Bachelor Computer Science & Datascience and Artificial Intelligence

A Linear Programming Approach to EV Charging Power Supply Optimization: Performance, Scalability, and Service Trade-Offs

Alejandro Mediavilla Sáez

First supervisor:
Dr. Thomas Moerland
Second supervisor:
Koen Ponse MSc

FIRST BACHELOR THESIS

Leiden Institute of Advanced Computer Science (LIACS)
redwww.liacs.leidenuniv.nl July 6, 2026

s3795055 Dr. Thomas Moerland Koen Ponse MSc
Bachelor Thesis
Data Science and Artificial Intelligence

A Linear Programming Approach to EV Charging Power Supply Optimization: Performance, Scalability, and Service Trade-Offs

Alejandro Mediavilla Sáez

July 6, 2026

Abstract

Operators of electric vehicle (EV) charging stations face a constrained revenue-maximisation problem: power must be allocated across vehicles while respecting per-port current limits, a shared grid power cap, and each vehicle's energy deadline. This paper develops a two-layer scheduling method that separates discrete admission from continuous allocation. Port admission and assignment are resolved combinatorially, after which a linear programme (LP) sets charging currents for all active ports at each five-minute timestep inside a receding-horizon model predictive control (MPC) loop. Because the discrete decisions are handled upstream, the LP contains no binary variables; deadline-driven priority is instead enforced through explicit must-serve constraints. The same solver also schedules an on-site battery energy storage system (BESS), allowing the controller to arbitrage energy over time alongside EV charging. We evaluate the controller across three station sizes (3, 6, and 12 ports), four lookahead horizons ($H = 1$, $H = 3$, $H = 6$, and $H = 12$ steps), and two pricing regimes: a fixed tariff (EUR 0.75/kWh) and a dynamic cost-plus tariff set at 1.5× the grid buy price. Three rule-based policies, MaxCharge, EqualShare, and Random, serve as baselines. Under the fixed tariff the LP earns mean net profits of EUR 208, EUR 412, and EUR 813 at 3, 6, and 12 ports, exceeding the EqualShare heuristic by 28%, 33%, and 30% respectively ($p < 0.001$, Cohen's $d \approx -1.4$). A greedy one-step horizon ($H = 1$) recovers roughly 99.9% of the profit available under a full 60-minute lookahead ($H = 12$) at a small fraction of the computational cost, and every configuration solves comfortably within the five-minute control interval. Taken together, these findings show that the LP-MPC controller substantially outperforms EqualShare and Random, while its advantage over MaxCharge at the nominal grid cap is small, being mostly attributable from profits derived from BESS scheduling.

Contents

1 Introduction 3
1.1 Operational Problem Setting . . . 3
1.2 Research Gap . . . 4
1.3 Research Questions and Thesis Structure . . . 5

2 Related Work 6
2.1 Optimization-based EV charging control . . . 6
2.2 Heuristic and rule-based charging strategies . . . 6
2.3 Dynamic pricing, storage, and profit-oriented scheduling . . . 7
2.4 Stochastic arrivals and service requirements . . . 7
2.5 Positioning of this paper . . . 8

3 Methodology 8
3.1 Conventions and Indexing . . . 9
3.2 Objective Function . . . 10
3.3 Constraints . . . 11
3.4 Rolling-Horizon Operation and Warm-Starting . . . 13
3.4.1 Must-Serve Constraints (C10) . . . 14
3.5 Experimental Design . . . 15

4 Baselines 16

5 Proof of Concept 16

6 Experiments 18
6.1 LP versus Rule-Based Baselines . . . 19
6.2 The Value of Lookahead . . . 20
6.3 Computational Scalability . . . 21
6.4 Profit versus Customer Service . . . 21
6.5 Effect of the Dynamic Tariff . . . 22
6.6 BESS Ablation Study . . . 23
6.7 Statistical Significance . . . 25

7 Conclusions and Further Research 26

8 Acknowledgments 28

A Full Grid-Cap Sensitivity Tables 29

References 32

Notation

The following symbols are used throughout this paper. All decision variables are continuous and non-negative. Port assignments, occupancy indicators, and SoC-dependent ratios are treated as fixed parameters.

Index Sets

$\mathcal{T} = \{1, \dots, T\}$	Set of time steps
$\mathcal{J} = \{1, \dots, J\}$	Set of EV charging ports
$\mathcal{J}^+ = \{1, \dots, J{+}1\}$	All ports (port $J{+}1$ is the BESS)
$\mathcal{I} = \{1, \dots, N\}$	Set of cars (EVs)

Parameters — Time & Grid

Δt	Time-step duration (h)
$P^{\max}$	Grid power cap (W)
p_t^{buy}	Electricity buy price at step t (€/kWh)
p_t^{sell}	Electricity sell price at step t (€/kWh)
L_t	Background (non-EV) load at step t (W)

Parameters — Ports & BESS

V_j	Voltage at port $j \in \mathcal{J}^+$ (V)
$I_j^{\max}$	Maximum current at EV port $j \in \mathcal{J}$ (A)
I^{high}	BESS maximum charging current (A)
I^{low}	BESS maximum discharging current (A)
$\text{SoC}_{\text{init}}^B$	Initial BESS state of charge (kWh)
$\text{SoC}_{\min}^B$	Minimum BESS state of charge (kWh)
$\text{SoC}_{\max}^B$	Maximum BESS state of charge (kWh)
r_t^{ch}	BESS SoC-dependent charging ratio at step t, $\in [0, 1]$
r_t^{dis}	BESS SoC-dependent discharging ratio at step t, $\in [0, 1]$
η	BESS one-way efficiency (round-trip $\approx \eta^2$)

Parameters — EVs

$z_{j,t}$	Occupancy indicator: 1 iff an EV occupies port j at step t
arr_i	Arrival time step of car i
dep_i	Departure time step of car i
$y_{i,j}$	Pre-computed assignment: 1 iff car i is assigned to port j
s_i^{init}	Initial state of charge of car i on arrival (kWh)
s_i^{cap}	Battery capacity of car i (kWh)
$s_i^{\min}$	Minimum state of charge of car i (kWh)
s_i^{target}	Target state of charge of car i at departure (kWh)
$P_i^{\text{car,max}}$	Maximum charging power of car i (W)
$r_{i,t}^{\text{car}}$	SoC-dependent charging ratio for car i at step t, $\in [0, 1]$

Parameters — Rolling-Horizon (MPC)

H	Planning horizon length (number of time steps)
$\mathcal{W}$	Current MPC window: $\{t_0, \ldots, \min(t_0{+}H{-}1,\, T)\}$
$\mathrm{SoC}^B_{\mathrm{now}}$	Measured BESS state of charge at current step (kWh)
s_i^{now}	Measured state of charge of car i at current step (kWh)

Decision Variables

$I^{\mathrm{ev}}_{j,t}$	Charging current at EV port j at step t (A)
I^{ch}_t	BESS charging current at step t (A)
I^{dis}_t	BESS discharging current at step t (A)
SoC^B_t	BESS state of charge at step t (kWh)
$s_{i,t}$	State of charge of car i at step t (kWh)

Objective Function Terms

ε	Urgency tiebreaker weight ($= 10^{-4}$)
p^{eff}_t	Effective price at step t, incorporating urgency bonus (€/kWh)
R^{ev}	Total EV charging revenue (€)
C^{grid}	Total grid electricity cost (€)

1 Introduction

The transition away from fossil fuels has made electric vehicles (EVs) one of the fastest-growing segments of the transport sector, and with their adoption comes a rapid expansion of the charging infrastructure that supports them [17]. As governments and international organizations commit to reducing greenhouse-gas emissions and actively promote EV uptake [6], the efficient operation of charging infrastructure becomes essential to sustaining that adoption. Furthermore, the transition to electric mobility does not only require more charging points; it also requires more intelligent use of existing charging capacity and, unlike conventional refueling, EV charging occupies the infrastructure for extended periods and converts customer service into a time-dependent power-allocation problem. These complexities mean that, as adoption increases, charging stations must therefore manage not only how many vehicles they can connect, but also how much power each connected vehicle should receive at each point in time.

This charging problem, however, can take different forms. Battery electric vehicles connect to the power system to recharge, and they can do so in very different ways. Charging at home or in a parking lot may take several hours, whereas dedicated fast-charging stations reduce that time to minutes. Achieving such short charging times requires drawing substantial power from the grid, which makes active energy management at the station level essential [12]. This distinction is especially important for public and semi-public charging stations. A single charger may be rated for high power, but a station with multiple chargers is usually constrained by a shared grid connection, transformer capacity, or contractual power limit. As a result, the sum of the individual charger capacities can exceed the power that the station is allowed to draw at any one time, which can incur a penalty. The station operator must therefore decide how to allocate a limited power budget among a continually changing landscape of connected vehicles [10].

Furthermore, the problem also has an economic dimension. Electricity cost varies over time, and a charging-station operator earns revenue by selling energy to customers. In the current market, many public charging tariffs are fixed or flat per kWh, setting unchanging prices for customers for each kWh sold, regardless of electricity costs incurred by the operator. Operators may also adopt dynamic tariffs, where the same unit of delivered energy can have different margins depending on when it is purchased from the grid [14, 1, 3]. Likewise, if the station includes an on-site battery energy storage system (BESS), the operator may additionally shift energy across time, charging the battery when electricity is cheaper and discharging it when grid energy is more expensive or when the grid cap binds [8, 13]. These factors add another layer of operational flexibility and complexity to the problem.

These characteristics make station operation a repeated decision problem under hard physical constraints. At every control step, the operator must decide how much current to send to each vehicle and whether to charge or discharge the on-site battery at a specific tariff, while respecting port limits, grid limits, vehicle energy requirements, and departure deadlines. This motivates the use of formal optimization methods rather than purely reactive or rule-based allocation rules.

1.1 Operational Problem Setting

The setting considered in this paper is a charging station with multiple EV charging ports and a shared power limit. Vehicles arrive over time, occupy a port if one is available, and leave after a specified dwell period. Each vehicle has an initial state of charge, a battery capacity, a target energy level, and a departure deadline. The controller operates at discrete five-minute intervals and determines the charging current assigned to each active port.

The scheduling problem is divided into two layers. A discrete admission and port-assignment layer determines which vehicles connect to which ports, making admission, assignment, and charging interdependent. Conditional on those assignments, the proposed linear programming (LP) controller solves the continuous allocation problem, determining how much current to deliver to each connected EV, and how to operate the BESS, during the current planning window.

The controller must respect several classes of constraints. Port-level current limits restrict the maximum current that can be delivered by each charger. Vehicle-level limits restrict charging according to battery capacity and acceptance rate. A shared grid-cap constraint limits the total power drawn by the station. The BESS is constrained by charge/discharge limits and state-of-charge bounds. Finally, service constraints encourage or enforce sufficient energy delivery before each vehicle departs.

The empirical question is therefore comparative rather than purely formulative. A mathematically structured controller may allocate power more efficiently, but it also introduces computational overhead and may behave differently from simple operational policies. This paper therefore evaluates whether a linear programming model predictive control (MPC) controller delivers meaningful gains over rule-based baselines, and whether those gains persist as the number of ports, tariff structure, and planning horizon change.

1.2 Research Gap

The literature reviewed shows that EV charging has been studied through optimization-based control, heuristic congestion management, dynamic pricing, local storage, and stochastic user modelling. However, there remains limited empirical evidence on how much practical value is gained by replacing simple rule-based allocation policies with a continuous LP-MPC controller under the same scalable simulation protocol. In particular, the trade-off between economic performance, computational tractability, and customer service quality remains underexplored.

This paper addresses that gap through four contributions:

1. It formulates a continuous receding-horizon LP controller for EV charging stations with a shared grid cap, vehicle state-of-charge requirements, deadline-aware service constraints, and on-site battery energy storage.

2. It separates discrete admission and port assignment from continuous current allocation, allowing the real-time scheduling problem to remain a linear programme without binary decision variables.

3. It benchmarks the LP-MPC controller against three operationally simple baselines, MaxCharge, EqualShare, and Random, across multiple station sizes, lookahead horizons, and tariff regimes.

4. It evaluates both economic and service-related outcomes, including net profit, solve time, customers served and rejected, and state-of-charge fulfillment at departure, thereby making the profit-service trade-off explicit.

Thus, the aim of this paper is to provide a controlled empirical evaluation of whether structured LP-based scheduling provides meaningful operational advantages over simpler rule-based charging policies, and under which operating conditions those advantages remain practically relevant.

1.3 Research Questions and Thesis Structure

Building on the research gap identified above, the guiding research question of this paper is:

> *How does a receding-horizon LP controller perform against operationally simple rule-based allocation policies, in particular equal power sharing among connected vehicles, for EV charging power-supply optimization?*

This question focuses on the practical value of structured optimization. The objective is not only to formulate an LP for the charging-station problem, but also to evaluate whether the resulting controller improves operational performance when compared with simpler policies that are easier to implement. The comparison is conducted in terms of economic performance, computational scalability, and service quality.

The main research question is divided into three sub-questions:

- **RQ1.** How do solver runtime and net economic performance scale with the number of charging ports?
- **RQ2.** How does the LP-MPC controller trade-off net profit against service quality, measured by customers served, rejected customers, and state-of-charge fulfillment at departure?
- **RQ3.** How do dynamic pricing and on-site battery energy storage affect controller performance, relative ranking, and computational tractability?

RQ1 evaluates whether the proposed controller remains computationally feasible as the station grows. This is tested by running the controller on stations with 3, 6, and 12 charging ports and recording both net profit and solver runtime. Since the controller is re-solved every five minutes, real-time feasibility requires each optimization step to solve well within the control interval.

RQ2 examines the operational trade-off between profitability and customer service. A controller may increase profit by allocating power more selectively, but this can affect how many customers are served and whether vehicles leave with sufficient charge. For this reason, the experiments report not only net profit, but also customers served, rejected arrivals, and mean state-of-charge fulfillment at departure.

RQ3 studies whether the formulation remains useful when additional operational components are introduced. Dynamic pricing changes the margin obtained from each unit of delivered energy, while on-site battery storage introduces intertemporal energy-management decisions. The experiments therefore compare controller performance under both fixed and dynamic tariff regimes and include a BESS ablation to isolate how much of the LP's advantage comes from storage scheduling rather than EV power allocation alone.

The remainder of the thesis is organized as follows. Section 3 presents the methodology and LP-MPC formulation, including the objective function, physical constraints, BESS modelling, and must-serve constraints. Section 3.5 describes the experimental design, including the station configurations, tariff regimes, random seeds, baselines, and evaluation metrics. Section 6 reports the empirical results, comparing the LP controller against the rule-based baselines in terms of profit, solve time, horizon sensitivity, and service quality. The final section discusses the implications of the results, the limitations of the modelling assumptions, and possible extensions for future work.

2 Related Work

The optimization of EV charging has been studied from several complementary perspectives, including grid-aware scheduling, model predictive control, heuristic congestion management, dynamic pricing, local storage, and stochastic customer behaviour. These strands of literature motivate the formulation used in this paper, but they also leave open a specific benchmarking question, which Section 1.2 states precisely and this paper sets out to answer.

2.1 Optimization-based EV charging control

Optimization-based charging control formulates station operation as a mathematical programme in which charging decisions are chosen to optimize an objective subject to physical and operational constraints. These constraints may include charger current limits, transformer or grid-connection limits, vehicle state-of-charge requirements, battery storage limits, and departure deadlines. Depending on the modelling choices, EV charging can be formulated as a linear programme, a convex optimization problem, or a mixed-integer programme [7].

Model predictive control is particularly relevant when charging decisions must be made online and future vehicle arrivals are uncertain. Tang and Zhang [22, 15] propose an MPC approach for EV charging scheduling that reduces computational complexity while allowing the controller to update decisions as new information becomes available. This supports the use of a receding-horizon control structure in the present thesis, where the controller repeatedly solves a finite-horizon optimization problem and executes only the first-step decision [18].

More detailed station-level formulations may use mixed-integer programming to represent discrete operating decisions, grid interaction, storage dynamics, and bidirectional energy flows. Šolic et al [24]. formulate EV charging-station power-supply optimization as a constrained optimization problem with V2X capabilities, explicitly accounting for transformer limits, grid interaction, and battery dynamics. Such formulations are expressive, but the inclusion of discrete variables can increase computational burden as the number of vehicles and time steps grows, and they are not benchmarked against the simple rule-based policies an operator might use instead. The present paper takes the complementary position: it strips the discrete decisions out of the optimisation entirely and asks how much a purely continuous LP then gains over those rule-based policies under a common protocol [23].

This paper follows a more restricted but computationally tractable direction. Instead of including admission and port assignment as binary decisions inside the optimization model, these discrete decisions are handled before the LP is invoked. Conditional on the resulting port assignments, the controller optimizes only continuous EV charging currents and BESS charge/discharge currents. This separation preserves the real-time tractability of a linear programme while retaining the main station-level constraints relevant to charging operation.

2.2 Heuristic and rule-based charging strategies

Rule-based charging strategies remain operationally attractive because they are simple to implement, require reduced computation, and are easy to explain. Equal power sharing distributes available capacity among connected vehicles and provides a natural fairness baseline under a shared power cap. MaxCharge prioritizes immediate delivery by charging vehicles as quickly as possible whenever capacity is available. These policies can perform well in lightly constrained settings, especially when the grid cap is not frequently binding, but they do not explicitly optimize over future prices, storage decisions, or deadline-aware service requirements.

Hussain et al. [9] approach charging-station utilization from a congestion-management perspective, proposing a fuzzy-inference optimization framework to reduce waiting times and improve charger utilization. Their method illustrates the appeal of heuristic and soft-computing approaches: they can adapt to operating conditions without requiring a fully specified mathematical programme. However, such methods rely on rule construction rather than explicit optimization of a global objective.

Recent comparative work also highlights the practical relevance of comparing rule-based control with predictive optimization. For example, rule-based controllers can be substantially faster and simpler than MPC, while MPC may offer advantages when performance-sensitive objectives are considered [20]. What such comparisons typically do not isolate is where the predictive controller's advantage originates, for instance whether it comes from power allocation itself or from auxiliary flexibility such as storage. Separating those sources is one of the aims of the BESS ablation in this thesis. This motivates the central empirical comparison in this paper: whether a structured LP-MPC controller produces meaningful gains over operationally simple baselines when all controllers are evaluated under the same simulated demand conditions.

2.3 Dynamic pricing, storage, and profit-oriented scheduling

A second line of work studies the economic operation of charging stations under time-varying electricity prices and local energy storage [4, 8]. In this setting, the controller is not only deciding how to divide power among vehicles, but also when to purchase electricity from the grid and whether stored energy should be used to support charging demand [13]. Dynamic tariffs make the margin on delivered energy time-dependent, while an on-site BESS can shift energy across time or relieve the grid cap during high-demand periods [1].

Luo et al. [14] study dynamic pricing and energy-management strategies for EV charging stations under uncertainty, considering the interaction between charging demand, electricity prices, storage, and operator profit. Their work is directly relevant to the economic framing of the present thesis, where controller performance is evaluated using net profit, defined as EV charging revenue minus grid electricity cost.

Mirzaei et al. [16] develop a two-step optimization framework for EV parking lots centred on profit maximization, integrating pricing structures and operational constraints to improve economic performance while respecting grid limits [11]. The present thesis adopts a similar profit-oriented perspective, but uses it in a comparative benchmark: the LP-MPC controller is evaluated against rule-based baselines under both fixed and dynamic tariff regimes.

2.4 Stochastic arrivals and service requirements

Charging-station operation is also shaped by uncertainty in customer behaviour. Vehicles arrive at different times, occupy chargers for different durations, and differ in their initial state of charge, battery capacity, and target energy requirement. These variables determine how urgent each charging session is and how much flexibility the controller has. In a real-time setting, future arrivals are uncertain, so a receding-horizon controller must repeatedly re-optimize as new vehicles arrive and active vehicles approach departure.

Alinejad et al. [2, 25] incorporate stochastic owner behaviour into charging and discharging schedules, recognizing that EV users are heterogeneous in their arrival times, deadlines, and charging requirements. This motivates the use of stochastic simulation in the present thesis, where different random seeds generate different arrival patterns, battery states, and target charge levels.

Service quality cannot be reduced to total energy delivered, because a vehicle that receives energy too late may still leave below its target state of charge. In this paper, service quality is therefore evaluated using operational metrics such as customers served, rejected arrivals, and state-of-charge fulfillment at departure, rather than through an explicit customer utility function.

2.5 Positioning of this paper

Taken together, prior work shows that EV charging can be approached through predictive optimization, mixed-integer station-level modelling, heuristic congestion management, stochastic behaviour modelling, and profit-oriented scheduling.

The contribution of this paper is therefore not the introduction of EV charging optimization itself, but a controlled empirical comparison of a continuous LP-MPC scheduler against operationally simple allocation policies. The comparison is conducted across multiple station sizes, lookahead horizons, and tariff regimes, and evaluates both economic and service-related outcomes. The empirical evaluation is performed in Chargax, a JAX-accelerated EV charging simulation environment that provides a common benchmark setting for comparing controllers under identical simulated station dynamics [21].

This positioning allows the thesis to assess not only whether optimization improves net profit, but also whether the improvement is computationally tractable and what service trade-offs arise relative to simpler control policies. Chargax was selected over alternative EV smart-charging simulators such as EV2Gym [19] for its JAX-accelerated architecture, which supports the repeated, high-frequency re-solving that a receding-horizon controller with many seeds and configurations requires.

3 Methodology

This section presents the complete linear programming (LP) formulation for the EV charging park with an on-site Battery Energy Storage System (BESS). The model is deployed as a rolling-horizon Model Predictive Control (MPC) controller: at each five-minute control step it determines charging currents for all EV ports and the BESS over a finite lookahead window so as to maximise profit subject to grid, hardware, and energy constraints, executes only the first step, and re-solves at the next step from a fresh observation. All decision variables are continuous and non-negative ($I^{\mathrm{ev}}_{j,t}, I^{\mathrm{ch}}_{t}, I^{\mathrm{dis}}_{t}, \mathrm{SoC}^{B}_{t}, s_{i,t} \in \mathbb{R}_{\geq 0}$), meaning that the LP chooses real-valued quantities greater than or equal to zero, such as charging currents and states of charge. The formulation contains no integer variables: every discrete decision, namely which car occupies which port, is resolved by a separate admission-control layer in the Chargax environment before the LP is invoked, leaving only continuous currents to optimise. This two-layer structure has the practical advantage of keeping the online controller computationally lightweight and suitable for repeated real-time re-optimization.

Its limitation is that admission, port assignment, and charging are not optimized jointly: the LP can improve current allocation and BESS scheduling for the vehicles already connected, but it cannot revise earlier discrete allocation choices that may have been affected by different charging-speed decisions.

Algorithmically, each MPC iteration follows the same sequence. At the beginning of a five-minute control step, the Chargax environment observes the current station state, updates which vehicles are present, and provides the fixed admission and port-assignment information used by the LP. The controller then computes the SoC-dependent charging ratios from the observed EV and BESS states, solves the LP over the lookahead window, and executes only the first-step charging and BESS decisions. The simulator then advances by one timestep, updating vehicle SoCs, BESS SoC, arrivals, departures, and port occupancy.

3.1 Conventions and Indexing

Four conventions, fixed once here, are relied on throughout the formulation.

Timing (end-of-step) convention. The variable $s_{i,t}$ is the state of charge of car i at the end of step t, after the energy delivered during step t has been added; SoC_t^B is defined analogously for the BESS. Energy deposited during a step is therefore reflected in that step's end-of-step SoC. This convention is load-bearing for the must-serve constraints of Section 3.4: it lets us fix the SoC entering a step to a measured value while leaving the end-of-step SoC as a decision. This makes sure that warm-starting and pacing do not conflict: warm-starting anchors the model to the observed pre-step SoC, whereas pacing imposes a minimum required increase by the end of the same step.

Sign convention. All currents are non-negative. The BESS direction is carried by two distinct variables, $I_t^{\mathrm{ch}} \geq 0$ (charging) and $I_t^{\mathrm{dis}} \geq 0$ (discharging), never by the sign of a single current. Charging always increases stored energy; discharging always decreases it. The net BESS current is $I_t^{\mathrm{ch}} - I_t^{\mathrm{dis}}$, positive when the battery is a net load on the grid.

Unit reduction. A product $V_j I_{j,t}^{\mathrm{ev}}$ $(\mathrm{V} \cdot \mathrm{A})$ is a power in watts. Multiplying by the step duration Δt (h) and dividing by 1000 converts watt-hours to kilowatt-hours, the unit of all SoC and price quantities. The recurring factor $\Delta t/1000$ performs exactly this $\mathrm{W} \cdot \mathrm{h} \to \mathrm{kWh}$ reduction, and appears in every energy term without exception.

Car–port linkage. The pre-computed assignment $y_{i,j} \in \{0,1\}$ is 1 iff car i is assigned to port j; the admission layer guarantees each docked car occupies exactly one port, so $\sum_{j \in \mathcal{J}} y_{i,j} \leq 1$ for every i. We write $j(i)$ for the unique port with $y_{i,j(i)} = 1$. The occupancy indicator follows from the assignment and the dwell window,

$$z_{j,t} \; = \; \sum_{i \in \mathcal{I}} y_{i,j} \, \mathbb{1}[\mathrm{arr}_i \leq t \leq \mathrm{dep}_i] \, , \qquad \forall \, j \in \mathcal{J}, \; t \in \mathcal{T}, \tag{1}$$

with the admission layer enforcing $z_{j,t} \in \{0,1\}$ (at most one car per port per step). All of $y_{i,j}$, $j(i)$, and $z_{j,t}$ are fixed parameters at solve time. The SoC-dependent derating ratios r_t^{ch}, r_t^{dis}, and $r_{i,t}^{\mathrm{car}}$

physically depend on SoC, which is itself a variable; to keep the model linear they are precomputed from the measured SoC at the start of the window and held fixed across it. The next re-solve refreshes them against the new observed SoC.

3.2 Objective Function

The objective maximises profit, EV charging revenue minus grid electricity cost, plus a vanishing urgency term that breaks ties in favour of earlier delivery.

Urgency tie-breaker. To stop the controller deferring charging unnecessarily ("procrastination"), a negligible bonus is added to the sell price. Let $\mathcal{W} = \{t_0, \ldots, t_{\text{last}}\}$ with $t_{\text{last}} = \min(t_0 + H - 1, T)$ be the current window. The effective sell price is

$$p_t^{\text{eff}} \;=\; p_t^{\text{sell}} + \varepsilon \, \frac{t_{\text{last}} - t + 1}{H}, \qquad \varepsilon = 10^{-4}. \tag{2}$$

The numerator $t_{\text{last}} - t + 1$ counts steps from t to the window end inclusive, so the multiplier $(t_{\text{last}} - t + 1)/H \in (0, 1]$ is largest at the window start and decays to the end. Because the multiplier never exceeds 1, the total per-kWh bonus is bounded by $\varepsilon = 10^{-4}$ /kWh regardless of H, and cannot overturn any genuine profit difference larger than that. Measuring the numerator relative to the window end t_{last} rather than the global horizon T is what keeps the bonus bounded; using $T - t + 1$ would let it grow with the horizon and risk distorting real decisions.

EV revenue. Total EV revenue sums, over every port and every step in the window, the energy delivered valued at the effective sell price p_t^{eff} from Eq. (2); the occupancy indicator $z_{j,t}$ zeroes out any port that is not currently serving a vehicle, so revenue is only accrued for energy actually sold.

The factor $\Delta t/1000$ converts the instantaneous power term $V_j I_{j,t}^{\text{ev}}$ into delivered energy in kWh. Specifically, $V_j I_{j,t}^{\text{ev}}$ is measured in watts, multiplication by Δt converts watts into watt-hours because Δt is measured in hours, and division by 1000 converts Wh into kWh.

$$R^{\text{ev}} \;=\; \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{W}} V_j \, I_{j,t}^{\text{ev}} \, z_{j,t} \, p_t^{\text{eff}} \, \frac{\Delta t}{1000}. \tag{3}$$

Grid cost and the role of efficiency. We treat the BESS one-way efficiency η as an economic (energy-accounting) quantity rather than a modifier of the physical power balance, a choice made deliberately so that the LP's incentives match the Chargax simulator it is evaluated in, which applies η the same way. Depositing a commanded charge current I_t^{ch} requires purchasing I_t^{ch}/η of grid current (losses paid at purchase); discharging a commanded I_t^{dis} offsets only $I_t^{\text{dis}}\,\eta$ of grid draw (losses incurred on delivery). The grid cost is therefore

$$C^{\text{grid}} \;=\; \sum_{t \in \mathcal{W}} \left(\sum_{j \in \mathcal{J}} V_j \, I_{j,t}^{\text{ev}} + V_{J+1} \, \frac{I_t^{\text{ch}}}{\eta} - V_{J+1} \, I_t^{\text{dis}} \, \eta \right) p_t^{\text{buy}} \, \frac{\Delta t}{1000}. \tag{4}$$

The physical grid-cap C3 (Eq. (9)) and the SoC dynamics C4 (Eq. (10)) instead use the raw commanded currents $V_{J+1}(I_t^{\text{ch}} - I_t^{\text{dis}})$, because the hardware draws and the bank stores the commanded quantities. We state plainly what this implies: the efficiency loss is realised in the money ledger but not in the energy ledger, where SoC tracks the raw current. No energy is destroyed in the SoC state; the asymmetry is an accounting penalty calibrated to the evaluation environment, not a claim about where physical losses occur. The next paragraph verifies this penalty is nonetheless sufficient to rule out artificial arbitrage, which is the property that actually matters for correctness.

Round-tripping is not a free money pump. Because η enters the cost ledger but not the SoC ledger, we must confirm that the LP cannot manufacture profit by cycling the battery. Consider charging and later discharging one unit of commanded BESS current at a constant grid price p. Under the raw-current SoC convention, the SoC returns exactly to its starting value when the charged and discharged current magnitudes are equal. Economically, however, the charge step requires purchasing p/η units from the grid, while the discharge step offsets only $p\eta$ units of grid cost. The net cost of the round trip is therefore $(p/\eta) - (p\eta) = p(\eta^{-1} - \eta) > 0$ for any $\eta \in (0, 1)$. A flat-price round trip therefore strictly loses money. More generally, BESS cycling is profitable only when the later avoided grid price is sufficiently higher than the earlier charging price to overcome the efficiency penalty. The accounting asymmetry thus disincentivises wasteful cycling: the two ledgers differ only by the paid-for loss, never in the direction of artificial arbitrage.

We emphasise that this is a fidelity choice, not a claim of physical conservatism: a physically conservative model would debit η from the stored energy as well, so our SoC trajectory is optimistic by the per-cycle loss relative to true physics. We adopt the simulator's convention so the controller's incentives match the environment it is scored in; a deployment against real hardware should move η into the SoC update (C4, Eq. (10)) and re-derive the round-trip penalty accordingly. The arbitrage-immunity result above is unaffected, since it depends only on η appearing in the cost ledger.

Optimisation problem. Combining the EV revenue of Eq. (3) and the grid cost of Eq. (4), at each control step the controller solves, over the current window $\mathcal{W}$,

$$\begin{aligned} \max_{\substack{I^{\text{ev}}_{j,t},\, I^{\text{ch}}_t,\, I^{\text{dis}}_t,\\ \text{SoC}^B_t,\, s_{i,t} \geq 0}} \quad & R^{\text{ev}} - C^{\text{grid}} \\ \text{s.t.} \quad & \text{C1–C9, Eqs. (6)–(15), and C10a–C10b, Eqs. (17)–(18).} \end{aligned} \tag{5}$$

The net profit reported in the experiments is $R^{\text{ev}} - C^{\text{grid}}$ accumulated over the executed steps; the urgency term contributes at most 10^{-4} /kWh and is negligible relative to the reported figures.

3.3 Constraints

The constraints below are stated over the current window $\mathcal{W} = \{t_0, \ldots, t_{\text{last}}\}$, which is rebuilt and re-solved at each control step.

C1 — EV port current upper bound. Each port's charging current is capped by its own hardware rating, regardless of which vehicle occupies it or its state of charge:

$$I^{\text{ev}}_{j,t} \;\leq\; I^{\max}_j, \qquad \forall\, j \in \mathcal{J},\; t \in \mathcal{W}. \tag{6}$$

C2 — BESS current bounds (SoC-dependent).

$$I_t^{\text{ch}} \le r_t^{\text{ch}} I^{\text{high}}, \qquad \forall t \in \mathcal{W}, \tag{7}$$

$$I_t^{\text{dis}} \le r_t^{\text{dis}} I^{\text{low}}, \qquad \forall t \in \mathcal{W}, \tag{8}$$

where I^{high} and I^{low} are the BESS maximum charge and discharge currents, respectively. (The names are retained for consistency with the implementation; both are positive upper bounds on the corresponding non-negative current.)

C3 — Grid power cap. The physical power drawn from the grid, using raw commanded currents (no efficiency factor, per Section 3.2), may not exceed $P^{\max}$ at any step. Background load L_t consumes headroom:

$$\sum_{j \in \mathcal{J}} V_j \, I_{j,t}^{\text{ev}} + V_{J+1}\big(I_t^{\text{ch}} - I_t^{\text{dis}}\big) + L_t \le P^{\max}, \qquad \forall t \in \mathcal{W}. \tag{9}$$

The term L_t is retained for generality. Chargax does not expose a background load, so $L_t \equiv 0$ in all experiments. We note that this makes the cap looser, not tighter: any real background load would consume headroom the model assumes is free, so the zero-load assumption is optimistic for feasibility and revenue, not conservative. It should be revisited if load data becomes available.

C4 — BESS SoC dynamics. The BESS state of charge evolves using raw commanded currents. At the first window step t_0 the SoC entering the step is fixed to the measured value $\text{SoC}^B_{\text{now}}$ (Section 3.4); thereafter it chains within the window:

$$\text{SoC}_t^B = \begin{cases} \text{SoC}^B_{\text{now}} + \big(I_t^{\text{ch}} - I_t^{\text{dis}}\big)\, V_{J+1} \dfrac{\Delta t}{1000}, & t = t_0, \\ \text{SoC}^B_{t-1} + \big(I_t^{\text{ch}} - I_t^{\text{dis}}\big)\, V_{J+1} \dfrac{\Delta t}{1000}, & t_0 < t \le t_{\text{last}}. \end{cases} \tag{10}$$

As explained in Section 3.2, η appears only in the grid-cost term; the SoC update uses raw currents, matching Chargax's internal battery model, and round-tripping remains strictly unprofitable at a flat price.

C5 — BESS SoC bounds. The battery's stored energy must remain within its allowed operating range at every step of the window, preventing over-charge or over-discharge:

$$\text{SoC}^B_{\min} \le \text{SoC}^B_t \le \text{SoC}^B_{\max}, \qquad \forall t \in \mathcal{W}. \tag{11}$$

C6 — Car SoC dynamics. Using the assigned port $j(i)$ and the end-of-step convention, with the SoC entering step t_0 fixed to the measured value s_i^{now} (Section 3.4):

$$s_{i,t} = \begin{cases} s_i^{\text{now}} + I^{\text{ev}}_{j(i),t}\, V_{j(i)} \dfrac{\Delta t}{1000}, & t = t_0 \le \text{dep}_i, \\ s_{i,t-1} + I^{\text{ev}}_{j(i),t}\, V_{j(i)} \dfrac{\Delta t}{1000}, & t_0 < t \le \text{dep}_i, \\ s_{i,t-1}, & t > \text{dep}_i. \end{cases} \tag{12}$$

EV charging is assumed loss-free (no η on the car side), consistent with C8, which only caps the same current.

C7 — Car SoC bounds. Analogously to the BESS, each car's state of charge must stay between a minimum floor and its physical battery capacity at every step:

$$s_i^{\min} \;\le\; s_{i,t} \;\le\; s_i^{\text{cap}}, \qquad \forall\, i \in \mathcal{I},\; t \in \mathcal{W}. \tag{13}$$

The bounds keep the SoC physical; deadline satisfaction is enforced separately by the must-serve constraints C10 (Section 3.4).

C8 — Car charging power limit (SoC-dependent). Beyond the port's own hardware cap (C1), the current delivered to a car is further limited by what the vehicle itself can accept, which tapers as it approaches full charge:

$$I^{\text{ev}}_{j(i),t}\, V_{j(i)} \;\le\; r^{\text{car}}_{i,t}\, P_i^{\text{car,max}}, \qquad \forall\, i \in \mathcal{I},\; t \in \mathcal{W} : t \le \text{dep}_i, \tag{14}$$

with $r^{\text{car}}_{i,t} \in [0,1]$ the precomputed acceptance ratio near full charge.

C9 — Occupancy enforcement. An unoccupied port delivers no current. Because $z_{j,t}$ is a parameter, this is a fixed equality, not a conditional (big-M) bound, which keeps the problem a pure LP and imposes no spurious upper bound on occupied ports:

$$I^{\text{ev}}_{j,t} \;=\; 0 \qquad \text{whenever } z_{j,t} = 0, \qquad \forall\, j \in \mathcal{J},\; t \in \mathcal{W}. \tag{15}$$

The two gating mechanisms are consistent by construction. From Eq. (1), $z_{j(i),t} = 1$ exactly when $\text{arr}_i \le t \le \text{dep}_i$, which is the same window over which C6 (Eq. (12)) and C8 (Eq. (14)) carry energy for car i. The two edges are enforced by distinct mechanisms. The departure edge is enforced solely by C9: for $t > \text{dep}_i$ the occupancy term gives $z_{j(i),t} = 0$ and C9 forces $I^{\text{ev}}_{j(i),t} = 0$, while C6 merely holds the SoC constant ($s_{i,t} = s_{i,t-1}$) and C8 is omitted, so neither car constraint zeroes the current itself. The arrival edge is enforced implicitly: car i enters the optimised set $\mathcal{I}$ only once docked, i.e. for $t \ge \text{arr}_i$, so no pre-arrival step admits a current variable for it. Since $z_{j,t} = 1$ holds precisely when some assigned car satisfies $\text{arr}_i \le t \le \text{dep}_i$, and at most one car occupies a port at any step, no step has a port deemed occupied by one mechanism and empty by the other; C9 binds only outside each assigned car's dwell window.

3.4 Rolling-Horizon Operation and Warm-Starting

The controller operates in a receding-horizon fashion. At each control step t_0 a window $\mathcal{W} = \{t_0, \ldots, \min(t_0 + H - 1, T)\}$ is optimised, the car set $\mathcal{I}$ contains only currently docked cars, and the solve is warm-started from the latest observation:

1. The BESS SoC entering step t_0 is fixed to the measured $\text{SoC}^B_{\text{now}}$, so the C4 recursion at $t = t_0$ reads $\text{SoC}^B_{t_0} = \text{SoC}^B_{\text{now}} + (I^{\text{ch}}_{t_0} - I^{\text{dis}}_{t_0})\, V_{J+1}\, \Delta t/1000$.
2. The car SoC entering step t_0 is fixed to the measured s_i^{now}, so $s_{i,t_0} = s_i^{\text{now}} + I^{\text{ev}}_{j(i),t_0}\, V_{j(i)}\, \Delta t/1000$.
3. The derating ratios $r^{\text{ch}}_t, r^{\text{dis}}_t, r^{\text{car}}_{i,t}$ are recomputed from the measured SoC at t_0 and held fixed across $\mathcal{W}$.

Under the end-of-step convention the warm-start fixes the SoC entering the window (the measured s_i^{now}, before step t_0 charges), whereas s_{i,t_0} is the SoC after step t_0 delivers energy. The two are distinct, which is what makes the pacing constraint below consistent rather than contradictory. Only the first step of each plan is executed; the window is then re-solved from a fresh observation. Re-solving every step rather than every H steps lets the controller track the actual car SoC, which drifts from the plan because the simulator delivers current in discrete levels, and lets it react immediately to new arrivals and departures.

3.4.1 Must-Serve Constraints (C10)

Two constraints prevent the receding-horizon solver from deferring all charging to the window end.

C10a - Pacing constraint. Let $needed_i = s_i^{\mathrm{target}} - s_i^{\mathrm{now}}$ and $K_i = \mathrm{dep}_i - t_0 + 1$ be the remaining deficit and steps until departure. The pacing floor is

$$\phi_i = \min\left\{ \frac{needed_i}{K_i},\ \frac{P^{\max}\,\Delta t}{|\mathcal{I}| \cdot 1000} \cdot 0.99,\ \frac{I_{j(i)}^{\max}\, V_{j(i)}\,\Delta t}{1000},\ needed_i \right\}, \tag{16}$$

the four terms being the per-step energy needed to finish on time, a fair grid-share budget (1% slack for numerical feasibility), the port hardware cap, and the total remaining need. The constraint requires the end-of-step SoC at t_0 to advance by at least ϕ_i:

$$s_{i,t_0} \ \geq\ s_i^{\mathrm{now}} + \phi_i, \qquad \forall\, i \in \mathcal{I} : needed_i > 0. \tag{17}$$

Since s_{i,t_0} is the post-charge SoC and s_i^{now} the pre-charge value, this forces a delivery of at least ϕ_i during step t_0 and is consistent with the warm-start.

The grid-share term $P^{\max}\,\Delta t/(|\mathcal{I}| \cdot 1000)$ budgets each car an equal slice of the cap, which guarantees the floors are jointly satisfiable with respect to C3 only when the EV ports are the sole consumers of grid power. When the BESS draws charging current or a background load L_t is present, the sum of floors can exceed available headroom; this residual joint infeasibility is handled by the bare-mode fallback of C10b (Eq. (18)).

Each term in ϕ_i is feasible for car i in isolation with respect to the remaining deficit, the average per-step energy needed to finish on time, the fair grid-share budget, and the port hardware cap. The floor does not explicitly include the SoC-dependent car acceptance limit; any residual infeasibility caused by car-side tapering, grid-cap contention, BESS charging, or background load is handled by the bare-mode fallback of C10b.

C10b - Deadline constraint. For cars departing inside the window whose target is physically achievable under both the car acceptance rate and the port limit:

$$s_{i,\,\mathrm{dep}_i} \ \geq\ s_i^{\mathrm{target}}, \qquad \forall\, i : \mathrm{dep}_i \leq t_0 + H - 1 \ \text{ and }\ s_i^{\mathrm{now}} + \frac{\min\Big(P_i^{\mathrm{car,max}},\ I_{j(i)}^{\max} V_{j(i)}\Big)\, K_i\,\Delta t}{1000} \geq s_i^{\mathrm{target}}. \tag{18}$$

The achievability test uses the binding minimum of the car acceptance rate and the port limit, so it never declares a target achievable that the port cannot deliver. It does not account for grid-cap contention among simultaneous cars; in the rare event that C10a and C10b jointly render the window infeasible, typically a tight grid cap with many nearly-full cars, the controller rebuilds and re-solves the window without C10a and C10b, so no step is ever skipped. This fallback abandons the service guarantee for that step only.

Note that at the shortest horizon $H = 1$ the window is the single step t_0, so the deadline guard $\text{dep}_i \leq t_0 + H - 1 = t_0$ activates C10b only for cars departing in the current step; all other cars are governed by the pacing floor C10a alone. This is part of why even a one-step horizon manages deadlines adequately.

3.5 Experimental Design

The experiments address the three research questions in turn. For RQ1, the station is evaluated at three sizes, 3, 6, and 12 charging ports, and for each configuration we record both net economic performance and computational runtime. The LP-MPC controller is tested with four lookahead horizons, $H = 1$, $H = 3$, $H = 6$, and $H = 12$, corresponding to planning windows of 5, 15, 30, and 60 minutes respectively. This allows us to assess whether longer planning horizons improve profit and whether solve time remains comfortably below the five-minute control interval as the problem size grows.

For RQ2, the experiments compare economic performance with service-related outcomes. A controller may increase profit by allocating power more selectively, but this can affect how many customers are served and the quality of charging. We therefore also record customers served, rejected arrivals, and mean state-of-charge fulfillment at departure. These metrics make the profit-service trade-off explicit and allow the LP-MPC controller to be compared against the rule-based baselines beyond profit alone. All demand, price, and service-outcome data are generated by the Chargax simulation environment [21], which advances the station state independently of the LP's internal planning window. Although each LP solve looks ahead over its window $\mathcal{W}$, only the first-step decision is ever executed (Section 3.4); the controller then re-observes the true, independently evolved state from Chargax and re-plans from scratch, so reported profit reflects performance under the environment's actual realized dynamics rather than under the single trajectory the controller happened to plan against.

For RQ3, the experiments evaluate the effect of dynamic pricing and on-site battery energy storage. Two tariff regimes are tested. Under the fixed tariff, EV energy is sold at a constant EUR 0.75/kWh. Under the dynamic cost-plus tariff, the EV charging price tracks the spot electricity price at $1.5\times$ the grid buy price. This dynamic tariff changes the absolute profit level but preserves a fixed proportional relationship between buy and sell prices. To isolate the contribution of the BESS, we additionally run a storage ablation at the nominal grid cap in which the battery is disabled while the EV-charging problem is otherwise unchanged. This comparison separates profit gains caused by EV power allocation from gains caused by storage scheduling.

Each controller is evaluated under the same simulated demand conditions. For each station size and tariff regime, we run 30 independent random seeds, where each seed produces a different realization of vehicle arrivals, dwell times, battery capacities, initial states of charge, and target charge levels. One episode corresponds to a full simulated day of 288 five-minute control steps. The same set of seeds is used across controllers so that performance differences can be attributed to the control policy rather than to differences in demand generation.

4 Baselines

Furthermore, the LP-MPC controller is benchmarked against three rule-based baselines: MaxCharge, EqualShare, and Random, where MaxCharge requests, for each connected vehicle, the maximum current its port and acceptance rate allow, then scales these requests down proportionally only if their sum would exceed the grid cap, EqualShare distributes available power among connected vehicles and serves as a proportional load-sharing baseline, and Random provides a low-information baseline for comparison. For the BESS-enabled experiments, the LP-MPC controller can schedule BESS charging and discharging. The rule-based baselines do not actively schedule the BESS; consequently, the BESS ablation is especially useful for identifying whether the LP's advantage comes from storage scheduling or from EV power allocation alone. We did not include a rule-based BESS baseline because doing so well requires designing a heuristic with its own lookahead-like logic, which would need separate justification and validation; instead we isolate the LP's storage-driven advantage via ablation, and flag a well-designed rule-based storage heuristic as the natural next comparison.

The main reported metrics are net profit, customers served, rejected arrivals, mean state-of-charge fulfillment at departure, and computational solve time. Net profit is computed as EV charging revenue minus grid electricity cost, accumulated over the executed control steps of each episode. Solve time is reported per control step and compared with the five-minute timestep budget. For the BESS ablation, we additionally report net profit with and without storage and decompose the LP-MaxCharge profit gap under both tariff regimes. Statistical significance of profit differences is assessed using pairwise Welch t-tests across the 30 seeds, with Cohen's d reported as the effect size.

5 Proof of Concept

Before evaluating the full model, we verified the LP's feasibility on a small, hand-built scenario designed to exercise the behaviours that matter most in practice: current allocation, port handover, deadline satisfaction, and attainment of the global optimum. The scenario uses 3 charging ports, 8 hourly time steps, and 5 car arrivals, with departure times chosen so that the optimal solution must fast-charge one of the cars. All EV ports are rated at 400 V / 32 A (12.8 kW each) under a 20 kW grid cap. Note that this proof-of-concept uses a deliberately separate, hand-specified hardware configuration chosen to exercise the constraint behaviours of interest; it is independent of the Chargax station parameters used in the main experiments (Section 6). The BESS is disabled here to keep the test purely EV-centric, since EV charging is the dominant revenue driver. To bring the scenario closer to market conditions and sharpen the LP's incentive to exploit cheap energy, the time-of-use prices fall from a morning peak (€0.20 / €0.45 per kWh buy/sell at t_1–t_2) to an off-peak trough (€0.12 / €0.22 per kWh at t_7–t_8). The scenario is solved as an oracle, where the model sees the entire horizon at once, so that any failure to exploit the price spread can be attributed to the formulation rather than to limited lookahead.

The scenario's objective mirrors the full model but retains only the EV charging current $I_{j,t} \geq 0$ at each port j and step t as a decision variable:

$$\max \sum_{j,t} V_j \, I_{j,t} \, z_{j,t} \, p_t^{\text{sell}} \, \frac{\Delta t}{1000} \; - \; \sum_t \Big(\sum_j V_j \, I_{j,t} \Big) \, p_t^{\text{buy}} \, \frac{\Delta t}{1000}, \tag{19}$$

where $z_{j,t} \in \{0, 1\}$ is the precomputed occupancy indicator, p_t^{sell} the per-kWh service tariff, and p_t^{buy} the time-of-use grid price. As in the full model, the small urgency term $\varepsilon(T - t + 1)/H$ on the sell price breaks degeneracy in favour of earlier delivery. Port assignments are computed by the admission-control layer before the LP is invoked, so the scheduling problem is a pure LP: Cars 1–3 arrive together at t_1 and fill all three ports, Car 4 arrives at t_1 but is rejected for lack of a free port, and Car 5 arrives at t_5 to take the port Car 1 has just vacated. The cases of interest are therefore the charging allocations themselves, especially for Car 1, which must hit its target before departing at t_4.

The output of the proof of concept was as follows:

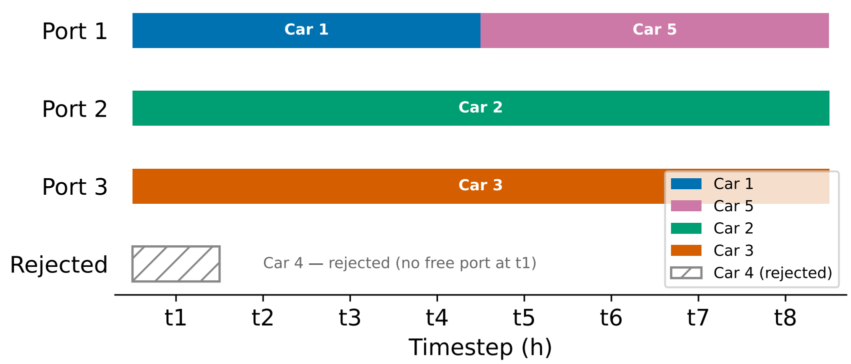

Figure 1: Vehicle-port assignment and rejection pattern for the tiny cost-case scenario. The Gantt chart shows the fixed port occupancy over an eight-hour scheduling horizon. Port 1 is assigned sequentially to Car 1 during t_1–t_4 and Car 5 during t_5–t_8, while Car 2 and Car 3 occupy Ports 2 and 3, respectively. Car 4 is rejected at t_1 because no charging port is available at its arrival time.

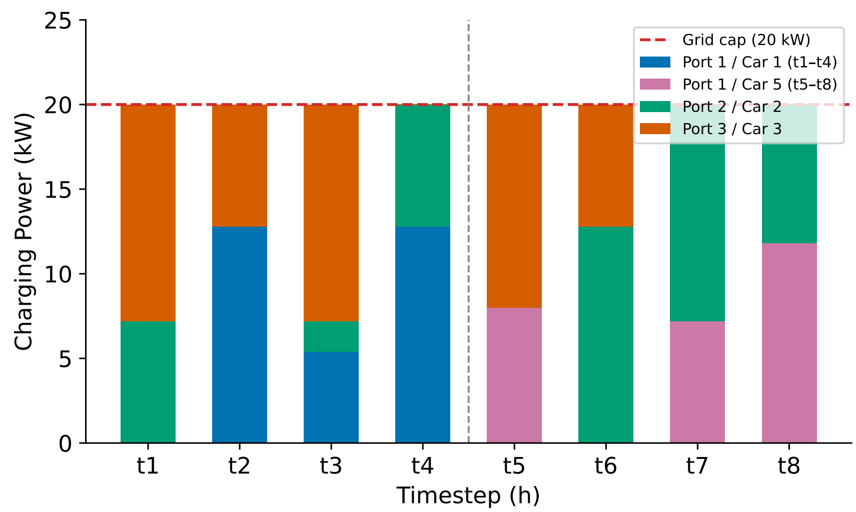

Figure 2: Optimised charging schedule for the tiny cost-case scenario. The stacked bars show the charging power allocated to each active vehicle-port assignment over eight hourly timesteps. The dashed horizontal line represents the 20 kW grid-cap constraint, which is respected throughout the horizon.

The test confirmed four expected behaviours, which can be read directly from Figures 1 and 2. First, the grid cap binds correctly: as shown in Figure 2, with three simultaneous cars whose individual maxima sum to 38.4 kW, the optimizer delivers exactly 20 kW in every slot, saturating the available capacity without ever exceeding it and confirming the tightness of C3, which is essential for replicating an industrial setting.

Second, the LP honours deadline-driven priority. As shown in Figure 2, Car 1 (departure t_4, target 20 kWh, initial SoC 5 kWh) must receive 15 kWh in four time steps and reaches its target by t_3, its SoC trajectory rising to 23.2 kWh, even though energy is cheaper later in the horizon. The LP thus sacrifices some cost-efficiency to meet the hard deadline. This is behaviour that, without an explicit must-serve constraint in the full model, would not survive the transition from oracle to rolling-horizon operation.

Third, the scenario demonstrates port reuse across sessions. As shown in Figure 1, Port 1 passes from Car 1 (occupying it t_1–t_4) to Car 5 (t_5–t_8) with no charging gap, and Car 5 climbs from a low initial SoC of 3 kWh to its 15 kWh target by t_7. This confirms that the occupancy model correctly encodes non-contiguous schedules within a single physical port index.

Fourth, the LP exploits slack capacity for non-urgent cars. As shown in Figure 2, Cars 2 and 3 carry targets well below their physical battery capacities (25 / 60 kWh and 22 / 60 kWh respectively). Once Car 1's deadline is satisfied, the LP continues charging Cars 2 and 3, making use of capacity that would otherwise sit idle, and front-loads that charging into the higher-margin morning slots wherever the grid cap and deadline constraints allow, converting the buy-sell spread into additional margin. Together these four results establish that the formulation behaves as intended, justifying the full-scale experiments that follow.

A limitation highlighted by Figure 1 is that port occupancy is treated as exogenous once a vehicle is admitted. In this proof-of-concept, and in the full model, a car is assigned to a port if one is available at arrival, otherwise it is rejected, and its departure time is assumed to be known at arrival. This keeps the online problem continuous and computationally tractable, but it abstracts away from charge-dependent dwell behaviour. In practice, some drivers may leave earlier if their vehicle reaches a satisfactory state of charge sooner, meaning that faster charging could itself create additional port availability. Modelling such endogenous departures would require coupling charging decisions with future occupancy and admission decisions, most likely through a mixed-integer or behavioural extension of the current two-layer framework.

6 Experiments

The main interpretation is deliberately comparative, and the choice of reference policy matters. Of the three baselines, MaxCharge is the one a real operator would plausibly deploy: it is simple, fast, and near-optimal whenever the grid cap is slack, so it is the demanding comparison against which the LP must justify its overhead. EqualShare and Random are weaker references included to bracket the performance range rather than to represent realistic operating practice. With that ordering in mind: the LP scheduler achieves the highest mean net profit in the tested configurations, but the source and size of this advantage differ across baselines. Relative to EqualShare and Random, the LP provides a clear and stable improvement. Relative to MaxCharge, however, the profit gap is small at the nominal grid cap, and the BESS ablation in Section 6.6 shows that this small edge is driven primarily by storage scheduling rather than by EV power allocation alone.

6.1 LP versus Rule-Based Baselines

Table 1 compares the best LP horizon against the three baselines under the fixed and dynamic tariff regimes. The LP scheduler obtains the highest mean profit in every configuration tested. The gain over EqualShare is substantial and stable, ranging from roughly 28% to 36% across station sizes and tariff regimes. This confirms that equal sharing leaves significant economic value unused when the controller can instead allocate current according to the station state, deadlines, grid cap, and storage opportunity. We report this EqualShare gain because it is large and stable, but caution against reading it as the headline result: it measures distance from a deliberately weak baseline. The operationally meaningful comparison is with MaxCharge, examined next.

Table 1: Best LP performance compared with rule-based baselines under fixed and dynamic tariff regimes. Net profit values are reported as mean episode profit over 30 random seeds.

Tariff	Ports	Best LP	EqualShare	MaxCharge	Random	Gain vs EqualShare	Gain vs MaxCharge
Fixed	3	208.09	163.09	205.56	180.99	27.6%	1.2%
Fixed	6	411.65	309.54	409.48	376.06	33.0%	0.5%
Fixed	12	813.40	627.65	812.09	760.92	29.6%	0.2%
Dynamic	3	59.30	45.79	57.76	50.95	29.5%	2.7%
Dynamic	6	116.52	85.90	115.06	105.61	35.6%	1.3%
Dynamic	12	229.42	174.77	228.18	213.89	31.3%	0.5%

Note. Fixed tariff = EUR 0.75/kWh. Dynamic tariff = 1.5× spot buy price.

The comparison with MaxCharge requires a more careful interpretation. At the nominal cap, MaxCharge is already close to optimal because the grid constraint binds only intermittently. On steps where the cap is slack, charging every connected vehicle at full rate is essentially the same action the LP would choose. As a result, the LP's advantage over MaxCharge is small in absolute terms, even though it is consistent across seeds. The BESS ablation later in this section shows that this edge mostly disappears when the battery is disabled, meaning that the LP's nominal-cap advantage over MaxCharge is mainly a storage-scheduling effect.

The LP-MaxCharge gap also decreases as the number of ports grows. Under the fixed tariff, the LP advantage over MaxCharge falls from EUR 2.53 at 3 ports to EUR 2.17 at 6 ports and EUR 1.31 at 12 ports. This should not be interpreted as a general loss of LP value, since the gain over EqualShare remains large. Rather, it suggests that MaxCharge is already close to optimal for EV charging at the nominal grid cap, and that the LP's remaining advantage is mainly linked to BESS scheduling. Since the same storage unit is used while EV throughput increases with station size, the economic influence of the BESS becomes diluted in relative terms. This interpretation is consistent with the BESS ablation in Section 6.6, where disabling storage reduces the LP-MaxCharge profit gap to approximately zero.

The full breakdown by horizon under the fixed tariff appears in Table 2. The ordering of controllers is stable across station sizes: the LP variants achieve the highest mean profit, MaxCharge follows closely, Random is lower, and EqualShare performs worst. The margins over EqualShare scale with the number of ports, while the margins over MaxCharge remain very small.

Table 2: Net profit under the fixed tariff regime, with EV charging revenue priced at 0.75 EUR/kWh. Values are reported as mean ± standard deviation over 30 seeds.

Controller	3 Ports	6 Ports	12 Ports
EqualShare	163.09 ± 38.94	309.54 ± 53.95	627.65 ± 93.38
MaxCharge	205.56 ± 18.77	409.48 ± 37.70	812.09 ± 76.14
Random	180.99 ± 16.03	376.06 ± 35.47	760.92 ± 72.13
LP, $H = 1$	207.67 ± 18.68	411.14 ± 37.93	813.18 ± 75.83
LP, $H = 3$	207.85 ± 18.77	411.25 ± 37.82	813.37 ± 76.02
LP, $H = 6$	208.09 ± 18.87	411.45 ± 37.82	813.40 ± 75.98
LP, $H = 12$	208.00 ± 18.77	411.65 ± 37.86	813.36 ± 76.10

The dynamic-tariff results in Table 3 show the same controller ordering at a lower absolute profit level. Because the dynamic tariff is defined as a fixed multiple of the spot buy price, it reduces the absolute margin but does not create a fundamentally different ranking of controllers.

Table 3: Net profit under the dynamic cost-plus tariff regime, where the EV charging price is set to 1.5 times the spot electricity price. Values are reported as mean ± standard deviation over 30 seeds.

Controller	3 Ports	6 Ports	12 Ports
EqualShare	45.79 ± 13.66	85.90 ± 17.15	174.77 ± 31.69
MaxCharge	57.76 ± 9.38	115.06 ± 18.67	228.18 ± 36.86
Random	50.95 ± 8.49	105.61 ± 17.11	213.89 ± 35.02
LP, $H = 1$	59.21 ± 9.55	116.36 ± 18.73	229.34 ± 37.05
LP, $H = 3$	59.29 ± 9.56	116.41 ± 18.78	229.28 ± 36.90
LP, $H = 6$	59.30 ± 9.54	116.43 ± 18.82	229.42 ± 36.99
LP, $H = 12$	59.30 ± 9.54	116.52 ± 18.78	229.41 ± 37.00

6.2 The Value of Lookahead

A natural expectation is that longer planning horizons yield better schedules. Table 4 shows that this effect exists but is practically very small. Extending the horizon from one step ($H = 1$) to twelve steps ($H = 12$), equivalent to a full hour of lookahead, improves average profit by only EUR 0.34 under the fixed tariff and EUR 0.11 under the dynamic tariff. In both tariff regimes, the greedy one-step LP captures approximately 99.9% of the profit of the longest horizon.

Table 4: Effect of LP lookahead horizon on average net profit. Results are averaged across the 3-port, 6-port, and 12-port configurations. Longer horizons provide only marginal profit improvements over the greedy one-step LP.

Horizon	Fixed Profit	Fixed vs $H = 1$	Dynamic Profit	Dynamic vs $H = 1$
$H = 1$	477.33	0.00	134.97	0.00
$H = 3$	477.49	0.16	134.99	0.02
$H = 6$	477.65	0.32	135.05	0.08
$H = 12$	477.67	0.34	135.08	0.11

The interpretation is that the tested environment offers limited value to multi-step planning at the nominal grid cap. Future arrivals are stochastic and unknown in advance, and the must-serve constraints already prevent the one-step controller from indefinitely postponing deadline-relevant charging. The BESS ablation in Section 6.6 further suggests that the limited value of lookahead is structural: without a strong storage or price-spread opportunity, most profitable decisions are effectively myopic.

6.3 Computational Scalability

Table 5 reports mean solve time per control step under the fixed tariff. All configurations are far below the 300,000 ms budget implied by a five-minute control interval. Even the largest tested case, $H = 12$ with 12 ports, solves in about 45 ms on average, leaving more than 99.98% of the timestep budget unused. The LP is therefore comfortably real-time at every station size and horizon tested.

Table 5: Mean LP computation time per control step under the fixed tariff regime. All solve times are far below the 5-minute simulation timestep budget of 300,000 ms. The $H = 3$, 3-port entry marked † is inflated by a single degenerate seed that reached the two-second solver guard; the median for that cell is 24.6 ms, in line with the neighbouring configurations.

Controller	3 Ports	6 Ports	12 Ports
LP, $H = 1$	14.69 ms	9.50 ms	11.95 ms
LP, $H = 3$	284.66 ms†	13.97 ms	18.19 ms
LP, $H = 6$	17.32 ms	20.31 ms	26.52 ms
LP, $H = 12$	25.09 ms	30.45 ms	45.09 ms

The one outlier, the 284.66 ms mean for $H = 3$ at 3 ports, is an artifact of a single outlier seed in which the simplex method hit a degenerate instance and ran up against the two-second time limit. The median for that cell is 24.6 ms, and every other $H = 3$ configuration solves in 14–18 ms, so the outlier reflects a rare numerical edge case rather than a genuine scaling effect.

6.4 Profit versus Customer Service

Profit is not the only operational concern: a station also cares how many drivers it serves and how fully it charges them. Table 6 reports customers served and mean state-of-charge fulfillment. MaxCharge reveals its main operational advantage here. Because it charges at full rate whenever capacity is available, it finishes sessions sooner and frees ports for new arrivals. As a result, MaxCharge serves noticeably more customers than the LP at 3 and 6 ports.

Table 6: Customer service metrics under the fixed tariff regime. The served and fulfilment figures are essentially identical under the dynamic tariff, because pricing alters neither port occupancy nor vehicle arrival dynamics.

Controller	3 Ports		6 Ports		12 Ports	
	Served	SoC Fulfill.	Served	SoC Fulfill.	Served	SoC Fulfill.
EqualShare	7.2	0.911	13.3	0.879	25.6	0.897
MaxCharge	10.4	0.947	20.7	0.898	39.2	0.924
Random	8.7	0.954	18.6	0.903	34.5	0.915
LP, $H = 1$	8.1	0.906	14.5	0.872	25.1	0.896
LP, $H = 3$	8.9	0.893	14.5	0.919	26.8	0.927
LP, $H = 6$	9.2	0.898	14.1	0.926	28.9	0.922
LP, $H = 12$	8.2	0.953	15.5	0.913	28.6	0.927

The LP's must-serve pacing, by contrast, spreads charging over time to satisfy deadlines and may keep a car connected longer than strictly necessary when its deadline is distant. State-of-charge fulfillment remains high across controllers, so the main service trade-off is one of throughput rather than charge quality, where the LP earns more per served customer, while MaxCharge serves more customers by finishing sessions faster.

This trade-off also helps explain the near-optimality of MaxCharge in the profit tables. At the nominal grid cap, the cap binds only intermittently. On the many steps where the cap is slack, full-rate charging is already close to the profit-maximizing action. The LP's economic edge appears mainly when controllable flexibility matters. Specifically when power must be rationed, when deadlines interact with scarce capacity, or when storage can be scheduled.

6.5 Effect of the Dynamic Tariff

Moving from the fixed tariff to the dynamic cost-plus tariff reduces absolute profit substantially for every controller. This is expected because the dynamic sell price tracks the grid buy price at a fixed multiplier, compressing the absolute euro margin relative to the fixed EUR 0.75/kWh tariff. Table 7 shows that the average profit reduction is almost uniform across controllers, lying between 71.7% and 72.1%. Importantly, the controller ranking is preserved: the LP remains best, MaxCharge remains close behind, Random follows, and EqualShare remains the weakest baseline.

The dynamic cost-plus tariff should be interpreted as a margin-scaling regime, not as a fully realistic dynamic-pricing mechanism. Because the dynamic tariff is a fixed cost-plus multiple of the buy price, it changes the profit scale but does not create a strong independent price-arbitrage (BESS charging) signal for EV charging. The BESS ablation in Section 6.6 confirms this interpretation: the LP's advantage over MaxCharge remains small and disappears when storage is removed.

This large uniform reduction is a direct, mechanical consequence of how the dynamic tariff is constructed, and it should not be read as evidence that dynamic pricing is economically unattractive in practice. Vehicle arrivals, dwell times, and target SoC are drawn from the same 30 seeds under

Table 7: Average net profit reduction when moving from the fixed 0.75 EUR/kWh tariff to the dynamic 1.5× spot tariff. Values are averaged across all port configurations.

Controller	Fixed Tariff	Dynamic Tariff	Reduction
EqualShare	366.76	102.15	-72.1%
MaxCharge	475.71	133.67	-71.9%
Random	439.32	123.48	-71.9%
LP, $H = 1$	477.33	134.97	-71.7%
LP, $H = 3$	477.49	134.99	-71.7%
LP, $H = 6$	477.65	135.05	-71.7%
LP, $H = 12$	477.67	135.08	-71.7%

both tariff regimes and are independent of the price signal: Chargax generates demand once, and only the sell price applied to that fixed demand changes between regimes. In particular, the model contains no demand-side price elasticity, so it does not represent extra arrivals being drawn in by cheaper off-peak charging, deferred sessions in response to high prices, or customers who would forgo charging altogether above some price threshold [3, 5]. Because the dynamic sell price is pinned to 1.5× the (generally lower off-peak, higher on-peak) grid buy price, the tariff is, on average, cheaper for the customer and therefore lower-margin for the operator than the flat EUR 0.75/kWh tariff, and with arrivals fixed there is no compensating volume effect to offset that. The result is best interpreted narrowly: a cost-plus dynamic tariff without demand response compresses realized margin relative to a flat tariff calibrated near the top of the price distribution. It says nothing about whether a real dynamic-pricing deployment, which would shift arrival volume and timing toward cheaper hours, would out- or under-perform a flat tariff.

6.6 BESS Ablation Study

The simulations include an on-site BESS. To isolate its contribution, every configuration was rerun with the battery disabled, leaving the EV-charging problem otherwise unchanged. Table 8 reports net profit with and without the BESS for all controllers under both tariffs. The comparison exposes a sharp asymmetry that the aggregate profit figures alone conceal.

Table 8: Net profit (EUR) with and without the on-site BESS, at the nominal grid-cap setting ($s = 1.00$), mean over 30 seeds. Δ is the incremental profit attributable to the BESS for that controller. The LP row uses the best horizon ($H = 12$).

Tariff	Controller	BESS on	BESS off	Δ
Fixed tariff (0.75 EUR/kWh)				
	EqualShare	627.65	627.65	+0.00
	MaxCharge	812.09	812.09	+0.00
	LP, $H = 12$	813.36	811.93	+1.43
Dynamic tariff (1.5× spot)				
	EqualShare	174.77	174.77	+0.00
	MaxCharge	228.18	228.18	+0.00
	LP, $H = 12$	229.41	228.16	+1.25

Two results are noteworthy. First, the BESS is used only by the LP, evidenced by both rule-based baselines earning identical profit with and without it because neither EqualShare nor MaxCharge schedules battery charge or discharge. Under these controllers, the battery is effectively inactive. Second, the LP's incremental gain from the BESS is modest at the nominal cap, with gains of EUR 1.43 under the fixed tariff and EUR 1.25 under the dynamic tariff at 12 ports. This is consistent with the grid cap binding only intermittently in the baseline setting.

The more consequential finding concerns where the LP's advantage over MaxCharge comes from. Table 9 isolates the LP-MaxCharge profit gap with and without the battery. With the BESS enabled, the LP holds a small but consistent edge. Once the BESS is disabled, that edge disappears under both tariffs: the LP becomes essentially profit-equivalent to MaxCharge.

Table 9: LP ($H = 12$) net-profit advantage over MaxCharge (EUR), with and without the BESS, at the nominal grid-cap setting ($s = 1.00$), by port count. A negative value means the LP trails MaxCharge.

	Fixed tariff		Dynamic tariff	
Ports	BESS on	BESS off	BESS on	BESS off
3	+2.44	−0.04	+1.54	−0.01
6	+2.17	−0.10	+1.46	−0.01
12	+1.27	−0.16	+1.23	−0.02

This ablation changes the interpretation of the headline LP-MaxCharge comparison. At the nominal cap, the LP's advantage over MaxCharge is not primarily caused by superior EV power allocation. Once the BESS is disabled, the LP and MaxCharge become essentially profit-equivalent. The optimizer's advantage therefore comes from controllable flexibility that MaxCharge does not exploit, especially storage scheduling.

The result also clarifies the practical operating regimes in which LP-MPC is most valuable. For a station with a loosely binding grid cap and no usable storage, a greedy full-rate policy can be close to optimal. The LP becomes more valuable when the controller has additional flexibility to exploit, such as on-site storage, tighter grid-cap conditions, or tariff structures that create meaningful price differences over time.

This decomposition also has a scope limitation worth stating explicitly. Neither rule-based baseline is given any storage policy: EqualShare and MaxCharge simply leave the BESS idle, so the entire battery-arbitrage margin accrues to the LP by construction. The reported LP-MaxCharge gap should therefore be read as an upper bound on the LP's advantage over a battery-aware rule-based operator, not as evidence that rule-based storage control is intrinsically incapable of capturing similar value. A simple heuristic, for instance, charging the BESS during the cheapest known price steps in a lookahead window and discharging during the most expensive, might recover much of this arbitrage margin without solving an LP at all. Testing such a heuristic is a natural extension: if it closes most of the gap, the LP's remaining advantage over a well-designed rule-based operator is even smaller than Table 9 suggests; if it does not, this would itself indicate that the LP's storage schedule already coordinates BESS timing with fluctuating EV demand and grid-cap pressure in a way that is nontrivial to replicate with a fixed rule.

6.7 Statistical Significance

Table 10 confirms that the main profit differences are statistically robust. The LP's advantage over EqualShare is large, and its advantage over Random is also substantial. The LP-MaxCharge difference is statistically significant in the BESS-enabled setting because it is highly consistent across seeds, but it should not be interpreted as a broad advantage of LP allocation over full-rate charging. The BESS ablation in Section 6.6 shows that this difference is primarily attributable to storage scheduling at the nominal cap. Without the BESS, the LP-MaxCharge profit gap collapses to approximately zero.

Table 10: Aggregate statistical significance of net-profit differences across all port configurations. Pairwise Welch t-tests are computed over 30 seeds and three station sizes.

Tariff	Comparison	Mean Difference	p-value	Cohen's d
Fixed	EqualShare vs LP, $H = 12$	-110.91	< 0.0001	-1.401
Fixed	MaxCharge vs LP, $H = 12$	-1.96	< 0.0001	-2.363
Fixed	LP, $H = 1$ vs LP, $H = 12$	-0.34	< 0.0001	-0.454
Fixed	LP, $H = 12$ vs Random	38.35	< 0.0001	3.199
Dynamic	EqualShare vs LP, $H = 12$	-32.92	< 0.0001	-1.432
Dynamic	MaxCharge vs LP, $H = 12$	-1.41	< 0.0001	-5.319
Dynamic	LP, $H = 1$ vs LP, $H = 12$	-0.10	< 0.0001	-0.544
Dynamic	LP, $H = 12$ vs Random	11.59	< 0.0001	3.204

The horizon differences are also statistically significant but practically negligible. The difference between $H = 1$ and $H = 12$ is measured in cents rather than operationally meaningful profit. The statistical result therefore reflects low variance across seeds, not a meaningful economic advantage of longer lookahead.

7 Conclusions and Further Research

This paper aimed to answer whether a structured LP controller can outperform simple rule-based policies for EV charging power allocation, and at what computational cost as the station scales. Across two tariff regimes, three station sizes, and thirty random seeds, the results are nuanced: the LP is reliably the most profitable controller and is comfortably real-time, but the magnitude and source of its advantage depend heavily on what flexibility the operating regime actually exposes.

The LP achieves the highest mean net profit in every configuration tested, beating the EqualShare heuristic by 28%-36%. This margin is large, stable across station sizes, and statistically robust ($p < 0.001$, Cohen's $d \approx -1.4$). Against MaxCharge, however, the results differ. The LP's edge is small at the nominal grid cap and, as the BESS ablation shows, almost entirely attributable to storage scheduling: with the battery disabled, LP and MaxCharge become profit-equivalent to within a few cents. The defensible claim is therefore that the LP's value lies precisely in the flexibility a greedy policy cannot exploit, namely storage arbitrage and the rationing of scarce power, and that this value is small whenever those levers are inactive.

On RQ1, the LP scales well on both axes. Profit advantage over EqualShare is preserved from 3 to 12 ports, and solve time grows sub-linearly: even a sixty-minute lookahead at twelve ports completes in roughly 45 ms, under 0.02% of the five-minute control interval. Nothing in the observed scaling suggests the controller would lose tractability at the next station size up. Real-time deployment is thus not a concern at any size tested.

On RQ2, the operative trade-off is throughput, not charge quality. State-of-charge fulfillment is uniformly high across all controllers, so no policy systematically sends customers away undercharged. The difference is in customers served, where MaxCharge finishes sessions quickly and frees ports, serving more drivers, while the LP paces charging to deadlines and earns more per session. This is a genuine operational lever rather than a defect, an operator prioritizing revenue per session and an operator prioritizing raw throughput would rationally choose different points on this frontier.

This service interpretation is also bounded by the behavioural assumptions of the model. Departure times are treated as fixed and known at arrival, and port occupancy is therefore exogenous once a vehicle has been admitted. In practice, dwell time may be state-dependent: some drivers may leave earlier once they reach a satisfactory SoC, while others may remain parked even after charging is complete. Customer satisfaction may also depend on more than SoC fulfillment, including waiting time, rejection probability, energy obtained, charging duration, fairness across users, and the price paid. Future work could therefore replace the current operational service metrics with a richer customer-utility model, allowing the controller to trade-off operator profit against a more realistic measure of customer welfare.

On RQ3, both extensions integrate without changing the model's character: it remains a pure linear program with no integer variables. The dynamic cost-plus tariff lowers absolute profit by about 72% but leaves the controller ranking untouched, because a fixed multiple of the buy price creates no independent arbitrage signal for EV charging. The BESS terms likewise enter the objective without nonlinearity or integrality. The ablation is the key result here: it localizes essentially all of the LP's nominal-cap advantage over MaxCharge to storage scheduling, which clarifies when the LP is worth deploying rather than merely confirming that it outperforms other baselines.

Expanding on the results, multi-step lookahead is not operationally significant in this environment: greedy one-step planning captures about 99.9% of the profit of a full-hour horizon, and full-rate MaxCharge captures roughly 99.6% of the LP's. We derive from this result that the nominal grid cap binds only intermittently, so allocation is dominated by current-state decisions rather than temporal planning, and the LP's advantage is concentrated on the minority of steps where power must be rationed.

The limited value of lookahead also suggests an important direction for future work. In this study, the controller re-optimizes frequently but does not optimize over a probabilistic forecast of future demand. Since future arrivals, dwell times, initial SoCs, and target charge levels are stochastic, optimizing deeply into a single assumed future may have limited value if the realized future differs from that trajectory. A natural extension would therefore be stochastic or scenario-based MPC, where the controller optimizes expected performance over several plausible future demand scenarios rather than over one deterministic lookahead path. This could make the policy more robust to forecast error, especially in settings with stronger congestion, tighter grid caps, larger storage capacity, or more volatile prices.

Furthermore, the conclusions are bounded by the study's modelling choices, several of which plausibly understate the LP's potential. The most important limitation is that customer behaviour is treated as exogenous. Vehicles arrive according to the simulator, occupy a port if one is available, and leave at a fixed departure time known at arrival. This makes the optimization problem tractable, but it removes an important feedback loop: in practice, departure may depend on the charging outcome itself. Some drivers may leave as soon as they reach a satisfactory state of charge, while others may remain connected even after charging is complete. A data-dependent departure model could therefore estimate the probability that a driver leaves as a function of current SoC, obtained charge, elapsed charging time, remaining dwell time, price paid, and station congestion. This would make port availability endogenous to charging decisions and would allow the controller to value faster charging not only as energy revenue, but also as a way to release capacity for future customers.

A second extension concerns arrival demand. In the current setup, arrivals are stochastic but not price-sensitive: the tariff affects the profit per kWh, but it does not affect whether customers choose to use the station. A more realistic model could make arrivals cost-sensitive by incorporating price elasticity. Under such a model, higher prices may increase margin per kWh but reduce demand, while lower prices may attract more users but worsen congestion or reduce profitability. This would change the operator's problem from pure power allocation to joint pricing and scheduling. The LP could then be embedded inside a higher-level demand model that chooses prices, or price multipliers, based on expected congestion, grid cost, and customer willingness to pay.

A third direction is to enrich the state information available to the controller. The experiments use simulated arrivals, battery states, targets, and electricity prices, but real charging stations may also have access to external predictors. Weather, day of week, holidays, local traffic, nearby events, and historical station utilization could all influence arrival rates and dwell times. Weather can also influence grid prices and charging demand, for example through heating or cooling load and through changes in driving patterns. Including these variables would not necessarily require changing the LP itself; instead, they could feed a forecasting layer that produces better scenarios for the MPC controller.

These extensions also clarify the role of LP relative to alternative methods such as machine learning. The strength of the LP formulation is that it is interpretable, auditable, fast, and able to enforce hard physical constraints such as grid caps, port limits, battery bounds, and service deadlines. This makes it attractive for real-time infrastructure control. However, LP requires the relevant quantities to be specified or forecasted in advance, and it becomes less natural when behaviour is endogenous, nonlinear, or only partially observed. Machine-learning methods could therefore play a complementary role by forecasting arrivals, dwell times, price response, departure probabilities, and future grid prices. Reinforcement learning could in principle learn a full pricing and charging policy directly, but it would require extensive training data or high-fidelity simulation, and it may be harder to guarantee safety, feasibility, and interpretability. A promising next step is therefore not to replace the LP, but to combine both approaches: use machine learning to predict uncertain behavioural and environmental variables, and use the LP-MPC layer to make the final constraint-respecting charging and storage decisions.

Lastly, the results point to a clear ordering of next steps, each targeting a regime where the LP's advantage should widen. First, tighten the grid cap further. The strain sweep in Appendix A already shows the LP–MaxCharge profit gap growing monotonically as the cap tightens: under the fixed tariff with the BESS enabled, the LP ($H = 12$) edge at 12 ports rises from €1.31 at the nominal cap ($s = 1.00$) to €2.62 at $s = 0.25$, with the same direction at 3 and 6 ports and under the dynamic tariff. This confirms the predicted trend, but with an important qualification: the BESS-disabled columns (Table 12) show LP and MaxCharge remaining profit-equivalent at every strain level, so the widening edge reflects storage arbitrage becoming more valuable under scarcity rather than any emerging advantage in raw power allocation. A regime in which allocation alone separates the two controllers was not reached even at $s = 0.25$. Probing tighter caps still, ideally with the BESS disabled to isolate allocation, is therefore the natural way to locate the boundary at which power rationing, rather than storage, begins to drive the LP's advantage. Second, introduce arrival forecasts or stochastic-programming planning, which is the most direct test of whether longer horizons recover value once anticipatory decisions matter. Third, raise physical fidelity by modelling the constant-voltage taper and moving η into the SoC update, then re-deriving the round-trip penalty. Beyond these, allowing V2G discharge and multiple storage units would expand the flexibility the LP is uniquely positioned to exploit. Most importantly, this controller was built to anchor exactly such comparisons: it provides the transparent, optimal-within-its-model reference against which learning-based controllers, in particular the reinforcement-learning agents this work was designed to benchmark [26], can now be measured.

8 Acknowledgments

I would like to thank my supervisor, Dr. Thomas Moerland, and second supervisor, Koen Ponse MSc, for their guidance and feedback throughout this thesis. I am also grateful to Shell Information Technology International B.V for their support during this work.

A Full Grid-Cap Sensitivity Tables

This appendix reports the full numerical results for the grid-cap strain sweep. The strain parameter $s \in \{0.25, 0.50, 0.75, 1.00\}$ scales the station grid cap relative to the nominal setting used in the main benchmark. Unless stated otherwise, all values are computed over 30 independent random seeds. The main results chapter reports only the summary tables needed to interpret the effect of grid-cap tightness on LP advantage, BESS value, profit scaling, service quality, and solve time.

Table 11: Net profit (€) — fixed tariff (0.75 €/kWh), BESS enabled, all controllers. Mean ± standard deviation over 30 seeds. **Bold** marks the highest mean per strain level.

		EqualShare		**MaxCharge**		**Random**		**LP,** $H = 1$		**LP,** $H = 3$		**LP,** $H = 6$		**LP,** $H = 12$	
Ports	**Strain**	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
3	0.25	45.45	8.31	51.51	4.69	51.35	4.68	54.75	4.69	54.78	4.68	54.85	4.69	**54.92**	4.68
	0.50	90.87	12.29	103.02	9.37	101.35	9.20	105.88	9.37	105.94	9.37	106.11	9.37	**106.34**	9.35
	0.75	125.27	23.17	154.47	14.06	146.31	12.95	156.72	14.28	156.95	13.93	**157.09**	14.13	157.05	14.07
	1.00	163.09	38.94	205.56	18.77	180.99	16.03	207.67	18.68	207.85	18.77	**208.09**	18.87	208.00	18.77
6	0.25	92.18	9.68	103.02	9.37	102.68	9.37	105.92	9.38	105.94	9.37	106.11	9.37	**106.34**	9.36
	0.50	168.21	28.64	205.60	18.77	204.21	18.79	208.08	18.81	208.11	18.81	208.20	18.86	**208.19**	18.81
	0.75	242.26	38.21	307.79	28.21	300.49	27.97	309.82	28.27	309.89	28.27	309.98	28.29	**309.99**	28.30
	1.00	309.54	53.95	409.48	37.70	376.06	35.47	411.14	37.93	411.25	37.82	411.45	37.82	**411.65**	37.86
12	0.25	185.95	17.73	205.60	18.77	204.58	18.83	208.08	18.81	208.12	18.81	208.17	18.80	**208.22**	18.82
	0.50	353.96	36.32	409.48	37.70	405.75	38.15	411.15	37.64	411.36	37.82	411.56	37.83	**411.65**	37.84
	0.75	485.20	70.65	611.65	56.89	601.17	56.68	613.00	56.95	613.13	56.90	612.97	56.65	**613.17**	56.93
	1.00	627.65	93.38	812.09	76.14	760.92	72.13	813.18	75.83	813.37	76.02	**813.40**	75.98	813.36	76.10

Table 12: Net profit (€) — fixed tariff (0.75 €/kWh), BESS disabled, all controllers. Mean ± standard deviation over 30 seeds. **Bold** marks the highest mean per strain level.

		Equal_share		**Max_charge**		**Random**		**LP,** $H = 1$		**LP,** $H = 3$		**LP,** $H = 6$		**LP,** $H = 12$	
Ports	**Strain**	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
3	0.25	45.45	8.31	**51.51**	4.69	51.35	4.68	51.50	4.69	**51.51**	4.69	**51.51**	4.69	**51.51**	4.69
	0.50	90.87	12.29	**103.02**	9.37	101.35	9.20	102.91	9.37	**103.02**	9.37	**103.02**	9.37	**103.02**	9.37
	0.75	125.27	23.17	**154.47**	14.06	146.31	12.95	154.30	14.04	154.45	14.07	154.43	14.06	154.45	14.07
	1.00	163.09	38.94	**205.56**	18.77	180.99	16.03	205.34	18.79	205.53	18.76	**205.56**	18.74	205.52	18.77
6	0.25	92.18	9.68	**103.02**	9.37	102.68	9.37	**103.02**	9.37	**103.02**	9.37	**103.02**	9.37	**103.02**	9.37
	0.50	168.21	28.64	**205.60**	18.77	204.21	18.79	205.58	18.74	**205.60**	18.77	**205.60**	18.77	**205.60**	18.77
	0.75	242.26	38.21	**307.79**	28.21	300.49	27.97	**307.79**	28.21	**307.79**	28.21	**307.79**	28.21	**307.79**	28.21
	1.00	309.54	53.95	**409.48**	37.70	376.06	35.47	**409.48**	37.70	409.23	37.62	409.35	37.60	409.38	37.61
12	0.25	185.95	17.73	**205.60**	18.77	204.58	18.83	**205.60**	18.77	**205.60**	18.77	**205.60**	18.77	**205.60**	18.77
	0.50	353.96	36.32	**409.48**	37.70	405.75	38.15	**409.48**	37.70	**409.48**	37.70	**409.48**	37.70	**409.48**	37.70
	0.75	485.20	70.65	**611.65**	56.89	601.17	56.68	**611.65**	56.89	611.62	56.88	611.62	56.88	**611.65**	56.89
	1.00	627.65	93.38	**812.09**	76.14	760.92	72.13	811.78	75.79	811.95	76.12	811.98	76.13	811.93	76.12

Table 13: Net profit (€) — dynamic tariff (1.5× spot), BESS enabled, all controllers. Mean ± standard deviation over 30 seeds. **Bold** marks the highest mean per strain level.

		EqualShare		**MaxCharge**		**Random**		**LP, $H = 1$**		**LP, $H = 3$**		**LP, $H = 6$**		**LP, $H = 12$**	
Ports	**Strain**	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
3	0.25	12.70	2.79	14.47	2.35	14.43	2.34	16.24	2.54	16.25	2.55	16.27	2.55	**16.29**	2.55
	0.50	25.67	5.88	28.95	4.70	28.48	4.66	30.61	4.88	30.62	4.88	30.67	4.89	**30.73**	4.90
	0.75	35.28	8.44	43.40	7.05	41.14	6.82	44.87	7.14	44.96	7.26	**44.99**	7.22	44.97	7.23
	1.00	45.79	13.66	57.76	9.38	50.95	8.49	59.21	9.55	59.29	9.56	**59.30**	9.54	**59.30**	9.54
6	0.25	25.91	4.93	28.95	4.70	28.85	4.69	30.61	4.88	30.62	4.88	30.67	4.89	**30.73**	4.90
	0.50	47.23	10.62	57.77	9.38	57.38	9.31	59.32	9.54	59.33	9.54	59.34	9.55	**59.35**	9.54
	0.75	67.38	13.03	86.49	14.04	84.42	13.69	87.91	14.18	87.93	14.18	**87.96**	14.18	**87.96**	14.18
	1.00	85.90	17.15	115.06	18.67	105.61	17.11	116.36	18.73	116.41	18.78	116.43	18.82	**116.52**	18.78
12	0.25	52.28	9.47	57.77	9.38	57.49	9.34	59.32	9.54	59.33	9.54	59.34	9.55	**59.35**	9.54
	0.50	99.37	19.68	115.06	18.67	114.01	18.41	116.38	18.83	116.44	18.78	116.51	18.78	**116.52**	18.79
	0.75	134.44	23.89	171.86	27.82	168.96	27.43	173.09	27.91	173.09	27.89	173.13	27.94	**173.14**	27.93
	1.00	174.77	31.69	228.18	36.86	213.89	35.02	229.34	37.05	229.28	36.90	**229.42**	36.99	229.41	37.00

Table 14: Net profit (€) — dynamic tariff (1.5× spot), BESS disabled, all controllers. Mean ± standard deviation over 30 seeds. **Bold** marks the highest mean per strain level.

		EqualShare		**MaxCharge**		**Random**		**LP, $H = 1$**		**LP, $H = 3$**		**LP, $H = 6$**		**LP, $H = 12$**	
Ports	**Strain**	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
3	0.25	12.70	2.79	**14.47**	2.35	14.43	2.34	**14.47**	2.35	**14.47**	2.35	**14.47**	2.35	**14.47**	2.35
	0.50	25.67	5.88	**28.95**	4.70	28.48	4.66	28.92	4.70	**28.95**	4.70	**28.95**	4.70	**28.95**	4.70
	0.75	35.28	8.44	**43.40**	7.05	41.14	6.82	43.36	7.05	43.39	7.05	43.39	7.05	43.39	7.04
	1.00	45.79	13.66	57.76	9.38	50.95	8.49	57.70	9.35	57.76	9.38	**57.77**	9.39	57.75	9.37
6	0.25	25.91	4.93	**28.95**	4.70	28.85	4.69	**28.95**	4.70	**28.95**	4.70	**28.95**	4.70	**28.95**	4.70
	0.50	47.23	10.62	**57.77**	9.38	57.38	9.31	**57.77**	9.39	**57.77**	9.38	**57.77**	9.38	**57.77**	9.38
	0.75	67.38	13.03	**86.49**	14.04	84.42	13.70	**86.49**	14.04	**86.49**	14.04	**86.49**	14.04	**86.49**	14.04
	1.00	85.90	17.15	**115.06**	18.67	105.61	17.11	**115.06**	18.67	115.03	18.69	115.05	18.67	115.05	18.67
12	0.25	52.28	9.47	**57.77**	9.38	57.49	9.34	**57.77**	9.38	**57.77**	9.38	**57.77**	9.38	**57.77**	9.38
	0.50	99.37	19.68	**115.06**	18.67	114.01	18.41	**115.06**	18.67	**115.06**	18.67	**115.06**	18.67	**115.06**	18.67
	0.75	134.44	23.89	**171.86**	27.82	168.96	27.43	**171.86**	27.82	**171.86**	27.82	171.85	27.84	**171.86**	27.82
	1.00	174.77	31.69	**228.18**	36.86	213.89	35.02	228.10	36.96	228.06	36.89	228.07	36.87	228.16	36.86

References

[1] An aggregator-based dynamic pricing mechanism and optimal scheduling scheme for electric vehicle charging. *Frontiers in Energy Research*, 2022.

[2] Mahyar Alinejad, Mohammad Javad Amiri, Mehdi Shafie-khah, and Jo ao P. S. Catalão. An optimal management for charging and discharging of electric vehicles in an intelligent parking lot considering vehicle owner's random behaviors. *Journal of Energy Storage*, 35:102245, 2021.

[3] Louise Bernard, Andy Hackett, Robert D. Metcalfe, Luca Panzone, and Andrew Schein. The impact of dynamic prices on electric vehicle public charging demand: Evidence from a nationwide natural field experiment. NBER Working Paper 34600, National Bureau of Economic Research, 2025.

[4] Shuo Chang et al. Coordinate scheduling of electric vehicles in charging stations supported by microgrids. *Electric Power Systems Research*, 199:107418, 2021.

[5] Demand Side Analytics. Price elasticity of demand for electric vehicle charging, 2017.

[6] European Environment Agency. New registrations of electric vehicles in europe, 2022.

[7] Lingwen Gan, Ufuk Topcu, and Steven H. Low. Optimal decentralized protocol for electric vehicle charging. *IEEE Transactions on Power Systems*, 28(2):940–951, 2013.

[8] Alberto Grimaldi, Francesco Demetrio Minuto, Jacob Brouwer, and Andrea Lanzini. Profitability of energy arbitrage net profit for grid-scale battery energy storage considering dynamic efficiency and degradation using a linear, mixed-integer linear, and mixed-integer non-linear optimization approach. *Journal of Energy Storage*, 95:112380, 2024.

[9] Shahid Hussain, Reyazur Rashid Irshad, Fabiano Pallonetto, Qasim Jan, Saurabh Shukla, Subhasis Thakur, John G. Breslin, Yun-Su Kim, Muhammad Ahmad Rathore, and Hesham El-Sayed. Enhancing the efficiency of electric vehicles charging stations based on novel fuzzy integer linear programming. *IEEE Transactions on Intelligent Transportation Systems*, 24(9):9150–9162, 2023.

[10] International Electrotechnical Commission. *IEC 61851-1:2017: Electric vehicle conductive charging system – Part 1: General requirements*. International Electrotechnical Commission, Geneva, Switzerland, 2017. International Standard IEC 61851-1:2017.

[11] Mohammad Javad et al. A two-step approach to optimal management of electric vehicle parking lots. *Sustainable Energy Technologies and Assessments*, 46:101258, 2021.

[12] M. C. Kisacikoglu, B. Ozpineci, and L. M. Tolbert. Ev/phev bidirectional charger assessment for v2g reactive power operation. *IEEE Transactions on Power Electronics*, 28(12):5717–5726, 2013.

[13] Dheepak Krishnamurthy, Canan Uckun, Zhi Zhou, Praveen R. Thimmapuram, and Audun Botterud. Optimal storage arbitrage under net metering using linear programming. In *IEEE Conference Publication*, 2019.

[14] Chao Luo, Yih-Fang Huang, and Vijay Gupta. Dynamic pricing and energy management strategy for ev charging stations under uncertainties. In *Proceedings of the International Conference on Vehicle Technology and Intelligent Transport Systems – VEHITS*, pages 49–59. INSTICC, SciTePress, 2016.

[15] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[16] Mahdi Mirzaei, Mehdi Shafie-khah, and Jo ao P. S. Catalão. Optimal management of electric vehicle parking lots considering vehicle owners' random behaviors. *Journal of Energy Storage*, 35:102245, 2021.

[17] R. B. Navesi, H. Tahami, M. H. Amini, F. Ghasemi, and A. Soroudi. Reliable operation of reconfigurable smart distribution network with renewable sources, energy storage, electric vehicles and demand-side management. *Electric Power Systems Research*, 248:112301, 2025.

[18] Daniel Ochoa et al. A systematic review of model predictive control for robust and efficient energy management in electric vehicle integration and v2g applications. *Modelling*, 6(1):20, 2025.

[19] Stavros Orfanoudakis, Cesar Diaz-Londono, Yunus Emre Yilmaz, Peter Palensky, and Pedro P. Vergara. Ev2gym: A flexible v2g simulator for ev smart charging research and benchmarking. *IEEE Transactions on Intelligent Transportation Systems*, 26(2):2410–2421, 2025.

[20] Zoltan Mark Pinter, Gunnar Rohde, and Mattia Marinelli. Comparative analysis of rule-based and model predictive control algorithms in reconfigurable battery systems for ev fast-charging stations. *Journal of Energy Storage*, 116:116008, 2025.

[21] Koen Ponse, Jan Felix Kleuker, Aske Plaat, and Thomas Moerland. Chargax: A jax accelerated ev charging simulator. *arXiv preprint arXiv:2507.01522*, 2025.

[22] Wanrong Tang and Ying Jun Zhang. A model predictive control approach for low-complexity electric vehicle charging scheduling: Optimality and scalability. *IEEE Transactions on Power Systems*, 32(2):1050–1063, 2017.

[23] Mehdi Veisi et al. Optimal charging station placement of electric vehicles in the smart distribution network based on the mixed integer linear programming. *International Journal of Electrical Power & Energy Systems*, 168:110675, 2025.

[24] Antonio Josip Šolić, Damir Jakus, Josip Vasilj, and Danijel Jolevski. Electric vehicle charging station power supply optimization with v2x capabilities based on mixed-integer linear programming. *Sustainability*, 15(22):16073, 2023.

[25] Zongfei Wang, Patrick Jochem, and Wolf Fichtner. A scenario-based stochastic optimization model for charging scheduling of electric vehicles under uncertainties of vehicle availability and charging demand. *Journal of Cleaner Production*, 254:119886, 2020.

[26] Zhonghao Zhao, Carman K. M. Lee, Xiaoyuan Yan, and Haonan Wang. Reinforcement learning for electric vehicle charging scheduling: A systematic review. *Transportation Research Part E: Logistics and Transportation Review*, 190:103698, 2024.