_{h\in[H]}\sum_{l=1}^{t}\sum_{i'\in\mathcal{N}_h}\frac{\mu_{i'}^{ls}}{\sigma_{i'}^{ls}}(\eta_{ihl1}^{tsm} - \phi_{ihl1}^{tsm}) \\
\text{s.t.} \quad & \pi_{ii'l1}^{tsm} - \pi_{ii'l2}^{tsm} + \pi_{ii'l3}^{tsm} - \pi_{ii'l4}^{tsm} + \frac{1}{\sigma_{i'}^{ls}}\sum_{h'\in\mathcal{H}(i')}(\eta_{ihl1}^{tsm} - \phi_{ihl1}^{tsm}) \\
& \qquad \leq z_{ii'l}^{ts} - p_s a_m x_{ii'l}^{ts1} \qquad \forall i' \in \Gamma(i), l \in [t-1] \qquad \text{(S.13a)} \\
& \pi_{ii'l1}^{tsm} - \pi_{ii'l2}^{tsm} + \pi_{ii'l3}^{tsm} - \pi_{ii'l4}^{tsm} + \frac{1}{\sigma_{i'}^{ls}}\sum_{h'\in\mathcal{H}(i')}(\eta_{ihl1}^{tsm} - \phi_{ihl1}^{tsm}) \\
& \qquad \leq z_{ii'l}^{ts} \qquad \forall (i',l) \in ([N]\times[T]) \setminus (\Gamma(i)\times[t-1]) \qquad \text{(S.13b)} \\
& \pi_{ii'l3}^{tsm} + \pi_{ii'l4}^{tsm} \leq \beta_{ii'l}^{ts} - p_s a_m x_{ii'l}^{ts2} \qquad \forall i' \in \Gamma(i), l \in [t-1] \qquad \text{(S.13c)} \\
& \pi_{ii'l3}^{tsm} + \pi_{ii'l4}^{tsm} \leq \beta_{ii'l}^{ts} \qquad \forall (i',l) \in ([N]\times[T]) \setminus (\Gamma(i)\times[t-1]) \qquad \text{(S.13d)} \\
& \eta_{ihl}^{tsm} + \phi_{ihl}^{tsm} \leq \gamma_{ihl}^{ts} - p_s a_m x_{ihl}^{ts3} \qquad \forall h \in \mathcal{H}(i), l \in [t-1] \qquad \text{(S.13e)} \\
& \eta_{ihl}^{tsm} + \phi_{ihl}^{tsm} \leq \gamma_{ihl}^{ts} \qquad \forall (h,l) \in ([H]\times[T]) \setminus (\mathcal{H}(i)\times[t-1]) \qquad \text{(S.13f)} \\
& \boldsymbol{\pi}_{i1}^{tsm}, \boldsymbol{\pi}_{i2}^{tsm}, \boldsymbol{\pi}_{i3}^{tsm}, \boldsymbol{\pi}_{i4}^{tsm} \in \mathbb{R}_+^{NT}, \boldsymbol{\eta}_i^{tsm}, \boldsymbol{\phi}_i^{tsm} \in \mathbb{R}_+^{HT} \qquad \text{(S.13g)}
\end{aligned}$$

We emphasize the index set of each constraint, which inherently follows the rules in affine functions $\mathcal{A}$. Consequently, for any $t \in [T], i \in [N], s \in [S]$ and $m \in [M]$, the RS constraint is equivalent to the following constraint:

$$\begin{cases}
\boldsymbol{\pi}_{i1}^{tsm}\underline{\boldsymbol{d}}^s - \boldsymbol{\pi}_{i2}^{tsm}\overline{\boldsymbol{d}}^s + (\boldsymbol{\pi}_{i3}^{tsm} - \boldsymbol{\pi}_{i4}^{tsm})\boldsymbol{\mu}^s + \sum_{h\in[H]}\sum_{l=1}^{t}\sum_{i'\in\mathcal{N}_h}\frac{\mu_{i'}^{ls}}{\sigma_{i'}^{ls}}(\eta_{ihl}^{tsm} - \phi_{ihl}^{tsm}) \\
\qquad\qquad \geq p_s a_m (x_i^{ts0} - \overline{\tau}_i) + p_s b_m k_i - o_i^{ts} \\
\text{(S.13a)} - \text{(S.13g)}
\end{cases}$$

□

B Extensions

B.1 Pareto optimization Procedure

The Problem ($\mathcal{RS} - \mathcal{SLM}$) minimizes the worst-off SLM over N swapping stations to address the fairness issue. It may yield multiple optimal solutions, but some may not be Pareto optimal. For better understanding, we provide an example. Consider $N = 3$ and two sets of individual SLMs given by:

$$\boldsymbol{\rho}_1^* = \{6, 5, 4\}, \quad \boldsymbol{\rho}_2^* = \{6, 3, 2\}$$

These two solutions are optimal in the sense of minimizing the worst SLM, which is 6. Nevertheless, $\boldsymbol{\rho}_1^*$ is not Pareto optimal since we can strictly decrease one ρ_i^* without increasing the other. Therefore, we present in Algorithm (1) an iterative algorithm to obtain Pareto optimal solutions, adapted from lexicographic optimization by Qi (2017) and Li et al. (2024). At each iteration, we minimize the worst SLM over set $\mathcal{L}_n$ while imposing upper bounds on individual SLM $\{\rho_{\underline{\boldsymbol{\tau}},\overline{\boldsymbol{\tau}}}(\tilde{\boldsymbol{x}}_i)\}_{i\in[N]}$. These upper bounds are reasonable due to Lemma S1, and they are optimal values obtained from the previous iterations. In iteration n', the SLM for stations in set $\mathcal{I}_{n'}$ cannot exceed the optimal value $\rho_{n'}^*$ attained. For stations in $\mathcal{L}_n$, we require SLM not to exceed the ρ_n^*. If the optimal SLM ρ_{n+1}^* is attained, we again find stations with SLM ρ_{n+1}^* and add them into the set $\mathcal{I}_{n+1}$. Then $\mathcal{L}_{n+1}$ is obtained by subtracting $\mathcal{I}_{n+1}$ from $\mathcal{L}_n$. A new problem is further solved over $\mathcal{L}_{n+1}$. This procedure repeated until $\mathcal{L}_{n+1} = \emptyset$, and the optimal SLM satisfies:

$$\rho_{\underline{\boldsymbol{\tau}},\overline{\boldsymbol{\tau}}}(\tilde{\boldsymbol{x}}_i) = \rho_{n'}^*, i \in \mathcal{I}_{n'}, n' \in [n+1]$$

To guarantee the closure of the feasible region, we need a technical condition that $\rho_{\underline{\boldsymbol{\tau}},\overline{\boldsymbol{\tau}}}(\tilde{\boldsymbol{x}}_i) \geq \epsilon$ to replace $\rho_{\underline{\boldsymbol{\tau}},\overline{\boldsymbol{\tau}}}(\tilde{\boldsymbol{x}}_i) > 0$. This condition shall not compromise the optimality since a sufficiently small positive ϵ can always be selected (Chen et al., 2015; Chow et al., 2022). This procedure keeps the updated optimal SLM $\{\rho_{\underline{\boldsymbol{\tau}},\overline{\boldsymbol{\tau}}}(\tilde{\boldsymbol{x}}_i)\}_{i\in[N]}$ no worse than that of the last iteration. Since the original optimal solution $(\boldsymbol{x}_n^*, \boldsymbol{u}_n^*, \boldsymbol{q}_n^*, \boldsymbol{y}_n^*)$ is always feasible for the updated problem, see Step 3. It is natural to solve the new problem from the solution $(\boldsymbol{x}_n^*, \boldsymbol{u}_n^*, \boldsymbol{q}_n^*, \boldsymbol{y}_n^*)$ instead of solving from scratch. In practice, we can preserve the current optimal basis and apply the simplex algorithm for LP (see Chapter 5 in Bertsimas and Tsitsiklis (1997)). This typically saves many iterations, thus accelerating the computation.

Algorithm 1 Pareto Optimization Procedure

1: Set $\epsilon > 0, n := 0, \mathcal{L}_0 := [N], \mathcal{I}_0 := \emptyset, \rho_0^* := \infty$ and any feasible solution $(\boldsymbol{x}_0^*, \boldsymbol{u}_0^*, \boldsymbol{q}_0^*, \boldsymbol{y}_0^*)$.
2: **while** $\mathcal{L}_n \neq \emptyset$ **do**
3: Solve the following problem, starting from solution $(\boldsymbol{x}_n^*, \boldsymbol{u}_n^*, \boldsymbol{q}_n^*, \boldsymbol{y}_n^*)$.

$$\begin{aligned} \rho_{n+1}^* = \inf \quad & \max_{i \in \mathcal{L}_n} \{\rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i)\} \\ \text{s.t.} \quad & \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i) \leq \rho_{n'}^*, \qquad i \in \mathcal{I}_{n'}, n' \in [n] \cup 0 \\ & \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i) \leq \rho_{n}^*, \quad 0 \quad i \in \mathcal{L}_n \\ & \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_i) \geq \epsilon, \qquad i \in [N] \\ & (\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{q}, \boldsymbol{y}) \in \mathcal{X} \end{aligned} \tag{S.14}$$

4: Let ρ_{n+1}^* and $(\boldsymbol{x}_{n+1}^*, \boldsymbol{u}_{n+1}^*, \boldsymbol{q}_{n+1}^*, \boldsymbol{y}_{n+1}^*)$ be optimal value and minimizer of Problem (S.14).
5: Set $\mathcal{I}_{n+1} = \{j \in \mathcal{L}_n | \rho_{\underline{\tau},\overline{\tau}}(\tilde{\boldsymbol{x}}_j) := \rho_{n+1}^*\}$ and $\mathcal{L}_{n+1} := \mathcal{L}_n \backslash \mathcal{I}_{n+1}$.
6: Increase $n := n + 1$.
7: **end while**
8: **return** $(\boldsymbol{x}_n^*, \boldsymbol{u}_n^*, \boldsymbol{q}_n^*, \boldsymbol{y}_n^*)$, and $\{\rho_{\underline{\tau},\overline{\tau}}(\boldsymbol{x}_i^*)\}_{i \in [N]}$

Qi (2017) proposed a lexicographic minimization algorithm to sequentially minimize the maximum *Delay Unpleasantness Measure* over N participants. We tailor their method in two ways: First, our proposed SLM preserves the model's linear structure, and hence, we can apply LP sensitivity analysis techniques to obtain a new optimal solution rather than solving the new problem from scratch (i.e., Step 3). This is more computationally efficient and appropriate for our problem. Second, compared to the subproblem solved in Qi (2017), we additionally impose upper bounds on SLM of $\mathcal{L}_n$ in Problem (S.14). Adding the constraints does not affect the optimality while enhancing the bounds.

B.2 The Relation to Budget-Driven Model

Hu et al. (2024) proposed a budget-driven DRO model for the multi-period hub location problem with uncertain periodic demands. They constrain each expected periodic cost within a budget and maximize robustness by maximizing the size of the ambiguity set. Our work differentiates from theirs in two ways. First, we focus on limiting the total costs under a budget instead of periodic costs. Second, the budget-driven DRO model is based on the DRO framework, whereas our model is proposed under the RS framework. Given the periodic operational costs $\boldsymbol{C} = (C_t)_{t \in [N]}$, we can constrain the expected cost in each period to be less than the budget target.

$$\sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}} \mathbb{E}_{\hat{\mathbb{P}}} \left[\sum_{i \in [N]} \left(c_t q_i^t + b(d_i^t - u_i^t) + \sum_{j \in [N], j \neq i} p_{ij}^t y_{ji}^t \right) \right] \leq C_t \quad \forall t \in [T]$$

where we limit the periodic operating cost. We can tune the operational budget for each period by $\sum_{t \in [T]} C_t = C - K(x_0 + \sum_{i \in [N]} x_i^0)$, where the initial battery investment cost $K(x_0 + \sum_{i \in [N]} x_i^0)$ is pre-determined before demand arrives. C_t for any period t may be considered uncertain and independent, see details in Hu et al. (2024). Moreover, we can modify our model to maximize the robustness following the same vein in Hu et al. (2024). Suppose the ambiguity set $\hat{\mathcal{P}}$ relates to the size parameter $r \geq 0$, e.g., probability-distance-based ambiguity set $\hat{\mathcal{P}}(r)$. We identify the maximal size of the ambiguity set concerning which the worst-case inventory meets the target. Hence, the budget-driven model ($\mathcal{BDO}$) can be written as:

$$\begin{aligned} \max_{r \geq 0} \quad & r \\ \text{s.t.} \quad & \text{Constraints } (1) - (5) \\ & \mathbb{E}_{\hat{\mathbb{P}}} \left[v_{\underline{\tau}_i, \overline{\tau}_i}(x_i^t) \right] \leq 0 \qquad \forall t \in [T], i \in [N], \hat{\mathbb{P}} \in \hat{\mathcal{P}}(r) \\ & \sup_{\hat{\mathbb{P}} \in \hat{\mathcal{P}}(r)} \mathbb{E}_{\hat{\mathbb{P}}} \left[\sum_{i \in [N]} \left(c_t q_i^t + b(d_i^t - u_i^t) + \sum_{j \in [N], j \neq i} p_{ij}^t y_{ji}^t \right) \right] \leq C_t \quad \forall t \in [T] \end{aligned} \tag{$\mathcal{BDO}$}$$

Both ($\mathcal{RS}$) and ($\mathcal{BDO}$) essentially maximize the model's robustness. ($\mathcal{RS}$) achieves this by minimizing the fragility by limiting the statistical distance the true distribution deviates from the reference distribution. Instead, ($\mathcal{BDO}$) achieves this by maximizing the ambiguity set size, which is defined by the deviation from the reference distribution. Intuitively, ($\mathcal{RS}$) maximizes robustness externally from the ambiguity set, whereas ($\mathcal{BDO}$) does so internally.

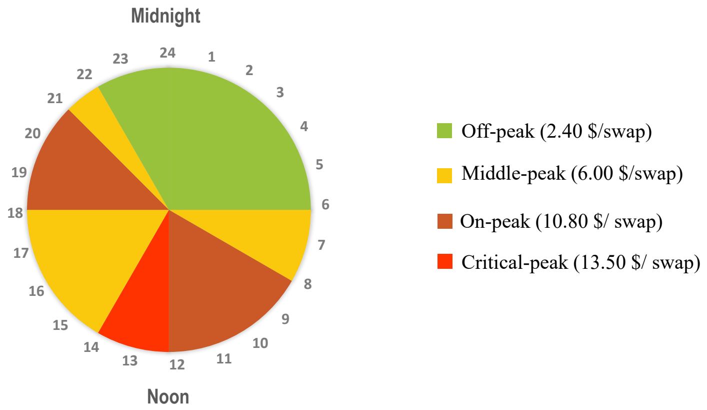

Figure S.1: Time-of-use electricity prices of Shanghai

C Supplementary Materials in Numerical Studies

C.1 Input Data and Parameters

The locations of swapping and charging stations are sampled from NIO's 45 swapping stations in Shanghai using Amap 2024 (Amap, 2024). We calculate the travel times between any two stations by dividing the distances by a 30 km/hour speed. The times are rounded to the next full period. We use Google Optimization Tools (Perron and Furnon, 2024) to solve the TSP solution for each instance and label the swapping station 1 through N by following the sequence of the TSP tour. Using the resulting label, we adopt the chaining structure such that two neighboring swapping stations form a pooling group.

Regarding daily battery swapping demand, a key factor is whether the day is a workday. We assume there are two event realizations $S = 2$; $\tilde{s} = 1$ indicates weekends and $\tilde{s} = 2$ represents workdays. The probabilities are $p_1 = 0.3$ and $p_2 = 0.7$ respectively. Intuitively, the demands of swapping stations within the same pooling group are likely correlated. Hence, we assume that the demands of a pooling group are jointly normally distributed. Specifically for a chaining structure, we consider a truncated bivariate normal distribution $(\tilde{d}_i^t, \tilde{d}_{i+1}^t) \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The marginal probability distribution function of d_i^t is a truncated normal distribution $N(\mu_i^{ts}, \sigma_i^{ts})$ with support $\{d_i^t : \mu_i^{ts} - 3\sigma_i^{ts} \leq d_i^t \leq \mu_i^{ts} + 3\sigma_i^{ts}\}$. The demands among pooling groups are i.i.d. According to empirical evidence, there is a higher traffic volume on weekdays than on weekends. We hence set $\mu_i^{ts} = \xi_t(5 + 5s)$ for all $i \in [N]$ and $t \in [T]$, where ξ_t is a scaling factor considering the distribution of hourly demand for batteries. As in Bertsimas et al. (2019), we let $\sigma_i^t = \mu_i^{ts} \cdot \epsilon_s$ in which $\epsilon_1 = \epsilon_2 = 0.1$. The covariance matrix $\boldsymbol{\Sigma}$ is given by

$$[\boldsymbol{\Sigma}]_{ij} = \begin{cases} \alpha\sigma_i\sigma_j & \text{if } i, j \in \mathcal{N}_h, \\ \sigma_j^2 & \text{otherwise,} \end{cases} \tag{S.15}$$

where $\alpha \in [0, 1]$ is the correlation coefficient. We approximate the non-positive definite matrix using the nearest positive definite matrix.

Based on the above distribution, we randomly sample 100,000 demand sample paths. The first 80% of the data is used as the training set to obtain the optimal solution, and the last 20% of the data is used as the test set for out-of-sample testing. Then the parameters for the ambiguity set are obtained from the observed samples. The empirical distribution $\hat{\mathbb{P}}$ is derived from samples, where $\hat{p}_s$ is the empirical frequency for event s in the observations. When the event information is not considered, the parameters of the ambiguity set are obtained by pooling the sample data.

The deterministic parameters are obtained from multiple resources, which we briefly show in the Table S.1. From NIO's report, the battery capacity at each swapping station is 20, thus we set the inventory window as $[5, 20]$. We consider the Standard Range Battery of 75 kWh capacity, widely equipped in models NIO eT5, eT7, and eS8. The standard fast charger will take $T^C = 3$ hours to charge a battery from SOC 20% to 90%. Each battery pack costs about USD$10,000 in 2024 (NIO, 2024). The battery swapping cost c_t is charged based on time-of-use electricity prices. Figure S.1 demonstrates the time-of-use prices of Shanghai (Shanghai Municipal Development & Reform Commission, 2022). We set the unit transshipment cost to be USD$1.1 per battery per km as in Qi et al. (2023). The penalty cost is hard to quantify since it depends on the specific contract between the operator and customer Schneider et al. (2018), we set it as USD$10000 and perform an extensive sensitivity analysis later.

Table S.1: Input parameters

Parameter	Symbol	Value
Number of peiods	T	12
Inventory window	$[\underline{\tau}_i^t, \overline{\tau}_i^t]$	$[5, 20]$
Charging time	T^C	3
Unit battery cost	K	10000
Charging cost	c^t	see Figure S.1
Transshipment cost	p	1.1
Penalty cost	b_t	10000

C.2 Details for Out-of-Sample Tests

To further examine the performance of four models in perturbed distributions to validate our model's robustness. In particular, for event $\tilde{s} = 1$, we fix the out-of-sample mean $\boldsymbol{\mu}_{out}^1 = \boldsymbol{\mu}^1$; for event $\tilde{s} = 2$, we vary the out-of-sample mean by $\boldsymbol{\mu}_{out}^2 = (1 + \Delta)\boldsymbol{\mu}^2$ where Δ take values in $\{-0.1, -0.05, 0, 05, 0.1\}$. In addition to the normal distribution, we consider generating demand samples independently of two other distributions: uniform and Poisson. The range of the uniform distribution is set as $[0.5\mu_i^{ts}, 1.5\mu_i^{ts}]$. The average arrival rate of the Poisson distribution is set as μ_i^{ts}.

C.3 Additional Out-of-Sample Test Results

Following the experiment setting in Section 6.1, we report the battery numbers of each station. Figure S.2 presents the number of initial batteries at the central charging station 0 and the total batteries at all stations when $\tilde{s} = 2$. We can see in Figure S.2(a) that RS-S, DRO-S, and DRO-NS stock almost the same number of batteries at the charging station 0, while RS-NS holds the least and RS-PE holds the most batteries at station 0 (68 and 432 respectively). It turns out that RS-NS holds the lowest battery stock in total among the five models but incurs more inventory violations than the others, as shown in Figure S.2(b). By contrast, RS-PE stocks the highest number of batteries, all of which are kept at the central charging station. This is mainly because the RS-PE overfits the in-sample data, and the result ELDR fails to adapt well to the out-of-sample demands. The DMs turn out to be overly risk-averse; they prioritize centralized inventory over distributed inventory.

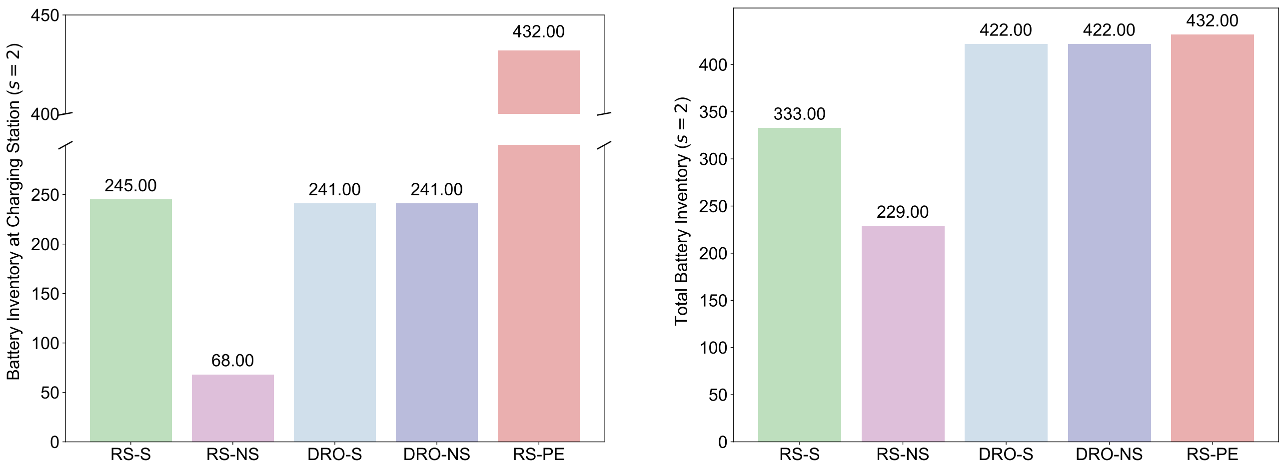

(a): Initial battery inventory at charging station 0 (b): Total battery inventory at all stations

Figure S.2: Comparison of initial battery inventory of five models ($\tilde{s} = 2$)

Figure S.3 demonstrates the battery dynamics at each swapping station when $\tilde{s} = 1$, which confirms the conclusion in Section 6.1. Our proposed RS-S model demonstrates slightly more stable inventory dynamics than the other four models.

Due to the limited space in the main body, we present the detailed out-of-sample results in the following tables. Table S.2 documents the detailed average performance of the four models under six distributions as in Figure 9. Table S.3 show the inventory violation statistics for RS-S under three flexibility structures. The observation in Section 6.2 is stable, the chaining structure is effective in lateral transshipment. Table S.4 displays the impacts of the Budget on the out-of-sample inventory violation of each model. Similar to the results in Section 6.4, our RS-S model dominates RS-NS in all metrics.

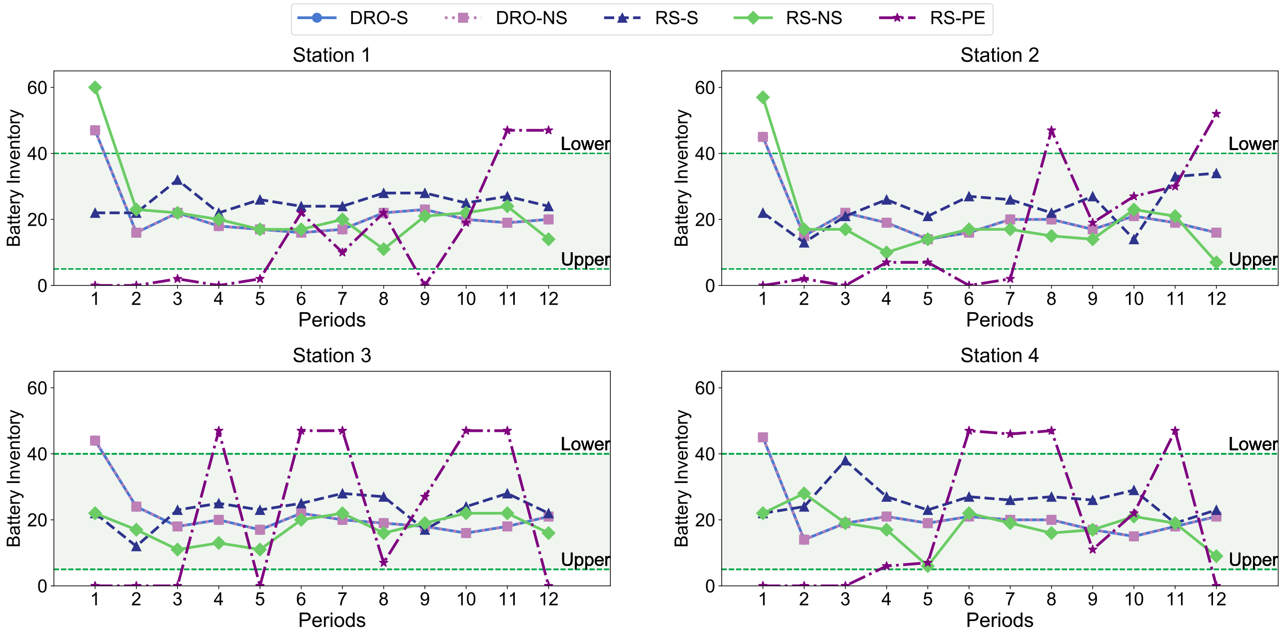

Figure S.3: Mean values of battery inventory dynamics at swapping stations for five models ($\tilde{s} = 2$)

Table S.2: Average Performance of RS-S and Benchmark Models Under Different Distributions

Model	Prob	Mean	Std.	VaR95%	CVaR95%
RS-S	1.00	1.00	1.00	1.00	1.00
DRO-S	2.03	2.20	1.50	3.46	1.87
DRO-NS	2.03	2.20	1.50	3.46	1.87
RS-NS	2.34	3.64	2.81	4.52	3.81

Table S.3: Comparison of Inventory Violation under RS-S for Three flexibility structures

$\mathcal{G}$	Prob %	Mean	Std.	VaR95%	CVaR95%
Dedicated	10.55	0.58	2.02	7.48	8.21
2-chain	8.05	0.38	1.80	4.37	5.70
Full chain	11.20	0.62	2.40	11.72	6.56

Table S.4: Comparison of Inventory Violation under RS-S and RS-NS Models for Different Budgets

Model	δ	Prob%	Mean	Std.	VaR%95	CVaR%95
RS-S	0	3.34	0.25	2.49	0.00	0.25
	1	0.00	0.00	0.00	0.00	0.00
	2	0.00	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00	0.00
	4	0.00	0.00	0.00	0.00	0.00
	5	0.00	0.00	0.00	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00
	7	0.00	0.00	0.00	0.00	0.00
	8	0.00	0.00	0.00	0.00	0.00
	9	0.00	0.00	0.00	0.00	0.00
RS-NS	0	13.83	1.14	4.01	9.13	17.76
	1	4.32	0.11	0.38	0.88	1.55
	2	5.22	0.13	0.40	1.03	1.63
	3	6.04	0.16	0.42	1.13	1.67
	4	5.25	0.14	0.39	1.03	1.55
	5	4.53	0.13	0.35	0.94	1.42
	6	4.09	0.12	0.34	0.90	1.36
	7	3.23	0.11	0.30	0.79	1.21
	8	2.67	0.10	0.28	0.72	1.13
	9	2.11	0.08	0.25	0.63	1.02

References

Amap. A Mapping and Navigation Service, 2024. URL https://lbs.amap.com/. (Accessed: 2024-08-11).

Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

Dimitris Bertsimas, Melvyn Sim, and Meilin Zhang. Adaptive distributionally robust optimization. *Management Science*, 65(2):604–618, 2019.

Lucy Gongtao Chen, Daniel Zhuoyu Long, and Melvyn Sim. On dynamic decision making to meet consumption targets. *Operations Research*, 63(5):1117–1130, 2015.

Vincent Tsz Fai Chow, Zheng Cui, and Daniel Zhuoyu Long. Target-oriented distributionally robust optimization and its applications to surgery allocation. *INFORMS Journal on Computing*, 34(4):2058–2072, 2022.

Zheng Cui, Daniel Zhuoyu Long, Jin Qi, and Lianmin Zhang. The inventory routing problem under uncertainty. *Operations Research*, 71(1):378–395, 2023.

Hans Föllmer and Alexander Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6:429–447, 2002.

Nicholas G Hall, Daniel Zhuoyu Long, Jin Qi, and Melvyn Sim. Managing underperformance risk in project portfolio selection. *Operations Research*, 63(3):660–675, 2015.

Jie Hu, Zhi Chen, and Shuming Wang. Budget-driven multiperiod hub location: A robust time-series approach. *Operations Research*, 73(2):613–631, 2024.

Hongming Li, Erick Delage, Ning Zhu, Michael Pinedo, and Shoufeng Ma. Distributional robustness and inequity mitigation in disaster preparedness of humanitarian operations. *Manufacturing & Service Operations Management*, 26(1):197–214, 2024.

Daniel Zhuoyu Long, Melvyn Sim, and Minglong Zhou. Robust satisficing. *Operations Research*, 71(1):61–82, 2023.

NIO. NIO releases new baas battery rental service policy, 2024. URL https://www.nio.cn/news/20240314002. (Accessed: 2024-09-16).

Laurent Perron and Vincent Furnon. OR-Tools: Google's Open Source Optimization Tools, 2024. URL https://developers.google.com/optimization/. (Version v9.10, Accessed: 2024-05-07).

Jin Qi. Mitigating delays and unfairness in appointment systems. *Management Science*, 63(2):566–583, 2017.

Wei Qi, Yuli Zhang, and Ningwei Zhang. Scaling up electric-vehicle battery swapping services in cities: A joint location and repairable-inventory model. *Management Science*, 69(11):6855–6875, 2023.

Frank Schneider, Ulrich W Thonemann, and Diego Klabjan. Optimization of battery charging and purchasing at electric vehicle battery swap stations. *Transportation Science*, 52(5):1211–1234, 2018.

Shanghai Municipal Development & Reform Commission. Notice on further improving the time-of-use electricity price mechanism in shanghai city, 2022. URL https://fgw.sh.gov.cn/fgw_jggl/20221216/e2652e3ab7ee49438d6e82af8880b160.html. (Accessed: 2024-09-13).

Wolfram Wiesemann, Daniel Kuhn, and Melvyn Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.