When Peak Usage Becomes Capacity: Optimal Charging of Battery Swapping Stations Under Demand Charges

Zimeng Li
Department of Service Science and Operations Management, School of Management, Zhejiang University, Hangzhou, China
zimeng@zju.edu.cn
Department of Decision Analytics and Operations, College of Business, City University of Hong Kong, Kowloon, Hong Kong
zimengli7-c@my.cityu.edu.hk

Huiyin Ouyang
Faculty of Business and Economics, The University of Hong Kong, Pok Fu Lam Road, Hong Kong
oyhy@hku.hk

Zhankun Sun
Department of Decision Analytics and Operations, College of Business, City University of Hong Kong, Kowloon, Hong Kong
zhankun.sun@cityu.edu.hk

Quan Yuan
Department of Service Science and Operations Management, School of Management, Zhejiang University, Hangzhou, China
quanyuan@zju.edu.cn

Battery swapping stations provide a promising alternative for rapid electric vehicle (EV) refueling, but their charging operations must coordinate stochastic demand, renewable energy availability, battery degradation, and demand charge tariffs. We formulate the charging problem of a renewable-integrated station as a finite-horizon dynamic program and uncover a fundamental cost-capacity tradeoff: while peak grid usage level incurs demand charges, it also creates reusable charging capacity that can benefit future operations. We analytically characterize the optimal single-period charging policy and establish a two-threshold structure that separates aggressive, economic, and conservative charging regimes. This structural result provides a mathematical explanation for why optimal charging decisions can be non-monotone in the peak usage level and even exceed a no-demand-charge benchmark. To address the curse of dimensionality in multi-period settings, we leverage these structural insights to develop structure-informed heuristics and a structure-guided reinforcement learning policy. Using a case study of a NIO battery swapping station in the Netherlands, we quantify the financial impact of ignoring demand charges and demonstrate the effectiveness of our proposed policies in improving operational performance.

Keywords: Battery Swapping Stations, Demand Charge, Battery Degradation, Electric Vehicle

1. Introduction

The rapid growth of electric vehicles (EVs) is reshaping the design and operation of charging infrastructure. Global sales of fully electric and plug-in hybrid vehicles exceeded 17 million in 2024, representing a 25% increase over 2023 (IEA 2025). China remains the largest and fastest-growing market, with about 11 million new energy vehicles sold in 2024, and EVs accounted for over 70% of newly registered private vehicles in Hong Kong by late 2024 (China Daily 2024). As EV penetration increases, the key infrastructure challenge is no longer only whether charging facilities are available, but whether refueling services can be provided quickly, reliably, and economically without creating excessive stress on the power grid.

Most existing EV refueling infrastructure relies on plug-in charging. Although fast charging technologies continue to improve, they do not fully eliminate operational frictions. For example, Tesla's V4 Superchargers can deliver up to 500 kW of power and charge a vehicle from 10% to 80% in about 30 to 35 minutes (Tesla 2026). This is still much slower than refueling a gasoline vehicle, and the high instantaneous power requirement of fast chargers can create substantial grid pressure. Range anxiety—the concern that a vehicle may run out of battery before convenient refueling is available—therefore remains an important barrier to wider EV adoption.

Battery swapping stations (BSSs) resolve this friction by exchanging a depleted battery for a fully charged one in under three minutes. This rapid turnaround is especially valuable for high-utilization commercial fleets (taxis, ride-sharing vehicles) where charging downtime is costly. Because of the quick turnaround, BSSs require much less space compared to charging stations, which often require parking areas to accommodate vehicles during long charging sessions. This space efficiency makes BSSs especially advantageous in compact, populated urban cities like Hong Kong and New York. Moreover, BSSs can store excess renewable energy in batteries, aligning EV charging with sustainable energy goals. Industry adoption reflects these advantages. As of early 2026, NIO operates over 3,750 swap stations across China and 61 stations in Europe, handling approximately 100,000 swaps per day (NIO 2026). CATL, the world's leading battery maker, has built more than 1,470 stations, targeting 3,000 stations by the end of 2026 (CATL 2026). This aggressive scaling underscores CATL CEO Robin Zeng's prediction that *by 2030, battery swapping, home charging piles, and public charging piles will each account for a third of the EV charging market*, establishing BSS as a pillar of future EV infrastructure.

The scaling of BSSs, however, creates a distinctive operational problem. A station must maintain enough fully charged batteries to serve stochastic swapping demand, while deciding when to recharge depleted batteries using intermittent renewable generation and grid electricity. These decisions are affected by time-varying energy prices, battery degradation from repeated charging, waiting costs when customers cannot be served immediately, and demand charges. Under many electricity tariffs, commercial users pay not only for total energy consumption but also for the maximum grid power draw, referred to as *peak usage level* thereafter, recorded during a billing cycle (Jin and Xu 2017, Chen et al. 2024). Because a BSS must recharge multiple high-capacity battery packs to maintain service readiness, uncoordinated charging can easily trigger large grid power spikes. These demand charges are far from trivial; empirical evidence shows that demand charges routinely account for 31% to 49% of a commercial station's total utility bill, and uncontrolled charging can inflate these peak fees by up to 335% compared to optimized scheduling benchmarks (Lee et al. 2020). Thus, managing peak usage level while preserving service quality is the central economic hurdle for modern BSS operations.

The existing operations management literature has extensively studied EV charging schedules, focusing primarily on minimizing time-of-use volumetric electricity costs (e.g., Wu et al. 2017, Jin et al. 2023) or optimizing battery availability (e.g., Mak et al. 2013, Schneider et al. 2018, Widrick et al. 2018). However, prior research rarely integrates the joint complexities of path-dependent demand charges, intermittent solar generation, and stochastic customer swapping demand. Incorporating demand charges fundamentally alters the mathematical structure of the problem because the peak usage level updates via a non-linear maximum operator. This introduces a highly non-linear, path-dependent state variable that prevents the direct application of standard arguments commonly used in operations management for proving structural properties of optimal value functions, leaving a critical gap in our theoretical understanding of how a BSS should coordinate battery inventory and peak grid usage over time.

In this paper, we address this gap by asking a fundamental question: *What is the structure of the optimal charging policy for a renewable-integrated BSS operating under demand charges, and how should operators leverage the resulting insights to develop implementable heuristic policies?* To answer this, we formulate this problem as a finite-horizon dynamic program over a billing cycle. In each decision period, the station observes its fully charged battery inventory, realized solar generation, grid energy price, and current peak usage level, and then decides how many depleted batteries to charge. The station earns swap revenue from served demand, incurs waiting costs for unserved customers, pays energy and battery degradation costs for repeated charging, and pays a demand charge based on the maximum grid usage level reached during the billing cycle. The objective is to maximize total expected profit through forward-looking charging decisions.

Our framework explicitly captures a compelling "cost or capacity" tension. The peak usage level is inherently a cost state because a higher peak increases the terminal demand charge. At the same time, once a peak has already been established, it acts as an implicit capacity endowment—the BSS can freely pull grid energy up to that peak usage level in later periods without incurring any additional demand charge. To this end, our main contributions are summarized as follows:

First, we provide the first exact analytical characterization of optimal charging decisions for a renewable-integrated BSS under demand charges. For the single-period charging problem, we establish structural properties of the value function and prove a renewable-first property. When solar generation is insufficient, we prove that the optimal policy exhibits a non-monotonic, three-regime threshold structure: an aggressive regime (deliberately driving up the peak usage level to "purchase" future capacity), an economic regime (maximizing inventory within the existing peak usage level), and a conservative regime (curtailing charging in very high peak usage levels to save empty batteries for future solar or off-peak hours). We also demonstrate an inventory-capacity substitutability property, showing that the optimal charging quantity decreases with the inventory of fully charged batteries, bounded by a marginal effect of one.

Second, by comparing our model to a no-demand-charge benchmark, we show that a natural intuition that adding a demand charge should uniformly suppress grid usage can fail. When the peak usage level is sufficiently high, a demand-charge-aware policy actually charges more aggressively than the benchmark because its marginal capacity cost has effectively dropped to zero. This demonstrates that peak usage level is not merely a terminal billing outcome, but an active operational state variable that alters the marginal value of inventory throughout the billing cycle.

Third, moving from theory to execution, we translate these structural principles into scalable decision rules for realistic multi-period charging settings where a battery requires multiple decision periods to become fully charged. Because the battery-pipeline state and peak usage level make the dynamic program high dimensional, solving the dynamic program becomes computationally impractical. Hence, we develop decision rules leveraging our theoretical insights. We develop structure-informed heuristics (Policy ST and Policy CB) alongside a structure-guided reinforcement learning algorithm (Policy ST-PPO).

Lastly, we validate our framework by a case study calibrated to a NIO station in the Netherlands, integrating real Dutch electricity prices and regional solar generation data. Our numerical results demonstrate that ignoring demand charges can lead to substantial profit losses of up to 77.90%. While solar energy sharply lowers peak exposure, station battery capacity exhibits a more nuanced effect: expanding station capacity can initially raise the optimal peak usage level to secure reusable charging capacity, though it eventually drives peak-smoothing benefits. Numerically, our structure-informed Policy ST virtually replicates the exact analytical optimum, yielding an optimality gap of just 0.47% for $L = 2$. We further show that the analytical threshold structure can improve reinforcement learning: ST-PPO, which trains PPO with guidance from Policy ST, consistently outperforms standard PPO in high-dimensional environments.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the base charging model and characterizes the optimal single-period charging policy. Section 4 extends the model to multi-period charging and develops structure-informed policies. Section 5 presents the case study and numerical experiments. Section 6 concludes. Proofs and algorithmic details are provided in the appendices.

2. Literature Review

This paper is closely related to three streams of literature: (i) infrastructure planning and operations for EV battery swapping stations, (ii) operations management under demand charges, and (iii) management of renewable energy and storage systems. Below, we review each area and highlight our analytical contributions.

2.1. Infrastructure Planning and Operations for Battery Swap Stations

A growing body of literature investigates the strategic and operational decisions of EV battery swap stations. The first sub-stream focuses on strategic infrastructure planning, encompassing location choices and capacity investments. An early work by Mak et al. (2013) develops optimization frameworks to determine optimal station locations and battery inventory under uncertain demand. Sun et al. (2019) utilize a fluid model to optimize the capacity of charging bays, the number of batteries, and the recharging decisions at a BSS. Furthermore, alongside infrastructure planning, Shi and Hu (2024) study the profitability of flexible battery leasing business models.

The second sub-stream emphasizes the operational decisions of BSSs, specifically service pricing and charging control. Avci et al. (2015) demonstrate that battery swapping can enhance EV adoption and optimize per-mile pricing, while Hu et al. (2023) compare pay-per-swap and subscription pricing schemes in a game-theoretic model. In a related EV charging-service setting, Lee et al. (2020) show that demand charges can account for a substantial share of operating costs even under optimized charging and pricing decisions.

Two studies are particularly close to our operational modeling approach. Schneider et al. (2018) use approximate dynamic programming to manage battery inventory and charging decisions across a network of battery swap stations, quantifying the value of lateral transshipments. Widrick et al. (2018) focus on single-station dynamics and analytically characterize monotone structural properties of optimal charging and discharging policies under nonstationary conditions. Our work shares their dynamic operational perspective, but differs by incorporating renewable generation and demand charges, which introduce the peak usage level as an additional state variable and lead to a non-monotone, two-threshold charging structure. Sun et al. (2017) frame the operation as a constrained Markov decision process to minimize costs while strictly guaranteeing quality-of-service (QoS). Extending this operational focus, Yu et al. (2024) employ a fluid control approach to optimize charging decisions by explicitly considering the battery state of charge (SoC). Furthermore, Wu et al. (2017) apply an evolutionary algorithm to minimize energy and battery degradation costs when scheduling charging under time-varying prices. Recently, Qi et al. (2023) propose a hybrid “swap-locally and charge-centrally” infrastructure, leveraging submodularity to exactly solve a joint location and repairable-inventory model.

Lastly, researchers have explored the integration of renewable energy and vehicle-to-grid (V2G) capabilities into BSS operations. Due to the high-dimensional complexity of stochastic demand, intermittent solar/wind supply, and volatile electricity prices, literature in this area leans heavily on predictive and algorithmic approaches. For instance, Chawrasia et al. (2024) design a solar-integrated BSS utilizing LSTM-assisted forecasting. Jin et al. (2023) employ deep reinforcement learning (DRL) specifically to manage real-time charging scheduling. Mao et al. (2024) utilize a hypergraph-based multi-agent reinforcement learning framework to tackle the joint EV routing and battery charging problem. Additionally, broader grid-integration studies propose bi-level optimization frameworks for coordinating microgrids and BSSs (Esmaeili et al. 2019) and contractual mechanisms for coordinating EV charging under uncertainty (Li et al. 2024).

While existing literature offers sophisticated capacity planning models and algorithmic solutions for battery recharging, our work stands apart by *analytically characterizing the structure of the optimal charging policy in a renewable-integrated BSS under demand charges.* Unlike purely learning-based algorithms, our exact structural analysis yields explicit, interpretable insights into how demand charges fundamentally alter the optimal charging decisions of BSSs.

2.2. Demand Charges in Operations Management

Demand charges have received increasing attention in electricity tariff design and energy-operations research because they link a customer's bill to peak power usage rather than only to total energy consumption. The tariff-design literature emphasizes that non-coincident demand charges can create incentives that differ from system-level grid costs (Borenstein 2016, Passey et al. 2017). In operations models, Jin and Xu (2017) characterize an optimal threshold policy for electric storage under time-varying energy prices and demand charges, and Lee et al. (2020) study pricing and charging decisions for EV charging services when demand charges are a major operating cost.

The papers most closely related to our setting study EV charging operations under demand charges. Yang et al. (2023) develop a block model predictive control approach for scheduling EV charging with demand charges. The closest paper to ours is Chen et al. (2024), who optimize EV charging management at a charging station with stochastic arrivals and demand charges using an exponential cone programming approach. Their model captures the operational importance of demand-charge-aware scheduling in plug-in charging stations. Our paper differs in both system structure and analytical focus. We study a battery swapping station, where charging decisions affect not only current grid usage but also future charged-battery inventory, renewable-energy utilization, and the path-dependent peak usage level. This inventory-coupled structure leads to a non-monotone, two-threshold optimal policy and reveals the dual role of the peak usage level as both demand-charge liability and implicit charging capacity. Thus, relative to the EV charging literature, our contribution is to characterize how demand charges reshape the optimal charging structure in a renewable-integrated BSS and to translate this structure into implementable policies.

2.3. Management of Renewable Energy and Storage Systems

Finally, our work relates to the literature on renewable energy and storage systems. In our setting, depleted batteries at a BSS serve as storage units that can be recharged using solar or grid power, and we study the resulting charging decisions under renewable generation and demand charges. One stream of this literature examines operational decisions in renewable energy and storage systems, such as generation scheduling and charging control. Early work by Wu and Kapuscinski (2013) analyzes generation operations with intermittent renewables and show that curtailment, together with storage, can improve system efficiency. Zhou et al. (2019) characterize optimal generation and storage policies for wind-based electricity generation with storage and transmission capacity constraints. Another stream focuses on capacity investment in renewable generation and storage systems. Kök et al. (2020) evaluate the capacity investment trade-offs between conventional and renewable energy and show that conventional flexibility complements renewable variability. Wu et al. (2023) examine storage location and capacity investment in a power distribution system, comparing centralized and localized storage. Furthermore, Kaps et al. (2023) study off-grid solar and storage investment decisions, revealing that generation and storage capacities are strategic complements at low investment levels but substitutes at higher levels.

While the above studies provide fundamental insights into the operation and capacity investment of renewable energy and storage systems, we focus on the short-term charging decisions of a renewable-integrated BSS under demand charges. We examine the operational frictions created by demand charges and how solar and battery capacities shape the peak usage level and system performance. Our paper thus extends the renewable energy and storage literature to the context of BSS operations.

3. Model of a Battery Swapping Station

We begin by modeling the operations of a battery swapping station using a dynamic programming formulation, assuming charging a depleted battery takes a single period. We characterize the structure of the optimal charging policy under demand charges, and provide insights. The results help devise heuristic policies that can apply to more realistic situations, including models in which charging a battery takes multiple periods.

3.1. Model Description

We study a battery swapping station that replaces depleted batteries with fully charged ones for EV owners. There are B batteries in the station. We assume that there is a charging bay for each battery because the cost of installing a charging bay is much lower than that of a battery. Moreover, due to space limitation, there is no incentive for BSS to install more charging bays than their batteries. The station operates over a fixed planning horizon (typically one month), which is divided into T

decision periods. For instance, when each period represents 1 hour, we have $T = 30 \times 24 = 720$ for a 30-day month. Decisions are made at the beginning of each period. Let $\mathcal{T}$ denote the set of decision epochs, i.e., $\mathcal{T} \equiv \{0, 1, 2, \cdots, T-1\}$. We assume that the charging times of any depleted battery are identical to simplify our model. This is a mild assumption because the state of charge (SoC) profiles of lithium-ion batteries (commonly used in EVs) exhibit a rapid increase during the initial charging phase, followed by a slower saturation phase. As a result, the total charging times for swapped-off batteries with varying SoC levels are generally similar. This assumption is also consistent with prior studies (Mak et al. 2013, Schneider et al. 2018). We further assume that any depleted battery can be fully charged within a single period (1 hour) for the sake of tractability. Our goal is to analyze the structure of the optimal policy and derive interesting insights from this simple model. We study a more realistic setting later by employing a finer time granularity (e.g., 15 minutes per period), in which case charging a battery takes multiple periods.

In each period $t \in \mathcal{T}$, an exogenous demand for battery swapping, denoted by $d_t \geq 0$, is generated. This demand is a non-negative random variable that follows a general distribution, which is allowed to be nonstationary over time as the swap demand for batteries depends on the time of day and even day of week. The BSS faces stochastic swapping demand and utilizes the fully charged batteries, whose quantity is denoted by x_t, to fulfill this demand. Each successful swap generates a revenue of $p > 0$. Any unmet demand is deferred to the next period, incurring a waiting cost of $s > 0$ per unit per period. Note that x_t can take negative values, in which case $-x_t$ represents the number of EVs waiting in line for swapping.

The station determines the number of batteries to charge in period t, denoted by q_t. Let $\bar{q}(x_t) \equiv B - x_t^+$ denote the maximum feasible charging quantity at state x_t, and let $\mathcal{Q}(x_t) \equiv [0, \bar{q}(x_t)]$ denote the feasible action set. Thus, $q_t \in \mathcal{Q}(x_t)$, where x_t^+ equals x_t if $x_t > 0$ and 0 otherwise. Because depleted batteries can be fully charged within a single period, the q_t batteries that start charging in the period t become available for swapping at the beginning of period $t+1$. Hence, we have

$$x_{t+1} = x_t - d_t + q_t, \ t \in \mathcal{T}. \tag{1}$$

The depleted batteries can be recharged using electricity from the power grid and solar energy generated by the solar panels installed at the BSS. The solar energy supply, however, is intermittent and uncertain. We assume that the solar energy generated during period t is $\Delta_t K$. Here, K denotes the installed capacity of the solar panels, measured by the maximum number of depleted batteries that can be fully charged by solar energy in one period. The random variable $\Delta_t \in [0, 1]$ captures the fraction of this capacity realized in period t. Notably, the specific distribution of Δ_t does not affect our analytical results. Moreover, Δ_t could also depend on the value in the previous period. In our numerical study, following Wu and Kapuscinski (2013) and Kouvelis et al. (2023), we model Δ_t as a Markovian process to better capture the fluctuations of solar generation in practice (see Section 5 for more details). Consistent with prior studies, we assume that the realized Δ_t, denoted by δ_t, can be observed at the beginning of period t (Wu and Kapuscinski 2013, Zhou et al. 2019, Kouvelis et al. 2023). This assumption is reasonable as the duration of each decision period is typically short (no more than 1 hour), allowing for accurate forecasting of the solar output in this period.

We assume that, once installed, solar panels generate power at zero marginal cost (Kaps et al. 2023, Wu et al. 2023). Hence, BSS prioritizes solar energy to fulfill the charging demand. The remaining energy requirement, $(q_t - \delta_t K)^+$, is met by electricity from the power grid. The tariffs for charging batteries using grid electricity have two components: the energy consumption cost which is based on the total amount of electricity consumed (measured in kilowatt-hours, kWh), and the demand charge which is based on the peak usage level, i.e., the highest rate of electricity usage that occurs at any point during the billing period (measured in kilowatts, kW). Let c_t denote the energy price for fully charging one depleted battery, which is exogenous and predetermined. We assume that the energy price can be time-dependent. Then, the energy consumption cost in period t is $c_t(q_t - \delta_t K)^+$. Let ϕ_t denote the peak usage level observed from 0 to period t. Then, we have

$$\phi_{t+1} = \max\left\{\phi_t, (q_t - \delta_t K)^+\right\}, t \in \mathcal{T}. \tag{2}$$

At the end of the billing period, i.e., $t = T$, the BSS pays a demand charge of $D(\phi_T)$, where $D(\cdot)$ is a convex and increasing function.

In addition to the electricity tariffs, repeated charging can accelerate battery wear, resulting in reduced battery capacity and lifespan. Battery degradation results from complex electrochemical and mechanical processes inside the battery. For tractability, we assume that recharging one depleted battery incurs a constant degradation cost, denoted by θ (Šepetanc and Pandžić 2020, Jin et al. 2023). For instance, the estimated degradation cost is \$1.6 for lithium-ion batteries with a nominal capacity of 40 kWh installed on Nissan Leaf EVs (Gowda et al. 2021). To avoid triviality and for the battery swapping business to be economically viable, we further assume $0 < \theta < p$.

Let $S_t = (x_t, \phi_t)$ denote the state of the BSS at the beginning of period t. The state space can be written as $\mathcal{S} = \{(x, \phi) : x \le B, 0 \le \phi \le B\}$. Let $V_t(S_t)$ denote the optimal value function of the system in state S_t, which captures the total expected profit from period t through the end of the billing cycle, assuming optimal charging actions are made in all subsequent periods. Without loss of generality, we assume that when $t = 0$, all batteries are fully charged, i.e., $x_0 = B$ and $\phi_0 = 0$.

The objective of the station is to maximize the total expected profit $V_0(B,0)$ by determining the charging decisions $q_t,\ t\in\mathcal{T}$. Then, for any $t\in\mathcal{T}$, the optimality equations can be written as follows[1]:

$$V_t(x_t,\phi_t)=\mathbb{E}_{\Delta_t}\left[\max_{q_t\in\mathcal{Q}(x_t)}\left\{\mathbb{E}_{d_t}\left[pd_t-c_t(q_t-\delta_tK)^+-\theta q_t-s(d_t-x_t)^++V_{t+1}(x_{t+1},\phi_{t+1})\right]\right\}\middle|\Delta_t=\delta_t\right], \tag{3}$$

$$V_T(x_T,\phi_T)=px_T-D(\phi_T). \tag{4}$$

The state transitions for x_{t+1} and ϕ_{t+1} are defined in (1) and (2), respectively. The first term in Equation (4) implies that all fully charged batteries in the terminal period are swapped out, and the second term is the demand charge paid at the end of the billing cycle.

Note that BSSs are non-terminating systems that provide swapping services continuously. However, to keep track of the demand charge, we model it as a finite-period problem for simplicity. As a result, we need to specify the terminal conditions in our formulation. We assume all fully charged batteries are swapped out and each swap generates a revenue of p. This assumption essentially implies infinite swapping demand in the terminal period, which is restrictive as it deviates from reality. However, as it becomes obvious from the proof, the structural results do not depend on the specific linear terminal value in Equation (4).

3.2. Characterization of Optimal Policy

In this section, we characterize the structure of the optimal charging policy by analyzing the dynamic programming formulation in Equations (3) and (4). We start by characterizing the structural properties of the optimal value functions.

PROPOSITION 1 **(Structural Properties of Optimal Value Functions)**. *For any $t\in\mathcal{T}$, the optimal value function $V_t(x_t,\phi_t)$ is (i) increasing in the number of fully charged batteries x_t and decreasing in the peak usage level ϕ_t; and (ii) jointly concave and submodular in (x_t,ϕ_t).*[2]

Proposition 1 first establishes the monotonicity properties of the optimal value function in x_t and ϕ_t, respectively. The joint concavity of the value function implies diminishing marginal returns in the expected total profit when the number of fully charged batteries x_t increases or when the peak usage level ϕ_t decreases. The submodularity of the value function indicates that the marginal benefit of increasing the number of fully charged batteries x_t diminishes when the peak usage level ϕ_t increases. The structural properties of the optimal value functions in Proposition 1 are critical for the proof of the optimal policy presented in the following results.

[1] If the optimizer is not unique, we select the smallest optimal solution. Throughout the paper, the optimal charging decision refers to this minimal optimizer.

[2] In this paper, we use *increasing* and *decreasing* in their weaker sense unless otherwise specified.

Proposition 2 **(Renewable-First Policy)**. *It is optimal to use all available solar energy generated in period t to charge depleted batteries whenever feasible; i.e., $q_t^* \geq \min\{\delta_t K, B - x_t^+\}$, $\forall t \in \mathcal{T}$.*

Because of $q_t^* \leq B - x_t^+$, Proposition 2 implies that the optimal charging quantity $q_t^* = B - x_t^+$ if $\delta_t K \geq B - x_t^+$. In other words, when the solar energy generated in period t is sufficient to fully charge all depleted batteries, it is optimal to charge all the depleted batteries with solar energy. This is intuitive because charging a battery using solar energy is free, except for battery degradation costs, which are assumed to be less than the revenue generated per battery swap. We refer to this as the Renewable-First Policy, under which BSS should prioritize renewable energy to charge the depleted batteries as much as possible. The red rhombuses in Figure 1 show the optimal charging quantities when the solar energy is enough to charge all the depleted batteries, i.e., $\delta_t K \geq B - x_t^+$.

Next, we examine the optimal charging policy when the solar generation is insufficient to fully replenish all depleted batteries, i.e., when $\delta_t K < B - x_t^+$.

Theorem 1 **(Optimal Policy)**. *Assume $\delta_t K < B - x_t^+$. The optimal policy is of threshold type. Specifically, for any given x_t, there exist two thresholds $\underline{\phi}_t(x_t)$ and $\overline{\phi}_t(x_t)$ such that $0 \leq \underline{\phi}_t(x_t) \leq \overline{\phi}_t(x_t) \leq B - x_t^+ - \delta_t K$. Then, the optimal charging quantity q_t^* is as follows:*

*(i) (**Aggressive Charging**) If $\phi_t \leq \underline{\phi}_t(x_t)$, then $q_t^* = \delta_t K + \underline{\phi}_t(x_t)$;*

*(ii) (**Economic Charging**) If $\underline{\phi}_t(x_t) < \phi_t < \overline{\phi}_t(x_t)$, then $q_t^* = \delta_t K + \phi_t$;*

*(iii) (**Conservative Charging**) If $\phi_t \geq \overline{\phi}_t(x_t)$, there exists $0 \leq q_t^c(\phi_t) \leq \overline{\phi}_t(x_t)$ such that $q_t^* = \delta_t K + \overline{\phi}_t(x_t) - q_t^c(\phi_t)$. Moreover, $q_t^c(\phi_t)$ is increasing in ϕ_t.*

Furthermore, both $\underline{\phi}_t(x_t)$ and $\overline{\phi}_t(x_t)$ are decreasing in x_t.

When solar energy generated in period t is not sufficient to fully charge all depleted batteries, the BSS must decide how much additional charging, if any, should be supplied by grid power. Theorem 1 establishes that the optimal charging policy in this scenario is characterized by two critical thresholds, $\underline{\phi}_t(x_t)$ and $\overline{\phi}_t(x_t)$. For any given inventory of fully charged batteries x_t, these thresholds partition the state space into three distinct charging regimes, detailed as follows.

First, when the peak usage level is low, i.e., $\phi_t \leq \underline{\phi}_t(x_t)$, the optimal charging quantity is constant and independent of the current peak usage level ϕ_t. In this regime, the policy dictates charging a target amount: utilizing $\delta_t K$ from solar energy and drawing $\underline{\phi}_t(x_t)$ from the grid. Despite the resulting increase in demand charges at the end of the billing period, the BSS intentionally increases the peak usage level ϕ_t to $\underline{\phi}_t(x_t)$. Hence, we refer to this as the *aggressive charging regime*. The orange triangles in Figure 1 illustrate the optimal charging quantities in this region.

Second, when the peak usage level is moderate, specifically in the range $\underline{\phi}_t(x_t) < \phi_t < \overline{\phi}_t(x_t)$, the optimal charging quantity q_t^* is characterized by the explicit formula: $q_t^* = \delta_t K + \phi_t$. This formula dictates that the strategy prioritizes utilizing all available solar energy $\delta_t K$. Subsequently, the BSS

Figure 1 **The optimal charging quantity q_t^***

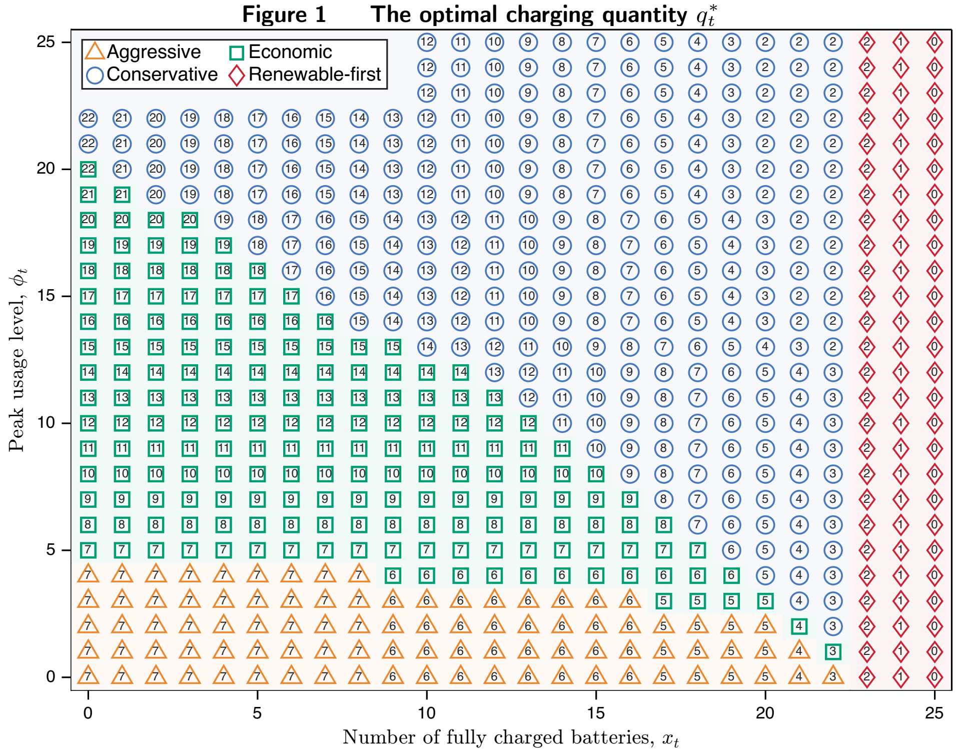

Notes. Parameters: $T=24$, $t=0$, $B=25$, $\delta_t K=2$, $d_t \sim N(10,20)$, $\Delta_t \sim N(0.5,0.7)$, $D(\phi_t)=400\phi_t$, $c_t=12$, $s=10$, $\theta=8$, and $p=35$. The numbers inside the markers show the optimal charging quantity q_t^*.

supplements this with grid power to maximize the inventory of fully charged batteries without exceeding the current peak usage level ϕ_t. This approach minimizes the risk of battery shortage while strictly avoiding additional demand charges. This behavior, which balances service quality and demand charge costs, is referred to as the *economic charging regime*. The green squares in Figure 1 show the optimal charging quantities under this regime.

Finally, when the peak usage level is sufficiently high, i.e., $\phi_t \geq \overline{\phi}_t(x_t)$, the optimal grid charging quantity is $\overline{\phi}_t(x_t) - q_t^c(\phi_t)$, which is less than the current peak usage level, as $\phi_t \geq \overline{\phi}_t(x_t)$ and $q_t^c(\phi_t) \geq 0$. Moreover, as the peak usage level increases, it is optimal to charge fewer depleted batteries. We refer to this strategy as *conservative charging*. The blue circles in Figure 1 depict the optimal charging quantities in this conservative regime.

Theorem 1 also establishes that the two thresholds, $\underline{\phi}_t(x_t)$ and $\overline{\phi}_t(x_t)$, decrease with the number of fully charged batteries x_t, implying a smaller region for the aggressive charging regime and a greater region for the conservative charging regime. Hence, with a higher inventory of fully charged batteries, the BSS transitions into the conservative charging regime at a lower peak usage level because it does not need to take risks with high demand charges to secure inventory.

Corollary 1 (**Sensitivity of q_t^* to ϕ_t**). *Assume $\delta_t K < B - x_t^+$. For all $t \in \mathcal{T}$, the relationship between the optimal charging quantity q_t^* and ϕ_t is non-monotone. At every point where q_t^* is differentiable, its derivative satisfies*

$$\frac{\partial q_t^*}{\partial \phi_t} = \begin{cases} 0, & \text{if } \phi_t \leq \underline{\phi}_t(x_t), \\ 1, & \text{if } \underline{\phi}_t(x_t) < \phi_t < \overline{\phi}_t(x_t), \\ \leq 0, & \text{if } \phi_t \geq \overline{\phi}_t(x_t), \end{cases}$$

where the thresholds $\underline{\phi}_t(x_t)$ and $\overline{\phi}_t(x_t)$ are defined in Theorem 1.

Theorem 1 and Corollary 1 reveal a non-monotone relationship between the optimal charging quantity q_t^* and the peak usage level ϕ_t, and Figure 1 illustrates how the three charging regimes appear across the state space. This non-monotone relationship is intriguing. Given a fixed inventory of fully charged batteries, the rationale for intentionally increasing the peak usage level during the aggressive charging regime seems counterintuitive, as raising the peak usage level increases the demand charge. Conversely, while one might expect a higher peak usage level ϕ_t to justify charging depleted batteries more aggressively, it is actually optimal to charge less in the conservative regime as ϕ_t increases.

The core insight driving this non-monotone behavior is that the peak usage level, ϕ_t, represents not only a financial cost but also an effective capacity for producing fully charged batteries. Consequently, the determination of q_t^* serves two distinct functions: replenishing the inventory of fully charged batteries and strategically acquiring (or maintaining) production capacity.

When ϕ_t is sufficiently low, the BSS effectively "purchases" additional production capacity by increasing the peak usage level to a higher fixed level. In contrast, when ϕ_t is sufficiently high, the BSS already possesses enough capacity to handle future demand variations. In this regime, it is optimal to charge less and reserve depleted batteries for future periods, allowing the system to capitalize on free solar energy or lower grid electricity prices. Finally, when ϕ_t is neither too high nor too low, the optimal strategy is to utilize the existing production capacity to its fullest extent without exceeding it, thereby avoiding additional demand charges.

Corollary 1 is a direct outcome of Theorem 1, and establishes strict bounds on the sensitivity of q_t^* to changes in ϕ_t: the marginal change is exactly zero in the aggressive charging regime, one in the economic charging regime, and non-positive in the conservative charging regime.

Next, we investigate how the optimal charging quantity q_t^* varies with the number of fully charged batteries x_t.

Proposition 3 (**Sensitivity of q_t^* to x_t**). *Assume $\delta_t K < B - x_t^+$. The optimal charging quantity q_t^* is decreasing in the number of fully charged batteries x_t, and $-1 \leq \frac{\partial q_t^*}{\partial x_t} \leq 0$;*

Proposition 3 shows that the marginal effect of increasing the inventory of fully charged batteries on the optimal charging quantity is non-positive. This result is intuitive: a higher inventory level serves as a hedge against future demand uncertainty, thereby reducing the immediate necessity for charging. Proposition 3 also shows that the optimal charging quantity decreases by at most one unit with every additional unit of fully charged battery inventory. This structural property is particularly valuable for computational purposes, as it facilitates the design of efficient heuristics to solve the dynamic programming model. We will discuss more details in Section 4.

3.3. Model Without Demand Charges

In this section, we consider a simplified version of the model introduced in Section 3.1. Charging a depleted battery still requires one full period, but no demand charges are imposed. With time-varying electricity prices and no demand charges, this formulation reflects BSS operations in China. Using the same notation as in Section 3.1, the optimality equations are given by:

$$V_t(x_t) = \mathbb{E}_{\Delta_t}\left[\max_{q_t \in \mathcal{Q}(x_t)} \left\{ \mathbb{E}_{d_t}\left[pd_t - c_t(q_t - \delta_t K)^+ - \theta q_t - s(d_t - x_t)^+ + V_{t+1}(x_{t+1})\right]\right\} \middle| \Delta_t = \delta_t\right], \tag{5}$$

$$V_T(x_T) = px_T, \tag{6}$$

which are similar to Equations (3) and (4) except that $D(\phi_T) = 0$ and we do not need to track the peak usage level ϕ_t. The dynamics of the system state x_t is given in (1).

If we further assume that renewable energy is unavailable for battery charging, our model in Equations (5) and (6) reduces to a class of inventory models with limited production capacities. Federgruen and Zipkin (1986) analyze inventory management under a constant production capacity constraint. They prove that a modified base-stock policy is optimal. Similarly, Ciarallo et al. (1994) show that a base-stock policy remains optimal even when the production capacity is uncertain. However, unlike these models where capacity is exogenous, the charging capacity (analogous to production capacity) in our model is endogenous and dependent on the number of fully charged batteries (analogous to the inventory level), i.e., $q_t \le B - x_t^+$. Interestingly, we show in the following proposition that a modified base-stock policy remains optimal, even in the general case with renewable energy. Let $q_t^{\#}$ denote the optimal charging quantity when demand charges are absent.

Proposition 4 **(Optimal Policy without Demand Charges)**. *When demand charges are absent, the optimal charging policy for the problem defined in* (5) *and* (6) *is a modified base-stock policy. More specifically, there exists a threshold, denoted by $Q_t^* \le B$, such that the optimal charging quantity is $q_t^{\#} = \min\left\{\max\left\{Q_t^* - x_t, \delta_t K\right\}, B - x_t^+\right\},\ \forall t \in \mathcal{T}$.*

Proposition 4 characterizes the optimal charging quantity $q_t^{\#}$ when demand charges are absent. The expression of $q_t^{\#}$ implies that it is optimal to exhaust all available solar energy generated in period t to charge the depleted batteries, similar to the result for the model with demand charges (Proposition 2). The optimal charging policy is a modified base-stock policy, where charging decisions are governed by an order-up-to level Q_t^*, adjusted to account for exogenous solar energy and available battery capacity. The threshold Q_t^* is independent of the system state x_t. Therefore, one can design efficient algorithms to solve it. In our numerical study in Section 5, the optimal charging policy for the model without demand charges will be used as a benchmark policy to demonstrate the benefit of considering demand charges in battery charging decisions.

Intuitively, one might expect that the introduction of demand charges would uniformly discourage grid charging. In other words, one might expect q_t^* to always be no greater than $q_t^{\#}$. Interestingly, our next proposition shows that this is not always the case.

Proposition 5 **(Comparison of Optimal Charging Quantities)**. *When the peak usage level is sufficiently high, specifically in the conservative charging regime ($\phi_t \geq \overline{\phi}_t(x_t)$), the optimal charging quantity under demand charges can exceed the optimal charging quantity in the model without demand charges, i.e., $q_t^* \geq q_t^{\#}, \forall t \in \mathcal{T}$. Moreover, we have $\lim_{\phi_t \to B} q_t^* = q_t^{\#}$.*

Proposition 5 reveals an interesting behavior: when the existing peak usage level is sufficiently high (e.g., when $\phi_t \geq \overline{\phi}_t(x_t)$), the optimal charging quantity under demand charges (q_t^*) can actually exceed the quantity chosen in a model where demand charges are absent ($q_t^{\#}$). This comparison complements the threshold structure in Theorem 1 and the sensitivity properties in Corollary 1 and Proposition 3.

This counterintuitive result again stems from the dual role of the peak usage level ϕ_t: it acts not only as a penalty cost, but also as an effective charging capacity for producing fully charged batteries. When demand charges are absent, the BSS determines its charging quantity by balancing the current inventory of fully charged batteries against future demand fluctuations. Conversely, under demand charges, a high ϕ_t implies that the station has already established a high peak usage level. Consequently, the marginal demand charge for drawing any additional grid power up to ϕ_t is effectively zero. Because this production capacity has already been “paid for”, the BSS has an incentive to utilize it to aggressively increase the inventory of fully charged batteries. This stockpiling serves as a valuable buffer against future swapping demand surges or intermittent drops in solar generation, ultimately leading to a higher optimal charging quantity.

Finally, Proposition 5 shows that as the peak usage level ϕ_t approaches the upper limit B, the two policies converge ($\lim_{\phi_t \to B} q_t^* = q_t^{\#}$). At this boundary, the demand charge becomes a fixed cost that can no longer be influenced by marginal charging decisions. Hence, the station's decision reverts to balancing inventory and energy prices, as in the model without demand charges.

4. The Multi-Period Charging Problem and Heuristic Policies

While the single-period model introduced in Section 3—hereafter referred to as the *base model*—is analytically tractable and provides fundamental insights into the structure of optimal charging policies, it abstracts away from certain operational realities. In practice, BSS operations often involve higher-frequency decision making; that is, a station may adjust its charging strategy every 10 or 15 minutes, whereas fully recharging a depleted battery typically requires approximately one hour.

Next, we extend our base model to a multi-period setting to capture these dynamics in Section 4.1. Because charging a depleted battery now spans multiple decision periods, the system must explicitly track the SoC of all batteries currently in the charging pipeline. This expansion significantly increases the dimensionality of the state space, which makes the problem analytically and computationally challenging. Hence, in Sections 4.2 and 4.3, we propose several heuristic policies that leverage structural insights identified from the base model to efficiently and effectively solve the multi-period charging problem.

4.1. The Multi-Period Charging Model

We first describe the multi-period charging model in detail. We divide each period of the base model in Section 3 into L shorter decision periods. As a result, a battery that requires one period to charge in the base model now requires L consecutive periods to become fully charged. Throughout this section, unless otherwise specified, a period refers to one such shorter decision period. This discretization increases the decision frequency by a factor of L. Let $\widehat{\mathcal{T}}$ denote the set of decision epochs, i.e., $\widehat{\mathcal{T}} = \{0, 1, \dots, \widehat{T} - 1\}$. For instance, when each period represents 10 minutes (i.e., $L = 6$), we have the cardinality of $\widehat{T} = 30 \times 24 \times 6 = 4,320$ for a 30-day month. The unit waiting cost, energy price, and demand in period t are denoted by $\hat{s}$, $\hat{c}_t$, and $\hat{d}_t$, respectively. The battery degradation cost per period per pack is denoted by $\hat{\theta}$. Hence, we have $\hat{\theta} = \theta / L$.

At the beginning of each period t, the station observes the system state $\mathbf{S}_t = (\mathbf{x}_t, \phi_t^m)$, where ϕ_t^m is the peak usage level, and $\mathbf{x}_t = (x_{0,t}, x_{1,t}, \dots, x_{L-1,t})$ represents the battery inventory, including those fully charged and those in the charging pipeline. More specifically, $x_{0,t}$ is the number of fully charged batteries available for swapping, while $x_{i,t} \geq 0$ for $i \in \{1, \dots, L-1\}$ denotes the number of batteries that require exactly i additional periods to become fully charged. Note that $x_{0,t}$ can take negative values, in which case $-x_{0,t}$ represents the number of EVs waiting in line for battery swapping. The station also observes the realized amount of renewable energy, $\delta_t K$. For notational compactness, define $Y(\mathbf{x}_t) \equiv \sum_{i=1}^{L-1} x_{i,t}$ as the total number of batteries in the charging pipeline at the beginning of period t.

Based on this information, the station determines the number of depleted batteries to initiate for charging, denoted by q_t^m. Let $\bar{q}^m(\mathbf{x}_t) \equiv B - x_{0,t}^+ - Y(\mathbf{x}_t)$ denote the maximum feasible number of newly initiated charges, and let $\mathcal{Q}^m(\mathbf{x}_t) \equiv [0, \bar{q}^m(\mathbf{x}_t)]$ denote the feasible action set. Thus, $q_t^m \in \mathcal{Q}^m(\mathbf{x}_t)$. Subsequently, the stochastic demand $\hat{d}_t$ is realized, resulting in swap revenue or incurring waiting costs if the demand exceeds the available supply $x_{0,t}$. The state transitions are governed by the following equations:

$$x_{0,t+1} = x_{0,t} - \hat{d}_t + x_{1,t}, \tag{7}$$

$$x_{i,t+1} = x_{i+1,t}, \quad \text{for } i = 1, \dots, L-2, \tag{8}$$

$$x_{L-1,t+1} = q_t^m, \tag{9}$$

$$\phi_{t+1}^m = \max\left\{\phi_t^m, \left(q_t^m + Y(\mathbf{x}_t) - \delta_t K\right)^+\right\}. \tag{10}$$

The BSS aims to maximize its total expected profit over the planning horizon by determining the charging quantity q_t^m in each period. Let $V_t^m(\mathbf{S}_t)$ denote the optimal value function at the beginning of period t. The Bellman equations are given by:

$$\begin{aligned} V_t^m(\mathbf{x}_t, \phi_t^m) = \mathbb{E}_{\Delta_t}\Bigg[\max_{q_t^m \in \mathcal{Q}^m(\mathbf{x}_t)} \Bigg\{ \mathbb{E}_{\hat{d}_t}\Big[p\hat{d}_t - \hat{s}(\hat{d}_t - x_{0,t})^+ - \hat{c}_t\Big(q_t^m + Y(\mathbf{x}_t) - \delta_t K\Big)^+ \\ - \hat{\theta}\Big(q_t^m + Y(\mathbf{x}_t)\Big) + V_{t+1}(\mathbf{x}_{t+1}, \phi_{t+1}^m)\Big]\Bigg\}\Bigg|\Delta_t = \delta_t\Bigg], \end{aligned} \tag{11}$$

$$V_{\widehat{T}}^m(\mathbf{x}_{\widehat{T}}, \phi_{\widehat{T}}^m) = \sum_{i=0}^{L-1} \frac{L-i}{L} p x_{i,\widehat{T}} - D(\phi_{\widehat{T}}^m). \tag{12}$$

The first term in Equation (12) represents the prorated terminal revenue of the battery inventory. Here, a battery that is i periods away from being fully charged is valued at a fraction $(L-i)/L$ of the full swap price p. The second term represents the demand charge incurred at the end of the billing cycle based on the peak usage level recorded across all periods.

The multi-period charging problem poses a significantly greater challenge than the single-period formulation for three reasons. First, the problem becomes high-dimensional and suffers from the well-known curse of dimensionality. Second, unlike the single-period model, where demand charges depend solely on contemporaneous actions, decisions in the multi-period setting affect both immediate operating costs and future demand charges through the pipeline of partially charged batteries. This is because the charging process spans multiple periods, the current distribution of batteries across charging stages and new charging decisions jointly determine future peak usage levels. Third, as the state space expands, the value function loses its submodularity, making it analytically difficult to fully characterize the optimal charging policy. Consequently, we can only establish some basic structural properties of the optimal value function.

Proposition 6 **(Properties of Optimal Value Functions under Multi-Period Charging)**. *For any $t \in \widehat{\mathcal{T}}$, the optimal value function $V_t^m(\mathbf{x}_t, \phi_t^m)$ of the multi-period charging problem is (i) increasing in the number of fully charged batteries $x_{0,t}$ and decreasing in the peak usage level ϕ_t^m; and (ii) jointly concave in $(\mathbf{x}_t, \phi_t^m)$.*

Proposition 6(i) shows that having more fully charged batteries increases the optimal value, while a higher peak usage level reduces it. This is consistent with the findings from the base model. Intuitively, a larger number of fully charged batteries improves the provider's ability to fulfill demand and thus improves value, whereas a higher peak usage level raises future demand charges and thus reduces the optimal value. Proposition 6(ii) further establishes that the optimal value function is jointly concave in $(\mathbf{x}_t, \phi_t^m)$. This implies that when battery supply is already sufficient, the incremental value of additional charged batteries declines, and similarly, the economic benefit of lowering peak usage level becomes less pronounced once the peak usage level is already low.

Although Proposition 6 establishes the monotonicity and joint concavity of the optimal value functions, calculating the exact optimal solution remains computationally prohibitive. Hence, we devise a suite of heuristic policies, which leverage the structural insights identified from our base model in Section 3.2 to reduce the computational complexity.

4.2. The Multi-Period Charging Problem without Demand Charges

In the absence of demand charges, the model defined by (11) and (12) simplifies, as it is no longer necessary to track the dynamics of the peak usage level ϕ_t^m governed by (10). This multi-period charging problem shares structural similarities with periodic-review inventory systems with lead times. In standard backlogged inventory models, it is well-established that the optimal policy depends solely on the aggregate inventory position (Karlin and Scarf 1958). However, this property does not directly carry over to the BSS context for two reasons. First, unlike traditional inventory models, where pipeline inventory incurs costs upon arrival, batteries currently in the charging process consume electricity continuously. This ongoing charging load, combined with new charging decisions and intermittent solar generation, results in a non-separable and convex charging cost. Second, the charging quantity is naturally bounded above by the number of depleted batteries, effectively framing the problem as a system with limited production capacity. To address these complexities, we adapt the analytical framework of Zipkin (2008) to establish the monotonicity of the optimal charging quantity.

Proposition 7 **(Optimal Multi-Period Charging Policy without Demand Charges)**. *When demand charges are absent, the optimal charging quantity, denoted by $q_t^{m\#}(\mathbf{x}_t)$, is decreasing in $x_{i,t}$ for $i = 0, 1, \ldots, L-1$ and $t \in \widehat{\mathcal{T}}$. Moreover, the following inequalities hold:*

$$-1 \leq \frac{\partial q_t^{m\#}}{\partial x_{L-1,t}} \leq \cdots \leq \frac{\partial q_t^{m\#}}{\partial x_{1,t}} \leq \frac{\partial q_t^{m\#}}{\partial x_{0,t}} \leq 0. \tag{13}$$

Proposition 7 yields two operational insights. First, the optimal charging quantity decreases with the number of batteries at any charging stage, with a maximum reduction of one unit per additional battery. Second, the inequalities in (13) imply that the optimal charging quantity is more sensitive (i.e., with greater magnitudes) to batteries at earlier charging stages (e.g., $x_{L-1,t}$) than those at later charging stages (e.g., $x_{0,t}$). Batteries at earlier charging stages occupy charging capacity for a longer period of time and have a greater effect on the future inventory of fully charged batteries. Thus, they exert a larger negative effect on the current optimal charging quantity.

The optimal policy derived under this relaxed setting can serve as a heuristic policy, which we hereafter refer to as the **No-Demand-Charge Policy (Policy ND)**. This monotonicity property is the key to efficiently computing Policy ND. Even without the peak usage level as a state variable, the state vector of batteries $\mathbf{x}_t$ remains high-dimensional. However, the bounded marginal effects in (13) provide strict analytical bounds on how the optimal charging quantity shifts across adjacent states. For instance, if the optimal charging quantity $q_t^{m\#}$ is already identified for a given state $\mathbf{x}_t$, the optimal charging quantity for a neighboring state with one additional battery at stage i is mathematically restricted: it must be either $q_t^{m\#}(\mathbf{x}_t)$ or $q_t^{m\#}(\mathbf{x}_t) - 1$. This significantly reduces the search space and thus the computational complexity, making Policy ND a highly tractable and scalable heuristic for the multi-period charging problem with demand charges.

Notes. We developed Policy ND as a heuristic for the setting with demand charges. However, Policy ND is in fact optimal for BSS operations in China, where electricity prices vary over time but no demand charges are imposed.

4.3. Heuristic Policies

In addition to the benchmark Policy ND, we consider two classes of heuristic policies for the multi-period charging problem. The first class consists of structure-informed policies that translate the single-period structural characterization into implementable multi-period decision rules. The second class consists of simple benchmark policies that are intuitive and easy to implement but are not derived from the structural properties of the optimal charging policy. We use the feasible action set $\mathcal{Q}^m(\mathbf{x}_t)$ and its upper bound $\bar{q}^m(\mathbf{x}_t)$, as defined in Section 4.1, in the heuristic descriptions below.

4.3.1. Structural-Threshold Policy (Policy ST): Policy ST is a multi-period extension of the single-period structural characterization in Theorem 1. Rather than computing the exact single-period thresholds, the policy uses the two heuristic cutoffs for the peak usage level, $\underline{\phi}_t^H$ and $\overline{\phi}_t^H$, defined in Appendix B, to reproduce the three-regime shape illustrated in Figure 2(a).

When the current peak usage level is low, $0 \le \phi_t^m \le \underline{\phi}_t^H$, the policy applies the aggressive charging action $q_t^m(\mathbf{x}_t, 0)$, which is the charging quantity computed for the same battery-pipeline state when the peak usage level is zero. This rule reflects the low peak usage level case in the single-period model, where the optimal charging quantity is constant in the peak usage level. When the peak usage level is intermediate, $\underline{\phi}_t^H < \phi_t^m \le \overline{\phi}_t^H$, the policy chooses $\min\{\hat{q}_t^m, \bar{q}^m(\mathbf{x}_t)\}$, where $\hat{q}_t^m = (\phi_t^m + \delta_t K - Y(\mathbf{x}_t))^+$.

This action uses the grid load already allowed by the established peak usage level while avoiding, whenever feasible, a higher peak usage level.

When the peak usage level exceeds $\overline{\phi}_t^H$, the single-period theory suggests that the optimal charging quantity should eventually decrease and converge to the no-demand-charge benchmark $q_t^{m\#}$. Because the exact decreasing path is difficult to characterize in the multi-period model, Policy ST uses the local search procedure in Algorithm 1. The local search is motivated by the bounded inventory sensitivity in Proposition 3: adding one fully charged or in-process battery reduces the optimal charging quantity by at most one unit. Therefore, instead of searching over the full feasible action set, the policy compares only two neighboring candidate actions, the action from a nearby state with one fewer battery in the first nonempty charging stage and that action minus one. Once the action reaches $q_t^{m\#}$, Policy ST keeps using $q_t^{m\#}$ for higher peak usage levels. The detailed implementation is provided in Algorithm 2 in Appendix B.

Figure 2 An Illustration of Policy ST and Policy CB

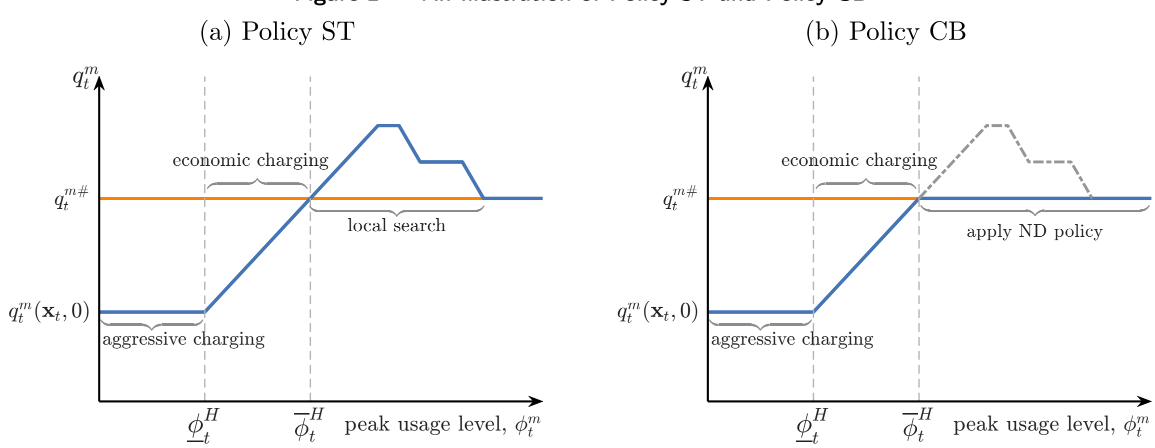

4.3.2. Capped-Baseline Policy (Policy CB): Policy CB is a simplified version of Policy ST. It preserves the same low- and intermediate-peak rules: it uses $q_t^m(\mathbf{x}_t, 0)$ when the peak usage level is below $\underline{\phi}_t^H$, and it uses $\min\{\hat{q}_t^m, \bar{q}^m(\mathbf{x}_t)\}$ when the peak usage level lies between $\underline{\phi}_t^H$ and $\overline{\phi}_t^H$. The simplification occurs in the high-peak region. Instead of using local search to trace the decreasing transition suggested by the single-period policy, Policy CB directly sets the charging quantity to the no-demand-charge benchmark $q_t^{m\#}$ once $\phi_t^m > \overline{\phi}_t^H$. Thus, as illustrated by the dotted line in Figure 2(b), Policy CB flattens the high-peak transition at the benchmark level. This design is computationally simpler than Policy ST, but it uses a coarser approximation of the conservative charging region. The detailed implementation is provided in Algorithm 3 in Appendix B.

We next introduce three simple benchmark policies. These policies provide transparent operational rules and useful comparison points in the numerical study. However, because they do not exploit the threshold structure, peak-sensitivity results, or bounded inventory-sensitivity results established above, their performance is not guaranteed by the structural analysis and must be assessed computationally.

- **Full-Capacity Policy (Policy FC):** This aggressive benchmark charges every feasible depleted battery in period t. The charging quantity is therefore $q_t^m = \bar{q}^m(\mathbf{x}_t)$. This rule is simple and maximizes immediate charging activity, but it does not account for the consequences of demand charges when establishing a new peak usage level.
- **Peak-Filling Policy (Policy PF):** This inventory-oriented benchmark attempts to raise the total number of batteries, including both fully charged and in-process batteries, to a target demand level. Let $d_{\max}$ denote an upper demand benchmark, calibrated from historical observations of high demand or bounded by the physical service capacity of the station. The policy sets $q_t^m = \min\left\{(d_{\max} - x_{0,t} - Y(\mathbf{x}_t))^+, \bar{q}^m(\mathbf{x}_t)\right\}$. The rule is easy to implement, but it is based on aggregate inventory coverage rather than the structural results considering demand charges.
- **Demand-Matching Policy (Policy DM):** This demand-charge-oriented benchmark avoids exceeding the current peak usage level whenever possible. Define $\hat{q}_t^m = \left(\phi_t^m + \delta_t K - Y(\mathbf{x}_t)\right)^+$ as the charging quantity that would match the current peak usage level after accounting for renewable generation and the existing charging pipeline. The policy sets $q_t^m = \min\{\hat{q}_t^m, \bar{q}^m(\mathbf{x}_t)\}$. Although this rule controls exposure to additional demand charges, it does not capture the aggressive, economic, and conservative transition implied by the structural results.

Finally, we consider two learning-based algorithms for settings in which the multi-period state space becomes too large for heuristics that are based on dynamic programming to remain computationally practical. These policies learn charging decisions from simulated system trajectories while enforcing action feasibility through the state-dependent action set $\mathcal{Q}^m(\mathbf{x}_t)$.

- **Proximal Policy Optimization (PPO):** PPO is a model-free deep reinforcement learning algorithm that learns a mapping from the system state $\mathbf{S}_t$ to a probability distribution over feasible charging quantities. The policy is trained using an actor-critic architecture and a clipped surrogate objective, which stabilizes policy updates while allowing exploration of high-dimensional state spaces (Jin et al. 2023). In our setting, PPO serves as a purely learning-based benchmark: it uses the feasible action mask implied by $\bar{q}^m(\mathbf{x}_t)$, but it does not use the structural results from the single-period model.
- **Structural-Threshold Proximal Policy Optimization (ST-PPO):** ST-PPO is a structure-guided variant of PPO. It augments the PPO objective with an imitation-learning loss that encourages the learned policy to choose actions close to those suggested by Policy ST. In this way, ST-PPO uses the threshold-based logic of the single-period problem as expert guidance while retaining PPO's flexibility to adapt to high-dimensional, multi-period dynamics. The detailed procedure is provided in Algorithm 4 in Appendix C.

5. A Case Study

In this section, we use a case study calibrated with battery-swapping demand observations, Dutch electricity prices, local solar generation, and the demand charge tariff faced by a NIO battery swapping station to evaluate the practical value of the proposed model and heuristic policies. The case study serves three purposes. First, it quantifies the economic loss from ignoring demand charges when making charging decisions. Second, it examines how renewable capacity and battery capacity affect the peak usage level and profitability, thereby illustrating the dual role of the peak usage level as both a cost state for demand charges and an implicit charging capacity state. Third, it evaluates whether the proposed structure-informed policies provide scalable and effective decision rules in a realistic multi-period setting.

5.1. Data and Parameter Estimation

NIO is among the most visible commercial operators of EV battery swapping. Its Power Swap network was first deployed at scale in China and has subsequently been extended to Europe, where battery swapping serves as a rapid refueling complement to conventional charging (NIO 2026). The Netherlands provides a relevant empirical setting because NIO has deployed swap infrastructure there, including the Tilburg station used as the basis for our calibration (Meadows 2022), and because Dutch commercial electricity tariffs combine volatile wholesale energy prices with capacity-based network charges. This setting therefore provides a realistic context in which charging decisions, local renewable generation, and the demand charge induced by peak usage interact.

We calibrate the base setting to the NIO BSS in Tilburg, the Netherlands. The Tilburg site anchors the geographic inputs used in the case study, including Dutch electricity prices, Dutch demand charges, and North Brabant solar generation. To evaluate operations at a larger station scale, we adopt the specifications of NIO's third-generation swap station, which can accommodate $B = 21$ batteries; each battery is assumed to have an energy capacity of 75 kWh (Meadows 2022). The battery degradation cost is estimated at \$0.04 per kWh (Gowda et al. 2021), which gives a unit degradation cost of $\theta = 0.04 \times 75 = \$3$ per charging cycle.

Grid energy prices are calibrated from a Dutch dynamic electricity contract. Under the contract structure offered by Vandebron, a Dutch energy supplier, the retail price combines the hourly wholesale electricity price, a service fee, and an energy tax. We use August 2024 as the representative billing month. During this month, the average energy tax and service fee are \$0.135/kWh and \$0.0615/kWh, respectively (Vandebron 2024), and the hourly wholesale prices are obtained from Ember's European Electricity Prices and Costs database (Ember 2025). We compute c_t by multiplying the sum of the wholesale price, service fee, and energy tax by 75 kWh. The resulting hourly charging costs over the month are shown in Figure 6 in Appendix D.

In addition to energy usage costs, the station is subject to demand charges. We calibrate this component using Enexis's commercial tariff (Enexis Netbeheer 2025), under which customers with contracted capacity above 125 kW pay a monthly demand charge of Euro 6.15 per kW. The BSS considered in our case study falls into this customer category. Using an exchange rate of 1 EUR = 1.10 USD, this rate is equivalent to \$6.77 per kW per month. In our model, one unit of peak usage level corresponds to one 75-kW charging bay drawing power from the grid during a period. Thus, under the baseline parameter setting, we use the linear demand charge function $D(\phi) = 6.77 \times 75 \times \phi = 507.75\phi$ USD.

We next calibrate the renewable generation process. Motivated by recent PV-enabled battery swapping projects at NIO stations (LONGi 2026), we assume that the station is connected to a local solar microgrid. In the numerical experiments, we vary K to evaluate the operational value of solar capacity. We estimate the stochastic process Δ_t using hourly solar generation data for North Brabant, the province in which Tilburg is located (Schepel et al. 2020). The data are normalized by the maximum observed generation to obtain the hourly solar availability coefficient Δ_t. Following the renewable energy literature (Wu and Kapuscinski 2013, Zhou et al. 2019, Kouvelis et al. 2023), we decompose the solar availability coefficient as $\Delta_t = a_t^d + a_t^m + A_t$, where a_t^d and a_t^m capture deterministic intraday and intramonth patterns, respectively, and A_t captures the remaining stochastic fluctuation. We model the residual component as an $AR(1)$ process, $A_t = \zeta A_{t-1} + \epsilon_t$, where $\zeta > 0$ and ϵ_t is normally distributed white noise with mean zero and standard deviation σ. Figure 3 reports the resulting decomposition. The estimated stochastic component is well described by an $AR(1)$ process with autoregressive coefficient $\zeta = 0.80$ and standard deviation $\sigma = 0.06$.

Figure 3 Estimation of Solar Energy Generation Coefficient (August 2024)

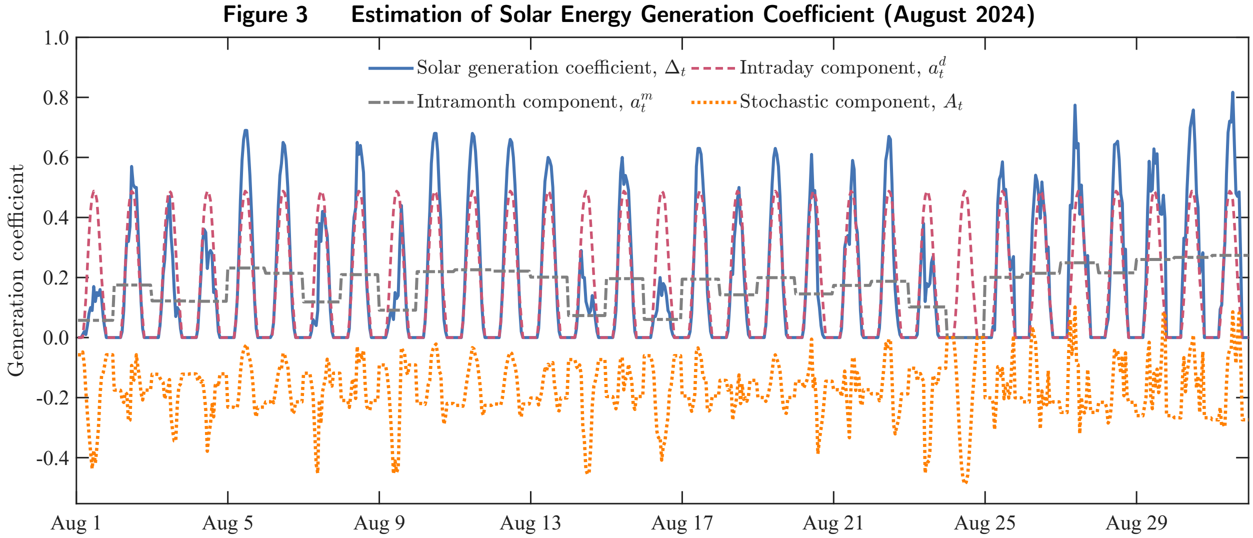

As of March 2025, NIO charges Euro 0.29/kWh for battery swaps at its stations in the Netherlands (Afonso 2025). Assuming a 75-kWh battery pack and an exchange rate of 1.1 USD/EUR, we set the revenue per swap to $p = 0.29 \times 1.1 \times 75 \approx \$24$. If demand in a period exceeds the available inventory of fully charged batteries, the station incurs a waiting cost for each unmet swap. Following Knoope (2023), who reports an average value of travel time of Euro 10.42 per hour per person for car users in the Netherlands, we set the waiting cost s to be \$11.5 per hour.

Finally, we calibrate the swapping demand process. Direct demand data for NIO's Tilburg swap station are not publicly available. To obtain a realistic intraday and day-of-week demand profile, we therefore use observed demand data from NIO swap stations in Shanghai, China, where battery swapping is more mature. We model arrivals as a time-dependent Poisson process. For each hour of the week, we estimate the average swap demand from the observed data and use this hourly average as the arrival rate for that hour. Figure 4 reports the estimated arrival rates.

Figure 4 **Hourly Swapping Demand to Our Study Swap Station by Time of Day and Day of the Week**

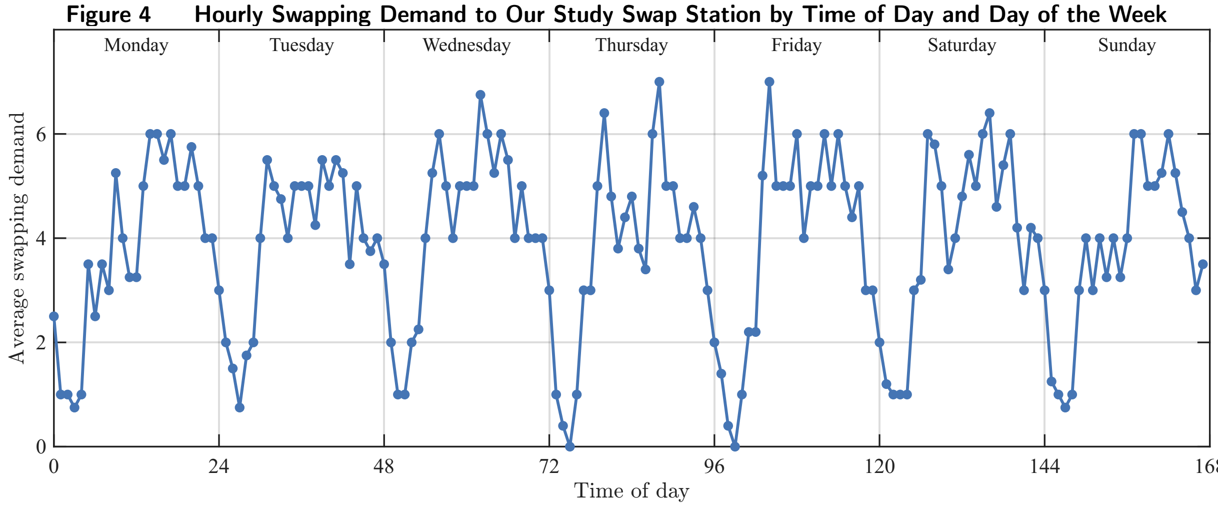

5.2. Value of Incorporating Demand Charges into Charging Decisions

We first isolate the operational value of incorporating demand charges into charging decisions. This experiment uses the empirical calibration from Section 5.1 and the single-period charging model from Section 3. We compare the optimal charging policy in Theorem 1, which considers demand charges, with Policy ND obtained from Proposition 4 by optimizing charging decisions as if demand charges were absent and is then evaluated under the actual model with demand charges. This design measures the economic loss from treating demand charges as an ex post billing item rather than as an operational state that should guide charging decisions.

Table 1 shows that demand charge awareness has substantial and heterogeneous operational value across station designs. The losses are generally small when the two policies have little room to differ, as in the smallest station with $B = 5$ and low solar capacity. Once the station has more operational flexibility, however, ignoring demand charges can be costly. Under the calibrated tariff

Table 1 Percentage Profit Loss from Ignoring Demand Charges across Station and Solar Capacity Levels

	$B=5$					$B=13$				
Solar Capacity K	0	5	10	15	20	0	5	10	15	20
$D(\phi)=200\phi$	0.00	0.00	0.02	1.08	0.98	7.84	1.90	1.18	1.16	1.15
$D(\phi)=400\phi$	0.00	0.00	0.32	19.01	5.71	25.22	4.83	3.06	2.86	2.82
$\boldsymbol{D(\phi)=507.75\phi}$	**0.00**	**0.00**	**0.95**	**76.30**	**10.18**	**39.54**	**6.85**	**4.12**	**3.89**	**3.85**
$D(\phi)=600\phi$	0.00	0.00	1.84	761.15	15.52	57.18	8.66	5.08	4.86	4.80
$D(\phi)=800\phi$	0.00	0.00	3.93	66.32	31.77	128.98	12.86	7.56	7.05	6.90
	$B=21$					$B=28$				
Solar Capacity K	0	5	10	15	20	0	5	10	15	20
$D(\phi)=200\phi$	17.69	5.70	2.30	1.63	1.34	22.49	7.56	3.91	1.88	1.24
$D(\phi)=400\phi$	44.01	12.78	5.11	3.75	3.13	55.76	17.14	8.37	4.16	2.87
$\boldsymbol{D(\phi)=507.75\phi}$	**62.19**	**17.07**	**6.81**	**4.94**	**4.15**	**77.90**	**22.81**	**10.84**	**5.44**	**3.79**
$D(\phi)=600\phi$	80.23	20.87	8.28	5.97	5.04	99.76	27.86	13.07	6.55	4.59
$D(\phi)=800\phi$	130.18	29.50	11.53	8.24	7.02	158.11	39.31	18.01	8.98	6.33

Notes. Each entry reports $100 \times (V - V^{\mathrm{ND}})/|V|$, where V is the expected profit under the optimal charging policy and V^{ND} is the expected profit under the policy that ignores demand charges (Policy ND). The initial conditions are $x_0 = 0$ and $\phi_0 = 0$. The values of B correspond to the capacity of NIO's first-generation BSS (5 packs), second-generation BSS (13 packs), third-generation BSS (21 packs), and fifth-generation BSS (27–28 packs). The fourth-generation BSS has 23 packs and is not included in the numerical experiments. Shaded entries indicate configurations in which $V < 0$, so that the percentage loss is less meaningful.

$D(\phi) = 507.75\phi$ and $K = 0$, the loss is 39.54% for $B = 13$, 62.19% for $B = 21$, and 77.90% for $B = 28$. These losses fall sharply as solar capacity increases: at $K = 20$, the corresponding losses decline to 3.85%, 4.15%, and 3.79%, respectively. Thus, solar capacity mitigates the profit loss from ignoring demand charges by reducing the need to rely on grid charging. The table also shows that the loss becomes more significant as the demand charge rate increases. For example, when $B = 21$ and $K = 0$, the loss rises from 17.69% under $D(\phi) = 200\phi$ to 130.18% under $D(\phi) = 800\phi$.

The operational mechanism behind these patterns aligns with the structural insights in Section 3.2. A policy that neglects demand charges treats peak usage as a passive accounting metric, leading the station to exploit low electricity prices whenever inventory capacity allows. In contrast, the optimal policy recognizes that current grid usage determines both immediate charging costs and future marginal demand charge exposure. These dynamics show that demand charges must be explicitly modeled in practical operations, particularly when solar generation is limited, station capacity offers high operational flexibility, or the demand charge tariff is high.

We next examine how the value of demand-charge-aware control varies with solar panel capacity and battery capacity. Figure 5 reports the profit loss from ignoring demand charges and the resulting peak usage levels under these capacity variations. Panels (a) and (b) show that solar panel capacity reduces the operational importance of demand charges. As K increases, the profit loss from ignoring demand charges declines, and the realized peak usage level also falls. Additional solar generation allows the station to charge more batteries without drawing from the grid, thereby directly reducing the need to establish high peak usage levels. This pattern is consistent with the renewable-first property: when renewable energy is available, the station can increase the inventory of fully charged batteries without increasing the peak usage level.

Panels (c) and (d) isolate the effect of station capacity while fixing $K = 10$. Panel (c) reports the absolute profit loss from ignoring demand charges, measured in thousand USD. The loss generally increases with B and with the demand charge rate. Panel (d) reports the resulting peak usage levels under the optimal charging policy that considers demand charges. The relationship between B and the peak usage level is non-monotone. As station capacity increases from a small level, the optimal charging policy may allow a higher peak usage level because the established peak creates charging headroom that can be reused later in the billing cycle. For larger values of B, however, the station has enough inventory flexibility to smooth charging more effectively, and the optimal peak usage level declines or stabilizes. Higher demand charge rates also lead to lower peak usage levels, reflecting the stronger incentive to avoid incurring costly peak usage from the grid.

5.3. Performance of Heuristic Policies

We next evaluate the heuristic policies proposed in Sections 4.2 and 4.3. The case study uses the demand charge function $D(\phi) = 507.75\phi$, which is calibrated to the Tilburg station. The solar capacity parameter is set to $K = 10$. Recall that K is measured by the number of batteries that can be fully charged by solar energy over one hour. When $L > 1$, $\Delta_t K$ should be interpreted as the number of active charging batteries whose energy consumption can be covered by solar generation during a decision period, rather than the number of batteries that can be fully charged within that period. Since each NIO battery has a capacity of 75 kWh (Meadows 2022), this corresponds to an installed solar capacity of $10 \times 75 = 750$ kW. PV-enabled NIO swap stations utilize commercial PV modules with a conversion efficiency of 24.3% (LONGi 2026). Under this specification, a 750 kW system requires approximately 3,086 m^2 of module surface area before accounting for spacing and site-specific installation constraints. Such capacity can be achieved through rooftop installations and adjacent parking-canopy PV deployment.

We first report exact dynamic programming benchmarks for $L = 1$ and $L = 2$, where the optimal policy can be computed exactly. We measure each heuristic by its optimality gap relative to the exact optimal policy, evaluated from the same initial state. Table 2 reports the expected profit, optimality gap, and runtime for both settings.

Table 2 shows that solving the dynamic programming remains feasible for $L = 1$ and $L = 2$, although the computation time rises sharply from 2.6 minutes to 49.4 minutes. Against this exact

Figure 5 Impact of Demand Charges under Varying Solar Panel and Battery Capacity

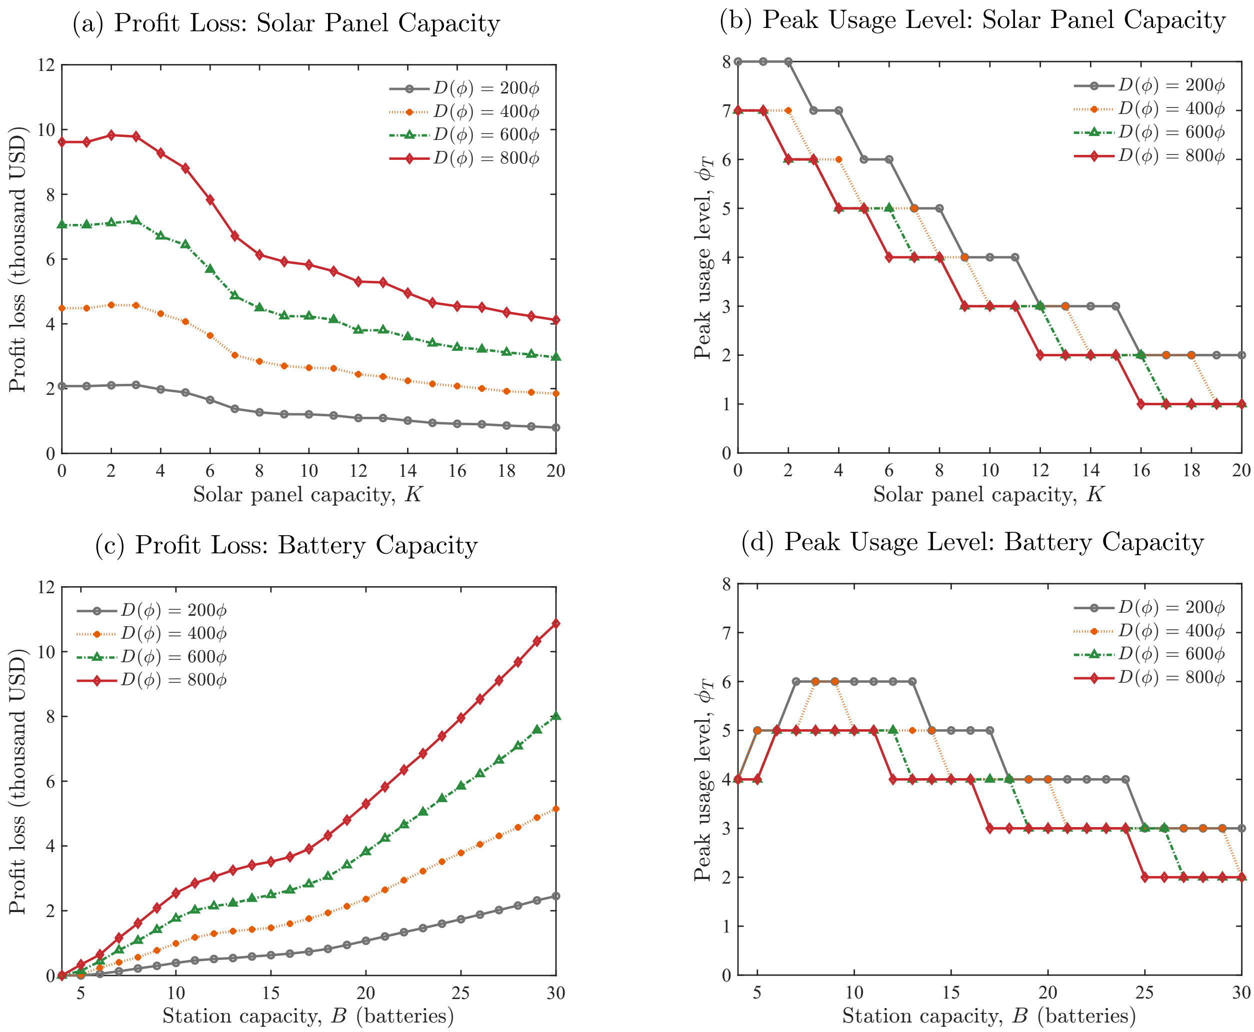

Notes. We set $x_0 = 0$, $\phi_0 = 0$, $B = 21$ for (a) and (b), and $K = 10$ for (c) and (d).

benchmark, the structure-informed policies perform especially well. Policy ST nearly replicates the optimal policy, with optimality gaps below 0.5% in both settings, while Policy CB remains within about 1%–2% of the optimal policy. In contrast, the simple benchmark policies perform substantially worse, especially when $L = 2$, with gaps of roughly 20%–31%. The learning-based policies also perform reasonably well relative to the simple benchmarks, but they require substantially longer runtimes than Policy ST and Policy CB. ST-PPO reduces the PPO optimality gap from 6.77% to 2.72% for $L = 1$ and from 6.58% to 2.61% for $L = 2$. Thus, when exact dynamic programming remains feasible, the structure-informed dynamic programming heuristics provide the best combination of accuracy and speed.

When the decision period is shortened further, exact dynamic programming is no longer computationally practical. For $L = 4$ and $L = 6$, corresponding to 15-minute and 10-minute decision

Table 2 Performance of Heuristic Policies at $L=1$ and $L=2$

Policy	$L=1$ Expected Profit (\$)	$L=1$ Optimality Gap (%)	$L=1$ Runtime (min)	$L=2$ Expected Profit (\$)	$L=2$ Optimality Gap (%)	$L=2$ Runtime (min)
Optimal Policy	51,753	-	2.6	53,489	-	49.4
Policy ST	51,745	0.02	0.3	53,237	0.47	4.5
Policy CB	50,928	1.59	0.2	52,940	1.03	4.2
Policy ND	48,035	7.18	0.1	48,739	8.88	4.1
Policy FC	39,901	22.90	0.0	40,003	25.21	0.0
Policy PF	38,165	26.25	0.0	37,018	30.79	0.0
Policy DM	-101,004	295.17	0.0	43,055	19.51	0.0
PPO	48,247	6.77	6.6	49,970	6.58	16.5
ST-PPO	50,347	2.72	7.8	52,095	2.61	24.2

Notes. $L=1$ corresponds to a one-hour decision period, and $L=2$ corresponds to a 30-minute decision period. The initial state is $\phi_0=0$ and B fully charged batteries.

periods, respectively, the state space becomes too large to compute the optimal policy. The structure-informed dynamic programming heuristics, Policy ST and Policy CB, also become impractical because their implementation requires solving and storing increasingly large dynamic programming tables. Therefore, Table 3 compares PPO and ST-PPO with the simple benchmark policies, rather than the optimal policy which is unavailable.

Table 3 Performance of Learning-Based Policies PPO and ST-PPO at $L=4$ and $L=6$

L	Policy	Expected Profit (\$)	Improvement over FC (%)	Improvement over PF (%)	Improvement over DM (%)	Runtime (min)
$L=4$	Policy FC	41,356	–	–	–	–
	Policy PF	38,589	–	–	–	–
	Policy DM	43,428	–	–	–	–
	PPO	52,119	26.07	35.11	20.03	32.6
	ST-PPO	53,163	28.59	37.82	22.44	37.1
$L=6$	Policy FC	41,717	–	–	–	–
	Policy PF	39,906	–	–	–	–
	Policy DM	43,068	–	–	–	–
	PPO	52,834	26.68	32.43	22.69	101.0
	ST-PPO	53,163	27.47	33.26	23.45	105.5

Notes. $L=4$ and $L=6$ correspond to 15-minute and 10-minute decision periods, respectively. Percentage improvements are computed relative to the expected profit of the corresponding benchmark policy within the same L setting. Runtime is reported for the learning-based policies.

The results in Table 3 show that both learning-based policies substantially outperform the simple benchmark policies when exact dynamic programming is no longer tractable. Across the two large-state settings, PPO improves expected profit by roughly 20%–35% relative to the simple benchmark policies, while ST-PPO raises these gains to roughly 22%–38%. The incremental value of structural guidance is positive but more moderate: ST-PPO improves expected profit over PPO by about 2.0% for $L=4$ and 0.6% for $L=6$, with only a small increase in runtime. Taken together, these results suggest that when $L=1$ or $L=2$, both exact dynamic programming and the structure-informed heuristics remain feasible, with Policy ST and Policy CB offering high accuracy at modest runtime. When the decision period becomes finer and dynamic-programming-based policies are no longer tractable, learning-based policies provide scalable alternatives, and structural imitation through ST-PPO yields consistent additional gains over PPO.

6. Conclusion

Battery swapping stations offer a promising solution for reducing EV refueling time, but their economic performance depends critically on how charging decisions are coordinated with renewable generation, stochastic swapping demand, and electricity tariffs. This paper studies the optimal charging problem for a renewable-integrated battery swapping station under demand charges. We formulate the problem as a finite-horizon dynamic program in which charging decisions affect not only current energy, degradation, and waiting costs, but also the peak usage level that determines the demand charge at the end of the billing cycle.

Our analysis shows that demand charges fundamentally change the structure of optimal charging policy. In the single-period charging model, we establish a renewable-first property and characterize a two-threshold policy structure that divides the state space into aggressive, economic, and conservative charging regimes. The economic mechanism behind this non-monotone structure is the dual role of the peak usage level ϕ_t: it is both a cost determinant and an implicit charging capacity state. When ϕ_t is low, increasing the peak usage level can be optimal because it creates additional charging headroom for the remainder of the billing cycle. When ϕ_t is intermediate, the station uses the existing peak usage level as fully as possible without increasing it. When ϕ_t is sufficiently high, the station already has enough charging headroom, so it becomes optimal to defer some charging and preserve depleted batteries for future periods with renewable generation or lower grid energy costs. This mechanism explains why the optimal charging quantity can be non-monotone in the peak usage level and, in some states, can exceed the corresponding no-demand-charge benchmark.

We then extend the analysis to a multi-period charging setting in which batteries require multiple decision periods to become fully charged. The resulting state space is high dimensional, making an exact structural characterization difficult. Nevertheless, the single-period structural results provide useful guidance for scalable decision rules. We develop structure-informed heuristic policies that translate the threshold logic of the single-period problem into the multi-period setting, and we further incorporate this structural guidance into a reinforcement-learning algorithm through

Structural-Threshold Proximal Policy Optimization. The numerical results show that the structure-informed policies achieve strong performance when $L = 1$ and $L = 2$, while the learning-based algorithms become valuable when finer decision periods make dynamic-programming-based methods computationally challenging.

The case study illustrates the operational value of modeling demand charges explicitly. Policies that ignore demand charges can establish unnecessarily high peak usage levels and suffer sizable profit losses. Renewable capacity mitigates demand-charge exposure by reducing reliance on grid charging, while station capacity has a more nuanced effect: additional batteries initially improve flexibility, but excessive capacity can also encourage aggressive charging if peak usage level is not managed carefully. These findings suggest that battery swapping operators should treat the peak usage level as an operational state variable rather than a passive billing outcome. Effective BSS operations require joint management of battery inventory, renewable generation, grid energy costs, and demand-charge exposure.

Several extensions remain promising for future research. One direction is to study networks of battery swapping stations under demand charges, in which battery inventory, customer demand, and charging capacity can be coordinated across locations. Another is to incorporate richer battery-level characteristics, such as heterogeneous states of charge, degradation histories, and battery health constraints. It would also be valuable to examine tariff designs with coincident peak charges, real-time demand charges, or demand-response incentives. These extensions would further clarify how battery swapping infrastructure can support both reliable EV service and efficient operations of electricity system.

References

C. Afonso. NIO permanently cuts battery swap and charging prices by 34% in the Netherlands. `https://eletric-vehicles.com/nio/nio-permanently-cuts-battery-swap-and-charging-prices-by-34-in-netherlands/`, March 2025. Accessed August 14, 2025.

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

S. Borenstein. The economics of fixed cost recovery by utilities. *The Electricity Journal*, 29(7):5–12, 2016.

CATL. BAIC Group and CATL join hands to build a new “integrated supercharging and battery-swapping” ecosystem. `https://www.catl.com/en/news/6891.html`, 2026. Accessed July 4, 2026.

S. Chawrasia, D. Hembram, D. Bose, and C. Chanda. Design of solar battery swapping station for EV using LSTM-assisted solar power forecasting. *Microsystem Technologies*, 30(9):1087–1098, 2024.

L. Chen, L. He, and Y. Zhou. An exponential cone programming approach for managing electric vehicle charging. *Operations Research*, 72(5):2215–2240, 2024.

China Daily. HK aims to be bridge for mainland automakers' global expansion. `https://www.chinadailyhk.com/hk/article/598824`, 2024. Accessed July 2, 2025.

F. W. Ciarallo, R. Akella, and T. E. Morton. A periodic review, production planning model with uncertain capacity and uncertain demand—optimality of extended myopic policies. *Management Science*, 40(3): 320–332, 1994.

Ember. European electricity prices and costs. `https://ember-energy.org/data/european-electricity-prices-and-costs/`, 2025. Accessed August 8, 2025.

Enexis Netbeheer. Periodieke tarieven elektriciteit 2025. `https://www.enexis.nl/tarieven/2025-elektriciteit-eenmalig`, 2025. Accessed August 11, 2025.

S. Esmaeili, A. Anvari-Moghaddam, and S. Jadid. Optimal operation scheduling of a microgrid incorporating battery swapping stations. *IEEE Transactions on Power Systems*, 34(6):5063–5072, 2019.

A. Federgruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands II. The discounted-cost criterion. *Mathematics of Operations Research*, 11(2):208–215, 1986.

S. N. Gowda, B. A. Eraqi, H. Nazaripouya, and R. Gadh. Assessment and tracking electric vehicle battery degradation cost using blockchain. In *2021 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference*, pages 1–5. IEEE, 2021.

X. Hu, Z. Yang, J. Sun, and Y. Zhang. Optimal pricing strategy for electric vehicle battery swapping: Pay-per-swap or subscription? *Transportation Research Part E: Logistics and Transportation Review*, 171: 103030, 2023.

IEA. Global EV outlook 2025. `https://www.iea.org/reports/global-ev-outlook-2025`, 2025. Accessed July 2, 2025.

J. Jin and Y. Xu. Optimal storage operation under demand charge. *IEEE Transactions on Power Systems*, 32(1):795–808, 2017.

J. Jin, S. Mao, and Y. Xu. Optimal priority rule-enhanced deep reinforcement learning for charging scheduling in an electric vehicle battery swapping station. *IEEE Transactions on Smart Grid*, 14(6):4581–4593, 2023.

C. Kaps, S. Marinesi, and S. Netessine. When should the off-grid sun shine at night? Optimum renewable generation and energy storage investments. *Management Science*, 69(12):7633–7650, 2023.

S. Karlin and H. Scarf. Inventory models of the Arrow-Harris-Marschak type with time lag. In K. J. Arrow, S. Karlin, and H. Scarf, editors, *Studies in the Mathematical Theory of Inventory and Production*, chapter 10. Stanford University Press, Stanford, CA, 1958.

M. Knoope. New values of travel time, reliability and comfort in the Netherlands. Summary report, KiM Netherlands Institute for Transport Policy Analysis, The Hague, 2023.

A. G. Kök, K. Shang, and Ş. Yücel. Investments in renewable and conventional energy: The role of operational flexibility. *Manufacturing & Service Operations Management*, 22(5):925–941, 2020.

P. Kouvelis, H. Matsuo, Y. Xiao, and Q. Yuan. Long-term service agreement in electricity supply chain with renewable energy penetration. *Production and Operations Management*, 32(6):1830–1845, 2023.

Z. J. Lee, J. Z. Pang, and S. H. Low. Pricing EV charging service with demand charge. *Electric Power Systems Research*, 189:106694, 2020.

L. Li, Y. F. Lim, and K. Pan. Integrating EV charging and discharging into power grid through bilateral negotiation. *Available at SSRN 4899806*, 2024.

LONGi. LONGi and NIO launch integrated PV and charging station as blueprint for green mobility. `https://eu.longi.com/press/longi-and-nio-launch-integrated-pv-and-charging-station-as-blueprint-for-green-mobility`, Feb. 2026. Accessed March 14, 2026.

H.-Y. Mak, Y. Rong, and Z.-J. M. Shen. Infrastructure planning for electric vehicles with battery swapping. *Management Science*, 59(7):1557–1575, 2013.

S. Mao, J. Jin, and Y. Xu. Routing and charging scheduling for EV battery swapping systems: Hypergraph-based heterogeneous multiagent deep reinforcement learning. *IEEE Transactions on Smart Grid*, 15 (5):4903–4916, 2024.

D. Meadows. NIO opens first battery exchange station in the Netherlands. `https://techzle.com/nio-opens-first-battery-exchange-station-in-the-netherlands`, December 2022. Accessed August 8, 2025.

NIO. NIO achieves 100 million battery swaps. `https://www.nio.com/news/nio-achieves-100-million-battery-swaps`, 2026. Accessed February 28, 2026.

R. Passey, N. Haghdadi, A. Bruce, and I. MacGill. Designing more cost reflective electricity network tariffs with demand charges. *Energy Policy*, 109:642–649, 2017. ISSN 0301-4215.

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

V. Schepel, A. Tozzi, M. Klement, H. Ziar, O. Isabella, and M. Zeman. The Dutch PV portal 2.0: An online photovoltaic performance modeling environment for the Netherlands. *Renewable Energy*, 154:175–186, 2020.

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

L. Shi and B. Hu. Frontiers in Operations: Battery as a Service: Flexible electric vehicle battery leasing. *Manufacturing & Service Operations Management*, 26(4):1269–1285, 2024.

B. Sun, X. Tan, and D. H. Tsang. Optimal charging operation of battery swapping and charging stations with QoS guarantee. *IEEE Transactions on Smart Grid*, 9(5):4689–4701, 2017.

B. Sun, X. Sun, D. H. Tsang, and W. Whitt. Optimal battery purchasing and charging strategy at electric vehicle battery swap stations. *European Journal of Operational Research*, 279(2):524–539, 2019.

Tesla. Supercharger for business. `https://www.tesla.com/supercharger-for-business`, 2026. Accessed July 10, 2026.

K. Šepetanc and H. Pandžić. A cluster-based operation model of aggregated battery swapping stations. *IEEE Transactions on Power Systems*, 35(1):249–260, 2020.

Vandebron. Dynamisch energiecontract. `https://vandebron.nl/energie/energiecontract/dynamisch`, August 2024. Accessed August 8, 2025.

R. S. Widrick, S. G. Nurre, and M. J. Robbins. Optimal policies for the management of an electric vehicle battery swap station. *Transportation Science*, 52(1):59–79, 2018.

H. Wu, G. K. H. Pang, K. L. Choy, and H. Y. Lam. An optimization model for electric vehicle battery charging at a battery swapping station. *IEEE Transactions on Vehicular Technology*, 67(2):881–895, 2017.

O. Q. Wu and R. Kapuscinski. Curtailing intermittent generation in electrical systems. *Manufacturing & Service Operations Management*, 15(4, SI):578–595, 2013.

O. Q. Wu, R. Kapuscinski, and S. Suresh. On the distributed energy storage investment and operations. *Manufacturing & Service Operations Management*, 25(6):2277–2297, 2023.

L. Yang, X. Geng, X. Guan, and L. Tong. EV charging scheduling under demand charge: A block model predictive control approach. *IEEE Transactions on Automation Science and Engineering*, 21(2):2125–2138, 2023.

J. Yu, M. Xue, H. Shen, and Y. Liang. Dynamic control of shared resources in multitype service system with an application in battery swap stations. *Working paper*, 2024.

Y. H. Zhou, A. Scheller-Wolf, N. Secomandi, and S. Smith. Managing wind-based electricity generation in the presence of storage and transmission capacity. *Production and Operations Management*, 28(4): 970–989, 2019.

P. Zipkin. On the structure of lost-sales inventory models. *Operations Research*, 56(4):937–944, 2008.