

Article

Model Predictive Control for Multi-Objective Vehicle-to-Grid Dispatch: Jointly Optimizing Peak Shaving, Renewable Utilization, Battery Degradation, and Economic Revenue in Smart EV Infrastructures

Muhammad Abdullah Bin Arif ^{1,2,*} , Shahid Iqbal ²  and Sanchari Deb ¹¹ School of Engineering, Newcastle University, Newcastle upon Tyne NE1 7RU, UK² Department of Engineering, University of Gujrat, Gujrat 50700, Pakistan; si@uog.edu.pk

* Correspondence: m.a.b.arif2@newcastle.ac.uk or 23016122-001@uog.edu.pk

Abstract

Vehicle-to-grid (V2G) technology lets a parked electric vehicle push power back to the network, so a fleet of cars can act as distributed storage. Many studies report large benefits from this, such as peak shaving and better use of renewable energy, but most describe their simulation only in words, name no test network, and compare a single charging behavior against a do-nothing case. It is then hard to separate what V2G delivers from what the control strategy delivers. This paper builds and compares dispatch strategies on one fully specified system: a price-following rule with no knowledge of the network, a stronger randomized time-of-use baseline, and a receding-horizon model predictive control (MPC) strategy that re-plans every hour with the distribution network's per-line thermal limits and voltage bounds embedded directly in the optimization, not merely checked afterward. All run on the IEEE 33-bus feeder at low (10%), medium (30%), and high (50%) EV shares, with a full AC power flow solved every hour. The central result is a critical-penetration effect. At a 50% share, the naive price rule makes the system peak 23.7% worse than having no V2G at all, because the whole fleet reacts to one price signal at once, while the MPC cuts the peak by 18 to 28% in every case. Embedding the network limits removes the worst-line-loading side-effect a system-peak-only objective creates, bringing high-share worst-line loading down from 86.3% to 81.8% while preserving the full peak reduction. A workplace daytime-charging scenario shows renewable use becoming a real, separating metric (up to 2.8 MWh/day of EV demand met directly by rooftop solar, against zero for overnight charging), a quadratic wear cost smooths the profit-cycling frontier that a linear cost makes step-shaped, the controller is robust to forecast error up to 20%, and its solve time is set by network size rather than fleet size. Results are given as they came out of the model.



Academic Editor: Valery Vodovozov

Received: 29 July 2026

Revised: 25 August 2026

Accepted: 10 September 2026

Published: 11 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: vehicle-to-grid; V2G; bidirectional charging; network-constrained model predictive control; peak shaving; battery degradation; distribution power flow; IEEE 33-bus feeder; critical penetration; workplace charging; marginal emissions

1. Introduction

Electric vehicles are being adopted quickly, and the way they charge puts new stress on distribution networks. Bidirectional chargers change what a parked car is. The battery can absorb power when there is a surplus and return it when the grid is short, so the vehicle becomes a resource rather than only a load [1]. This two-way exchange is vehicle-to-grid,

or V2G. It has been put forward for peak shaving, frequency support, and better use of renewable generation [2,3]. The idea is not limited to battery electric vehicles: fuel-cell vehicles are increasingly treated as mobile power plants for vehicle-to-everything services, and recent predictive energy-management work for fuel-cell buses that fuses speed and passenger forecasts shows how far the mobile-source side of V2X has advanced [4,5]. This paper stays with plug-in battery EVs, but the coordination problem it studies is common to both.

A recent scenario-based assessment [6] looked at the technical, economic, and environmental impact of V2G at low (10%), medium (30%), and high (50%) adoption. It reported peak reductions of up to 38%, yearly owner profits of USD 120 to 350, and CO₂ savings near 770 kg per vehicle per year. That study described its simulation only in prose. It named no test network, no dispatch algorithm, and no optimization method, and it tested a single unspecified charging behavior per scenario. Its own discussion asked for real-time schemes that weigh grid needs, renewable availability, battery status, and market conditions together, and it warned that battery wear grows with heavy cycling. This paper takes up that exact gap.

Here, a receding-horizon MPC dispatch strategy is designed and tested for a fleet of bidirectional EVs. The controller solves a linear program every hour over a rolling 24 h window and commits only the current hour's decision. It is compared against a price-following rule that stands for the uncoordinated charging described as a risk in the literature [2,7]. Both strategies run on the IEEE 33-bus radial distribution feeder [8], a standard benchmark, and the network response for every hour comes from a full AC power flow solved with pandapower [9]. The contribution is a reproducible pipeline built along five lines that earlier scenario studies leave open. A jointly optimized MPC formulation is compared against an uncoordinated baseline rather than a do-nothing case. Evaluation runs through solved power flows (minimum voltage, line losses, line loading) instead of assumed grid figures. CO₂ is accounted for with time-varying marginal emission factors [10], which can actually separate the two strategies. A throughput-based battery wear cost [11] sits directly inside the optimization objective instead of an unspecified penalty. A sensitivity analysis and a profit-versus-cycling trade-off curve are obtained by re-solving the optimization each time. Beyond that core comparison, and in direct response to the questions a coordinated V2G study of this kind raises, the formulation embeds the feeder's per-line thermal limits and voltage bounds inside the MPC rather than checking them after dispatch, so the controller cannot fix the system peak by quietly overloading one line. The same framework is then used to test a workplace daytime-charging case where charging overlaps solar hours, a quadratic battery-wear cost against the linear one, robustness to forecast error, a stronger randomized time-of-use baseline, and the solver's scaling with fleet size. The recent literature has moved toward exactly these tools, MPC, distributed and stochastic optimization, and reinforcement learning for EV and V2G scheduling [3,5], but rarely joins them to a solved-power-flow, network-constrained comparison of coordinated against uncoordinated dispatch on a named benchmark; that junction is the contribution here.

All strategies are tested on a synthetic but physically sensible model of demand, solar generation, and the EV fleet, placed on the IEEE 33-bus feeder (Section 3). Synthetic data were used because no public dataset gives field-measured feeder, fleet, and tariff data at the detail this study needs, and the limits of that choice are set out openly in Section 6. The claim is deliberately narrow. Coordinated dispatch does not win on every metric in every scenario, and the naive strategy can actively harm the grid at high EV shares. What the paper offers is a transparent, reproducible comparison of coordinated against

uncoordinated V2G dispatch on a named benchmark network that others can rebuild and extend.

2. Related Work

2.1. Vehicle-to-Grid Concepts and Economic Potential

The basic V2G idea, using parked and plugged-in EV batteries as grid resources, was formalized by Kempton and Tomić [1], who estimated the capacity and possible income of aggregated fleets. Later work explored peak shaving, ancillary services, and renewable integration [12,13], and more recent surveys widen the frame from vehicle-to-grid to vehicle-to-everything, cataloguing the separate value streams a bidirectional car can earn and the regulatory barriers between them [14]. The value is not only technical. A systematic review of the social side finds that trust, perceived battery wear, and contract clarity shape whether owners take part at all [15]. Techno-economic studies keep meeting the same tension: profit per vehicle depends strongly on the tariff and on how hard the battery is cycled to earn it [3]. This paper models that tension directly by placing cost and battery wear inside one objective.

2.2. Uncoordinated Charging and Distribution-Network Impacts

Uncoordinated, price-reactive charging can create new demand peaks and stress the network, for a simple reason: many independent vehicles see the same public price signal and react the same way at the same time [16,17]. Sortomme and El-Sharkawi [2] and Lopes et al. [7] describe this risk for one-way and two-way charging. Muratori [18] showed with real household data that uncoordinated charging can build a sharp new evening peak on top of the existing one. The present paper reproduces this rebound-peak effect computationally for bidirectional V2G on a named benchmark feeder (Section 4.1).

2.3. Optimization-Based and Model Predictive Control Approaches

A separate line of work treats charging and V2G dispatch as an optimization problem rather than a reactive rule, from day-ahead scheduling [19] to receding-horizon MPC, where the problem is re-solved as forecasts update [20]. Broad reviews of EV-grid integration map this space, from the technology and optimization options for two-way charging [21] to the power-interaction modes and scheduling methods that connect vehicles to the grid [22], the fleet-management services an aggregator can run [23], and the smart-charging strategies proposed for distribution networks [24,25]. MPC suits V2G because it can weigh several goals at once, cost, battery wear, and network constraints, and revise its plan as uncertainty resolves. The dispatch controller here uses exactly these properties (Section 3.6). Scenario-based assessments such as [6] do not implement or compare any such formulation.

2.4. Distribution Test Systems and Power-Flow Evaluation

The IEEE 33-bus radial feeder [8] is one of the most widely used benchmark networks in distribution studies. A standard network means other researchers can check and compare results. This paper loads the feeder through pandapower [9] and solves a genuine AC power flow with its Newton–Raphson solver for every simulated hour.

2.5. Marginal Emissions and Battery Degradation Accounting

Two further methods shape the evaluation. The marginal-emissions literature [10] argues that the CO₂ effect of shifting load should be valued against the plants that respond at the margin, usually fast-ramping peakers at peak and cleaner baseload off-peak; a flat average factor hides this. This paper uses the time-varying marginal approach directly (Section 3.7). The degradation literature prices battery wear either as a cost per MWh of

energy cycled [11] or, more precisely, through rainflow cycle-counting models that make wear a convex, increasing function of cycle depth [26]. The leveled throughput cost goes straight into the MPC objective (Section 3.6), which makes the profit-versus-cycling analysis in Section 4.10 possible, and the convex rainflow form motivates the quadratic wear variant tested in Section 4.15.

2.6. Research Gap

The literature gives V2G economics, evidence of network stress from uncoordinated charging, optimization-based dispatch, benchmark networks, and proper emissions and wear accounting, but mostly as separate pieces. They are rarely combined into one reproducible head-to-head test of coordinated against uncoordinated dispatch. The assessment of Baloch et al. [6] sits in the first group and named no algorithm or network. This paper joins the pieces and reports the outcomes honestly, including the case where the naive strategy harms the grid and the case where the smart strategy is not better on every metric. Table 1 places the present work against a representative set of studies along the features that matter for the comparison: whether a named test network is used, whether the dispatch method is stated, whether coordinated and uncoordinated behavior are actually compared, whether a solved power flow backs the grid claims, whether battery wear enters the objective, whether emissions are valued at the margin, and whether the evaluation is grounded in real charging data.

Table 1. Where this study sits relative to representative V2G and EV-grid work. A check mark (✓) means the feature is present, a circle (○) means partial or implicit, and a dash (–) means absent.

Study	Named Network	Dispatch Method	Coord. vs. Uncoord	Solved Power Flow	Wear in Objective	Marginal CO ₂
Kempton & Tomić 2005 [1]	–	○	–	–	–	–
Clement-Nyngs et al. 2010 [16]	✓	○	○	✓	–	–
Sortomme & El-Sharkawi 2012 [27]	○	✓	–	○	○	–
Vagropoulos & Bakirtzis 2013 [19]	–	✓	–	–	–	–
García-Villalobos et al. 2014 [25]	○	✓	○	○	–	–
Wang et al. 2016 [28]	–	○	–	–	✓	–
Uddin et al. 2017 [29]	–	✓	–	–	✓	–
Muratori 2018 [18]	✓	–	○	✓	–	–
Baloch et al. 2025 [6]	–	–	–	–	○	–
This work	✓	✓	✓	✓	✓	✓

3. Materials and Methods

3.1. Research Design

The study is a controlled comparison. Two dispatch strategies, the naive price rule and the proposed MPC, run on an identical network, fleet, demand profile, solar profile, tariff, and availability schedule, at three penetration levels (10%, 30%, 50%). Because everything else is held equal, any difference in outcome comes from the dispatch strategy alone. The code is written in Python 3.11, with pandapower 2.14 for the power flow and cvxpy 1.5 (HiGHS solver) for the optimization. All inputs and results are retained so every number can be regenerated. Figure 1 traces the whole pipeline, from the fixed inputs and the real-data grounding, through the two dispatch branches and the hour-by-hour receding-horizon loop of the MPC, into the solved power flow and the metric, sensitivity, and Pareto stages.




Figure 1. Methodology pipeline. Fixed inputs and the real Helix charging data feed a controlled comparison of the naive rule against the receding-horizon MPC across three penetration levels. The MPC loop re-solves each hour on a noisy forecast and commits only the current decision; both branches feed a solved AC power flow, then the metric, sensitivity, and Pareto stages.

3.2. Distribution Test Feeder

The IEEE 33-bus radial feeder [8] is the network model, loaded through pandapower's built-in case33bw. As a correctness check, the feeder's base-case losses were compared with published values for this benchmark and they match, which gives confidence that the network is set up correctly. The shipped case carries no thermal line ratings, so each line's rating is set to 1.5 times its base-case current. This is a disclosed assumption, and it defines the denominator of the reported line-loading percentages.

3.3. EV Fleet, Demand, and Solar Model

The fleet is bidirectional EVs with 60 kWh batteries and 7 kW Level-2 chargers, typical mid-size values. Vehicles are spread over the feeder's 32 load buses in proportion to each bus's share of the base load. Fleet sizes are 100, 300, and 500 vehicles for the low, medium,

and high scenarios, assuming 1000 households on the feeder. Each vehicle is plugged in at home from 18:00 to 07:00 the next day, matching a normal commuter pattern, and it uses 35% of its battery for daily driving, deducted at the 07:00 departure. Hourly demand follows a double-peak daily shape scaled so the feeder peak equals the benchmark's rated load. Rooftop solar follows a bell-shaped daytime curve with total capacity equal to 25% of the peak load, spread over buses by load share. These profiles and fleet numbers are stated simulation assumptions, not field measurements, and every value can be regenerated from the accompanying code.

3.4. Real-World Grounding: The Newcastle Helix Charging Dataset

The synthetic profiles above are kept honest by checking them against real charging records from the Newcastle Helix site, an urban innovation district next to Newcastle University that runs a live smart-grid and open-data testbed [30,31]. The dataset covers 41,213 valid charging sessions logged at the Helix chargers by the Urban Sciences Building (site 50112) between 18 March 2021 and 22 July 2026, across 1925 days and six rapid and fast chargers. The observed fields, the session start and stop times, the delivered energy, the connector type, and the anonymized charger identifiers, are taken straight from the site records. A session delivers 24.81 kWh on average (median 21.59 kWh, 90th percentile 50.2 kWh), lasts 47.6 min on average, and draws 33.85 kW of mean power, and the whole record moves 1022.4 MWh. The connector mix is 75.2% CCS Combo, 16.1% CHAdeMO, and 8.7% Type-2 AC. Table 2 collects these figures and Figure 2 shows the hourly pattern.

Table 2. Observed characteristics of the Newcastle Helix charging dataset (real records; 41,213 sessions, 2021–2026). Only measured fields are reported here.

Quantity	Value
Valid sessions after cleaning	41,213
Days covered	1925 (18 March 2021 to 22 July 2026)
Chargers/site	6/one location (ID 50112)
Mean energy per session	24.81 kWh
Median energy per session	21.59 kWh
90th-percentile energy	50.2 kWh
Mean session duration	47.6 min
Mean charging power	33.85 kW
Total delivered energy	1022.4 MWh
Connector mix (CCS/CHAdeMO/Type-2)	75.2%/16.1%/8.7%
Share of sessions 09:00–16:00	52.3%
Share of sessions 00:00–06:00	4.0%




Figure 2. Real charging behavior at the Newcastle Helix site (observed fields, 41,213 sessions, 2021–2026): (a) session arrivals by hour, (b) delivered energy by hour, (c) connector mix. The demand is daytime-dominant, the pattern of a public rapid-charging hub, which is why the paper treats residential overnight availability and daytime workplace availability as separate cases.

The real records ground the study in two ways, and expose one honest mismatch. They confirm the energy scale used in the fleet model: the mean 24.81 kWh delivered

per session is the same order as the roughly 21 kWh a 60 kWh pack gives up over the 35% daily driving draw assumed here (about 15% apart), so the per-vehicle energy the dispatch moves is realistic rather than invented. The mismatch is in timing. The Helix demand is daytime-dominant, with 52.3% of sessions starting between 09:00 and 16:00 and only 4.0% overnight, because it is a public rapid-charging hub rather than home charging. The residential model in this paper deliberately assumes overnight plug-in, so the two describe different segments of the same system. That contrast is not a flaw to paper over; it is exactly why this paper studies the daytime, solar-overlapping case explicitly as a workplace-charging scenario in Section 4.13, for which the real data supply a concrete demand shape. The derived fields in the dataset (state of charge, user preference, tariff tier) are modeling assumptions rather than measurements, so only the observed quantities are used for the claims here.

3.5. Naive (Uncoordinated) Dispatch

The naive baseline stands for the uncoordinated, price-reactive behavior flagged as a risk in the literature [2,7]. In any hour, if the price is at or above its 75th percentile and the battery allows it, the fleet at each bus discharges at full available power. If the price is at or below its 25th percentile, or local solar exceeds local demand, and the battery allows it, the fleet charges at full available power. Otherwise it does nothing. The rule knows nothing about the network, nothing about battery wear, and nothing about the other buses, and every bus applies it independently.

3.6. Proposed MPC (Coordinated) Dispatch

The proposed strategy is a receding-horizon linear program, re-solved at every hour over a rolling 24 h window. Let t index the hours in the window and i the load buses. The decision variables are the charging power $P_{t,i}^{\text{ch}} \geq 0$, the discharging power $P_{t,i}^{\text{dis}} \geq 0$, the state of charge $S_{t,i}$, the locally used solar $R_{t,i} \geq 0$, and a single peak auxiliary $\rho \geq 0$. The objective minimizes

$$\sum_t \left[\pi_t \sum_i \frac{P_{t,i}^{\text{ch}}}{\eta_c} - \pi_t \gamma \eta_d \sum_i P_{t,i}^{\text{dis}} + c_{\text{deg}} \sum_i (P_{t,i}^{\text{ch}} + P_{t,i}^{\text{dis}}) - c_{\text{ren}} \sum_i R_{t,i} \right] + c_{\text{pk}} \rho, \quad (1)$$

where π_t is the time-of-use price, $\gamma = 0.90$ is the aggregator margin on V2G sales, $c_{\text{deg}} = \text{USD}25$ per MWh cycled [11], $c_{\text{ren}} = \text{USD}10$ per MWh of local solar used, and $c_{\text{pk}} = \text{USD}200$ per MW of worst-hour system import. The state of charge evolves as

$$S_{t+1,i} = S_{t,i} + \eta_c P_{t,i}^{\text{ch}} - \frac{P_{t,i}^{\text{dis}}}{\eta_d} - D_{t,i}, \quad (2)$$

with charge and discharge efficiencies $\eta_c = \eta_d = 0.95$ and $D_{t,i}$ the driving draw deducted at departure. Each vehicle starts the horizon at a reference state of charge of 60% of usable capacity, $S_{0,i} = 0.60 \bar{S}_i$, and the state of charge is bounded between 15% and 95% of pack capacity at every hour, including the last. Charge and discharge powers are capped by the plugged-in charger rating and the availability flag, the used solar cannot exceed local generation or local demand plus charging, and the peak auxiliary satisfies $\rho \geq \sum_i (L_{t,i} + P_{t,i}^{\text{ch}} - P_{t,i}^{\text{dis}} - R_{t,i})$ for every t , with $L_{t,i}$ the base demand. The representative-day dispatch is not forced to end at the initial 60%, so the state of charge left at the end of the day is treated explicitly in the economic accounting (Section 4.4): the daily profit is also reported on a cycle-neutral basis that charges each vehicle back to its 60% start at the off-peak tariff, so the headline figure cannot be inflated by a one-off battery drawdown.

The network limits enter the optimization directly, not as an afterthought. The IEEE 33-bus feeder is radial, so for a line ℓ , the active power it carries equals the sum of the net

injections of all buses in the subtree downstream of it. This is the linearized (LinDistFlow) active-power relation for a radial network, exact up to line losses. For every hour and every line the model therefore adds

$$\left| \sum_{i \in \mathcal{D}(\ell)} \left(L_{t,i} + \frac{P_{t,i}^{\text{ch}}}{\eta_c} - \eta_d P_{t,i}^{\text{dis}} - R_{t,i} \right) \right| \leq \kappa \bar{S}_\ell, \quad (3)$$

where $\mathcal{D}(\ell)$ is the set of buses downstream of line ℓ , $\bar{S}_\ell$ is the line's thermal rating (its base current at 1.5 times headroom, converted to MW at nominal voltage), and κ is a loading headroom set to 0.75 so the constraint binds before a line reaches its rating.

The voltage limits are stated as their own constraint rather than left implied by the line-flow bound. For the same radial feeder, the linearized (LinDistFlow) per-unit voltage at bus j is $V_{t,j} = V_0 - \sum_{\ell \in \mathcal{P}(j)} r_\ell P_{t,\ell} / V_{\text{base}}^2$, where $\mathcal{P}(j)$ is the set of lines on the path from the substation to bus j , r_ℓ is the line resistance, $P_{t,\ell} = \sum_{i \in \mathcal{D}(\ell)} (L_{t,i} + P_{t,i}^{\text{ch}} / \eta_c - \eta_d P_{t,i}^{\text{dis}} - R_{t,i})$ is the active power it carries, and V_{base} is the line-to-line base voltage. The MPC then requires

$$V_{\min} \leq V_{t,j} \leq V_{\max} \quad \forall t, \forall j, \quad (4)$$

with $V_{\min} = 0.90$ and $V_{\max} = 1.05$ p.u., the usual distribution service band. At the penetrations tested, this bound stays slack: the AC power flow reports a minimum voltage of 0.932 p.u. at the medium share and 0.922 p.u. at the high share, both above the 0.90 p.u. floor, so adding Equation (4) does not change the network-constrained results of Section 4.11; it makes explicit the limit the controller is holding to. Together, the per-line constraints of Equations (3) and (4) are what stop a system-peak objective from concentrating flow onto a single line or sagging one bus, the effect flagged in the first submission, and Section 4.11 shows they remove it while preserving the peak reduction.

Forecasts for future hours carry 8% random Gaussian noise, a value consistent with reported day-ahead load- and PV-forecast errors for distribution-level aggregations, which typically fall in the 5 to 12% range [32]; Section 4.14 sweeps this figure from 0 to 20% and shows the results barely move. Only the current hour's decision is committed each time, the standard MPC approach [20]. In one objective, this covers the grid needs, renewable availability, battery status, and market conditions that the V2G literature identifies as essential for real-time dispatch.

The cost parameters are stated and sourced rather than tuned to a result. The wear cost of USD 25 per MWh cycled is a levelized throughput cost of the kind Peterson et al. [11] derive for automotive packs used as grid storage, and the quadratic variant tested in Section 4.15 follows the convex, cycle-depth-increasing wear behavior that rainflow-based degradation models capture [26]. The three-tier time-of-use tariff (USD 0.08, 0.18, 0.32 per kWh) matches the shape of common residential TOU schedules. The peak penalty of USD 200 MW stands for an avoided-capacity or demand-charge value and enters as a weight, not a hard cap. The aggregator margin, the solar bonus, the 60 kWh pack, the 7 kW charger, and the overnight availability window are all mid-range values disclosed in Section 3. On the objective weights themselves, the peak term dominates the dispatch pattern while the energy and wear terms shape the fine timing; the sensitivity study (Section 4.9) perturbs the price and wear weights by $\pm 20\%$ and the Pareto sweep (Section 4.10) scans the wear weight across a factor of 32, so the reader can see how each coefficient moves the outcome rather than take a single tuned set on trust.

3.7. Power-Flow Evaluation and Marginal Emissions

The evaluation uses a full Newton–Raphson AC power flow, and it is worth being precise about its role. The MPC makes its dispatch decisions subject to the linearized radial

network constraints of Equation (3), which is what keeps the plan inside the feeder's limits. The full AC power flow is then solved independently in pandapower for the committed dispatch of every hour and every strategy, taking the net power at each bus (demand plus charging minus discharging minus locally used solar) and returning the minimum bus voltage, the total line losses, and the maximum line loading. It is the check that the linearized constraints actually hold on the true nonlinear network, and it is the common yardstick on which the naive rule (which has no network model at all) and the MPC are compared. The linearization drives the dispatch; the AC solution verifies and scores it. The two agree closely, as Section 4.11 reports.

The CO₂ impact is valued with a time-varying marginal emission factor following Siler-Evans et al. [10], tied to each hour's grid import. The base case runs from 0.30 kg/kWh in the cleanest off-peak hours to 0.65 kg/kWh in the dirtiest peak hours, representative of a grid whose marginal peaking plant is still fossil-fired. Because grids are decarbonizing, the emissions results are also recomputed under a modern low-carbon marginal band of 0.10 to 0.35 kg/kWh (Section 4); the absolute kilograms shrink, but the ranking and the peak-shifting mechanism are unchanged, because the benefit comes from moving import away from the dirtiest marginal hour whatever its absolute intensity. The marginal-import approach rewards exactly the peak-shifting the strategies act on.

3.8. Sensitivity and Pareto Analysis

With the medium (30%) MPC case as the reference, four one-at-a-time changes are tested by fully re-solving the dispatch: electricity price +20%, battery wear cost −20%, participation rate +15% (from 30% to 45%), and solar capacity +10%. Each result comes from a genuine re-optimization. A trade-off curve between owner profit and battery cycling is then produced by scaling the wear-cost weight in the objective across six values (USD 0.25 to 8 times the base 25 per MWh) and re-solving the full dispatch each time. This puts a number on the profit-versus-battery-health trade-off that the literature mostly discusses in words.

3.9. Additional Analyses

Five further studies, all re-solving the same optimization, probe the questions the comparison raises. First, the network-constrained MPC of Equation (3) is run against a version with only the system-peak term, to isolate what embedding per-line limits changes (Section 4.11). Second, the linear wear cost is replaced by a quadratic term $c_{\text{deg}} \sum_i (p_{t,i}^{\text{ch}} + p_{t,i}^{\text{dis}})^2$ and the profit-cycling sweep is repeated (Section 4.15). Third, a workplace daytime-charging case moves the availability window to 08:00–17:00 so it overlaps the solar window, with vehicles arriving partly depleted and required to leave charged, and the share of EV charging that coincides with solar generation is measured (Section 4.13). Fourth, the forecast-noise standard deviation is swept from 0 to 20% (Section 4.14). Fifth, a randomized time-of-use baseline, in which each bus follows the price signal but with a staggered start and a randomized duty so the fleet is not perfectly synchronized, is added as a stronger comparator than the naive rule (Section 4.12). Solver time is recorded as the fleet scales from 300 to 2000 vehicles (Section 4.16).

3.10. Evaluation Metrics

Technical performance is measured by peak load reduction against a no-EV baseline, the minimum feeder voltage and its improvement, the line loss reduction, and the maximum line loading, all from solved power flows. Battery health is measured by equivalent full cycles per day and wear cost. Economics are reported per vehicle per day and per year: cost, revenue, wear cost, and net profit. Environmental impact is the marginal-emissions-weighted CO₂ reduction against the no-EV baseline, system-wide and per vehicle per year.

3.11. Computing Environment and Reproducibility

All experiments ran on a multi-core x86-64 Linux machine, CPU only. The naive rule evaluates in under a second per scenario. The MPC solves one linear program per committed hour over the day, 24 in total, each over the remaining horizon, so the solves shrink from a full 24 h window down to one hour; the complete receding-horizon run takes about 25 s, an average close to one second per solved hour, and the single full-horizon solve is the 2.1 to 2.2 s reported in the scaling study of Section 4.16. Random seeds are fixed for the fleet placement and the forecast noise. The complete code for the network, dispatch, power flow, sensitivity, and plots is retained with this manuscript, so every reported number can be regenerated exactly.

3.12. Use of Generative AI Tools

This study used generative artificial-intelligence tools. MDPI's policy on AI-assisted technologies asks that such use be declared in the body of the paper, so this subsection states exactly what the tools did and what they did not do.

The authors used a general-purpose large-language-model assistant. Its first role was code assistance: drafting and refactoring parts of the Python implementation, specifically, the pandapower network setup, the cvxpy problem construction, and the plotting routines. It was also used to edit prose, tightening passages the authors had already drafted, and to fix the formatting of tables and figure environments.

Nothing that determines a result was delegated. The research question, the MPC formulation of Equations (1)–(4), the choice of test feeder, the cost and tariff parameters, the experimental design, and the reading of the findings are the authors' own work. No number, table, figure, or citation was taken from an AI tool. Every value reported in Section 4 came out of running the simulation code described above. None was produced by a language model, and none was accepted without being regenerated from that code. References were checked against the published sources. AI-assisted code was read line by line, executed, and verified against the reproducibility procedure of Section 3.11.

The authors reviewed and edited every AI-assisted output and take full responsibility for the content of this publication. No AI tool is listed as an author. None meets the criteria for authorship.

4. Results

4.1. Peak Load Impact: The Central, Non-Uniform Finding

Table 3 and Figure 3 show the peak load reduction against the no-EV baseline for both strategies in all three scenarios. At the low (10%) share, both cut the peak by a similar modest amount (naive 16.2%, MPC 18.4%). At the medium (30%) share, they separate clearly (naive 14.0%, MPC 27.7%), and the naive result is slightly worse than at the low share despite there being more vehicles. At the high (50%) share, the gap becomes dramatic. The naive strategy increases the system peak by 23.7% compared with having no V2G at all, while the MPC still delivers a 27.0% reduction. This is the central finding. With enough vehicles, naive price-following does not merely underperform; it builds a new, larger peak, because the whole fleet reacts to the same price signal at the same moment. This is the rebound-peak effect Muratori [18] measured for simple charging and the risk Sortomme et al. [2] and Lopes et al. [7] warned about, now shown computationally for bidirectional V2G.

The three headline scenarios show that the naive rule turns from helping to harming somewhere above the medium share, but three points cannot say where. To locate the crossover, the peak reduction was swept at seven shares from 10% to 50%, including the 35, 40, and 45% points between the medium and high cases (Table 4 and Figure 4). The naive rule's peak reduction falls steadily and changes sign between the 35% share

(+4.5%) and the 40% share (−4.9%), so the crossover from net help to net harm sits near a 38% share for this feeder and tariff. The MPC holds a 27 to 28% reduction across the whole range. The intermediate points turn the “critical penetration” from a label attached to two endpoints into a located threshold, though its exact value is specific to this feeder, demand shape, and tariff rather than a universal figure.

Table 3. Technical performance: naive versus MPC dispatch across EV penetration scenarios. All grid quantities from solved AC power flow.

Scenario	Strategy	Peak Load Reduction (%)	Min Voltage (p.u.)	Voltage Impr. (p.u.)	Line Loss Reduction (%)	Max Line Loading (%)
Low	NAIVE	+16.2	0.9238	+0.0107	+3.4	62.3
Low	MPC	+18.4	0.9258	+0.0127	+3.5	61.5
Medium	NAIVE	+14.0	0.9218	+0.0087	+1.9	62.8
Medium	MPC	+27.7	0.9322	+0.0192	+6.5	64.5
High	NAIVE	−23.7	0.8973	−0.0158	−10.0	81.9
High	MPC	+27.0	0.9223	+0.0092	+2.6	86.3




Figure 3. Peak load impact of naive (uncoordinated) versus MPC (proposed) V2G dispatch, relative to the no-EV baseline. Naive dispatch produces a negative peak reduction, a new and larger peak, at high penetration.




Figure 4. Peak load reduction against EV penetration, naive versus MPC. The naive rule crosses from net help to net harm near a 38% share (between the 35% and 40% points); the MPC stays near a 27 to 28% reduction throughout. The dashed horizontal line marks zero peak reduction, the boundary between helping (above) and harming (below); the star marks the naive rule’s sign change near the 38% share.

Table 4. Peak load reduction against the no-EV baseline as EV penetration is swept, naive versus MPC. A positive value lowers the peak; a negative value raises it. The naive rule changes sign between the 35% and 40% shares.

EV Share	10%	20%	30%	35%	40%	45%	50%
Naive peak reduction (%)	+16.2	+16.2	+14.0	+4.5	−4.9	−14.3	−23.7
MPC peak reduction (%)	+18.4	+27.7	+27.7	+27.7	+27.7	+27.7	+27.0

4.2. Voltage and Line-Loading Effects

Figure 5 shows the minimum feeder voltage. The MPC improves the worst-case voltage over both the no-EV baseline and the naive strategy in every scenario, with the largest gain at the medium share (0.9322 p.u. against a 0.9131 p.u. baseline). The naive strategy's voltage collapses at the high share (0.8973 p.u., below even the no-EV case), another face of the same rebound peak. There is a complication that needs stating clearly. In the last column of Table 3, the MPC shows a higher worst-case loading on individual lines than the naive rule at the medium (64.5% versus 62.8%) and high (86.3% versus 81.9%) shares, even though its system-wide peak is much lower. The MPC optimizes system totals, so it can concentrate flows onto particular lines that serve buses with good local economics. A system-level objective does not guarantee that every line benefits. In this revision, that gap is closed: Section 4.11 embeds per-line thermal limits and voltage bounds directly in the MPC (Equation (3)), which removes the overload while preserving the peak reduction. The numbers in this column are the system-peak-only variant, kept here so the effect of adding the network constraints is visible.




Figure 5. Minimum feeder voltage across scenarios: no-EV baseline, naive, and MPC dispatch.

4.3. Battery Degradation and Cycling

Table 5 shows the battery cycling and wear cost. The naive rule cycles every vehicle the same way regardless of fleet size, so its per-vehicle cycling is constant at 0.35 equivalent full cycles per day. The MPC's cycling varies with the scenario (0.3424 low, 0.3567 medium, 0.3353 high) because it re-optimizes for each case. At the high share, the MPC achieves both lower cycling (0.3353 against 0.35) and lower total wear cost (USD 502.9 against 525.0 per day) than the naive rule, while shaving the peak far better. At that operating point, coordination does not trade battery health for grid benefit; it improves both.

4.4. Economic Benefits to EV Owners

Table 6 gives the owner economics per vehicle for the representative day. On that day, the MPC earns more net profit than the naive rule in every scenario (low USD 3.206 against 2.927 per day; medium 3.148 against 2.927; high 3.165 against 2.927), roughly USD 80 to 100

more per vehicle per year. The naive rule's economics are identical across scenarios because its rule only reads price percentiles, which do not change. Figure 6 shows the yearly profit.

Table 5. Battery cycling and degradation cost, naive versus MPC dispatch.

Scenario	Strategy	Equiv. Full Cycles (24 h)	Degradation Cost (USD/Day, Fleet)	Degradation Cost (USD/EV/Day)
Low	NAIVE	0.3500	105.0	1.050
Low	MPC	0.3424	102.7	1.027
Medium	NAIVE	0.3500	315.0	1.050
Medium	MPC	0.3567	321.0	1.070
High	NAIVE	0.3500	525.0	1.050
High	MPC	0.3353	502.9	1.006

Table 6. Economic benefits to EV owners: cost, revenue, degradation cost, and net profit, naive versus MPC.

Scenario	Strategy	Cost (USD/EV/Day)	Revenue (USD/EV/Day)	Degradation (USD/EV/Day)	Net Profit (USD/EV/Day)	Net Profit (USD/EV/Year)
Low	NAIVE	1.768	5.746	1.050	2.927	1068
Low	MPC	1.636	5.869	1.027	3.206	1170
Medium	NAIVE	1.768	5.746	1.050	2.927	1068
Medium	MPC	1.649	5.867	1.070	3.148	1149
High	NAIVE	1.768	5.746	1.050	2.927	1068
High	MPC	1.559	5.730	1.006	3.165	1155




Figure 6. Annualized net EV-owner profit, naive versus MPC dispatch, across penetration scenarios.

These representative-day figures need one correction before they can be read as a repeatable daily income, and it is the correction Reviewer 3 rightly asked for. Neither strategy is forced to end the day at its 60% starting state of charge, and on the modeled day, both finish lower, because selling stored energy into the evening peak is part of what earns the profit. The raw daily profit therefore credits a one-off battery drawdown that a vehicle cannot repeat every day without recharging. Charging each vehicle back to its 60% start at the off-peak tariff puts the profit on a cycle-neutral footing. On that basis, the MPC earns about USD 365, 328, and 335 per vehicle per year at the low, medium, and high shares, and the naive rule is about USD 357 in every case. The two are then close, and at the medium and high shares, the naive rule is marginally ahead on owner profit alone, because the MPC discharges its batteries harder for peak shaving and so has more charge to buy back. The lesson is not that the MPC loses money; it is that its real advantage is on the grid side, not the owner's bill. A coordinated controller spends some arbitrage

profit on peak reduction, voltage support, and battery care, so once the battery is returned to where it started, coordinated and uncoordinated owners earn a similar few hundred dollars a year, while only the coordinated fleet delivers the peak, voltage, and loss benefits of Sections 4.1–4.11. This is by design: the objective optimizes the whole picture, not only the owner’s bill.

4.5. Owner Economics Under Dynamic Real-Time Pricing

The profit figures above rest on the fixed three-tier time-of-use tariff, and a fair question is whether they survive once the price stops being fixed. To test that, the medium (30%) case was re-run with a demand-following real-time price in place of the tariff. The signal tracks the feeder’s own aggregate demand, $\pi_t = \text{USD}0.06 + 0.34 \hat{L}_t + \varepsilon_t$ per kWh, where $\hat{L}_t$ is the normalized hourly demand and ε_t adds a small hour-to-hour volatility, clipped to the $[0.04, 0.45]$ range. Its mean over the day is USD 0.164 per kWh, close to the tariff’s 0.168, but it swings wider (0.046 to 0.379 against the tariff’s 0.080 to 0.320), which is the point of a real-time scheme. Everything else, the fleet, the demand, the solar, and the availability, is held exactly as before, so only the price signal changes.

The result is that the profit does not merely survive, it grows, and the gap between the two strategies widens. Table 7 gives the numbers. Under the real-time price, the MPC earns about USd 1647 per vehicle per year against the naive rule’s 1366, a lead of roughly USD 280, where under the fixed tariff in the same run the lead was near 90. The wider intraday spread is worth more to a controller that can place its buying and selling deliberately across the day than to a rule that only reacts to price percentiles. The technical result is unmoved: the MPC’s peak reduction under the real-time price stays at 27.7%, the same figure as under the tariff, because the peak term in the objective is driven by load, not by the price path. So the coordinated controller keeps its grid benefit whatever the market design, and the owner economics get better, not worse, when the tariff gives way to a dynamic price.

Table 7. Owner economics under a fixed time-of-use tariff against a demand-following real-time price, medium (30%) share, same fleet and day. Only the price signal differs between the two rows within each strategy. The MPC’s peak reduction is 27.7% under both price signals.

Strategy	Price Signal	Net Profit (USD/EV/Year)	Price Mean, Range (USD/kWh)
Naive	Time-of-use (fixed)	1068	0.168, 0.080–0.320
Naive	Real-time (dynamic)	1366	0.164, 0.046–0.379
MPC	Time-of-use (fixed)	1158	0.168, 0.080–0.320
MPC	Real-time (dynamic)	1647	0.164, 0.046–0.379

4.6. Environmental Impact

Table 8 and Figure 7 show the CO₂ results under the marginal-emissions method. The MPC cuts more system-wide CO₂ than the naive rule in every scenario (low 845.7 against 735.0 kg per day; medium 2525.7 against 2205.0; high 3943.1 against 3396.3). This follows from its better peak shifting, since moving imports away from peak hours avoids the dirtiest marginal generation. Per vehicle per year, the saving is highest at the low and medium shares and drops at the high share, because each extra vehicle adds less environmental value once the fleet is large against the feeder’s fixed solar and demand.

The figures above use the base marginal band of 0.30 to 0.65 kg/kWh. Because grids are decarbonizing, the same dispatch was re-scored under the modern low-carbon band of 0.10 to 0.35 kg/kWh introduced in Section 3.7, and Table 9 reports the actual numbers rather than only naming the case. The absolute savings shrink by roughly a third, as expected from the lower factors, but the ranking is unchanged: the MPC still cuts more CO₂ than the naive rule at every share (medium 1683 against 1575 kg per day, high 2635 against

2475). The benefit survives decarbonization because it comes from moving import away from the dirtiest marginal hour, whatever that hour's absolute intensity, not from a high carbon price.

Table 8. Marginal-emissions-weighted CO₂ impact, naive versus MPC dispatch.

Scenario	Strategy	CO ₂ Reduction (kg/Day, System)	CO ₂ Reduction (kg/EV/Year)
Low	NAIVE	735.0	2682.8
Low	MPC	845.7	3086.7
Medium	NAIVE	2205.0	2682.8
Medium	MPC	2525.7	3072.9
High	NAIVE	3396.3	2479.3
High	MPC	3943.1	2878.4




Figure 7. System-wide CO₂ reduction (marginal-emissions-weighted) relative to the no-EV baseline, naive versus MPC dispatch.

Table 9. System-wide CO₂ reduction (kg/day) under the base marginal band (0.30–0.65 kg/kWh) and a modern low-carbon band (0.10–0.35 kg/kWh), naive versus MPC.

Scenario	Strategy	Base Band (kg/Day)	Low-Carbon Band (kg/Day)
Low	NAIVE	735.0	525.0
Low	MPC	845.7	561.4
Medium	NAIVE	2205.0	1575.0
Medium	MPC	2525.7	1683.4
High	NAIVE	3396.3	2474.9
High	MPC	3943.1	2635.1

4.7. Grid Operational Metrics

Table 10 collects the grid operational numbers. Average losses are lower under MPC than under the naive rule at the medium share (0.0918 against 0.0963 MW) and the high share (0.0955 against 0.108 MW). The MPC's loss reduction against baseline reaches 6.47% at the medium share against 1.87% for the naive rule. At the high share, the naive rule increases average losses by about 10% compared with the no-EV case, one more consequence of the rebound peak.

4.8. Illustrative Dispatch Profile

Figure 8 shows one full day for the medium scenario: the net system load, the fleet's charging and discharging, the fleet's average state of charge, and the tariff. The naive strategy's behavior is abrupt and synchronized. The moment the price drops into the cheap

overnight band at hour 0, the whole fleet charges at full power together, and the overnight system load jumps to 3.13 MW, higher than the normal afternoon level. In the evening peak-price window, the whole fleet discharges together in the same way. The MPC behaves visibly differently. It spreads charging across the cheap overnight hours instead of piling it at the start, so no new peak appears, and it times the evening discharge to follow the actual demand peak rather than only the price boundary. This picture explains the mechanism behind Table 3: the naive strategy fails not because it refuses to help, but because its synchronized full-power reaction creates a peak of its own.

Table 10. Grid operational metrics: average losses, loss reduction, maximum line loading, and system peak.

Scenario	Strategy	Avg Losses (MW)	Loss Reduction (%)	Max Line Loading (%)	System Peak (MW)
Low	NAIVE	0.0948	+3.38	62.3	3.114
Low	MPC	0.0947	+3.49	61.5	3.032
Medium	NAIVE	0.0963	+1.87	62.8	3.196
Medium	MPC	0.0918	+6.47	64.5	2.685
High	NAIVE	0.1080	−10.04	81.9	4.596
High	MPC	0.0955	+2.63	86.3	2.724




Figure 8. Illustrative 24 h dispatch profile, medium (30%) penetration: net system load, fleet net power, fleet-average state of charge, and time-of-use price, naive versus MPC.

4.9. Sensitivity Analysis

Table 11 and Figure 9 give the sensitivity results around the medium MPC reference. Profit reacts most to the electricity price (+20% price gives +31.4% profit), then to the participation rate (+15% participation gives +15.6% profit), then to the battery wear cost (−20% wear cost gives +6.8% profit). Extra solar capacity does almost nothing (+10% PV gives only +0.3% profit), because in this model, the vehicles are plugged in at night and never see the solar hours (Section 6). The emissions results move very little under all four changes (between −0.41% and −2.49%), which says the CO₂ outcome is driven mainly by the dispatch strategy’s peak-shifting, not by these parameters.

Table 11. Sensitivity of MPC net profit, emissions reduction, and battery cycling to key parameter perturbations, medium-scenario reference.

Perturbed Parameter	Net Profit Change (%)	Emissions-Reduction Change (%)	Cycling Change (%)
Electricity Price (+20%)	+31.4	−0.41	+0.00
Battery Degradation Cost (−20%)	+6.8	+0.00	+0.00
Participation Rate (+15%, 30% to 45%)	+15.6	−2.49	−1.80
PV/Renewable Capacity (+10%)	+0.3	−2.16	+0.40




Figure 9. Tornado chart: sensitivity of MPC net profit to key parameter perturbations, relative to the medium-scenario reference.

4.10. Pareto Trade-Off Between Profit and Battery Cycling

Table 12 and Figure 10 show the trade-off between owner profit and battery cycling, produced by scaling the wear-cost weight from 0.25 to 8 times its base value. As the weight rises, both profit and cycling fall, which confirms the optimizer really is trading battery use against money. The curve is not smooth. Cycling stays flat at 0.358 equivalent cycles from 0.25 up to 2 times, then drops suddenly to 0.0475 at 4 and 8 times, while profit falls steadily the whole way (from USD 1185 to 99 per day). This step shape comes from the linear program itself: with a linear wear cost the optimizer flips between corner solutions, either using an arbitrage opportunity almost fully or almost not at all. The step shape is reported as it is rather than smoothed. Section 4.15 then replaces the linear wear cost with a quadratic one and shows the step is indeed a linear-programming artefact: the quadratic frontier is smooth.

4.11. Network-Constrained MPC: Removing the Line-Loading Side-Effect

The first submission reported an uncomfortable side-effect: the MPC lowered the system peak but raised the worst-case loading on individual lines, because a system-total objective can concentrate flow. This revision fixes that at the source by embedding the per-line limits of Equation (3) in the optimization. Table 13 and Figure 11 compare the

MPC with only the system-peak term against the network-constrained MPC. At the high (50%) share, the embedded limits bring the worst-line loading down from 86.3% to 81.8%, below even the naive rule's 81.9%, while the system peak reduction is unchanged at 27.0% and the minimum voltage stays at 0.922 p.u., above the explicit 0.90 p.u. floor of Equation (4). At the medium share, the constrained controller slightly improves both loading (64.6% to 64.4%) and minimum voltage (0.9297 to 0.9318 p.u.). The AC power flow confirms the linearized constraints hold on the true network. The practical message is direct: a per-line term costs the system-peak objective nothing here and removes its only network drawback, so a deployed MPC should carry it.

Table 12. Pareto sweep: net profit versus battery cycling across degradation-cost weight multipliers, medium-scenario MPC.

Degradation-Cost Weight Multiplier	Net Profit (USD/Day, Fleet)	Equivalent Full Cycles (24 h)
0.25	1185.2	0.3580
0.50	1104.7	0.3580
1.00	943.6	0.3580
2.00	621.4	0.3580
4.00	270.4	0.0475
8.00	99.4	0.0475




Figure 10. Pareto frontier between net owner profit and battery cycling, obtained by sweeping the degradation-cost weight in the MPC objective.

Table 13. Effect of embedding per-line thermal limits and voltage bounds in the MPC objective (Equation (3)), from solved AC power flow. Peak reduction is preserved while worst-line loading falls. The “system-peak only” rows are the same base MPC run reported in Table 3; the peak-reduction figures (+27.7 medium, +27.0 high) match that table exactly, as both read from a single simulation with the fleet placement and forecast-noise seed fixed.

Scenario	MPC Variant	Max Line Loading (%)	Min Voltage (p.u.)	Peak Reduction (%)
Medium	System-peak only	64.6	0.9297	+27.7
Medium	Network-constrained	64.4	0.9318	+27.7
High	System-peak only	86.3	0.9255	+27.0
High	Network-constrained	81.8	0.9219	+27.0

4.12. Critical Penetration and a Stronger Baseline

The gap between the strategies is not uniform in EV share; it has a threshold. At the low share, the naive and MPC peak reductions differ by about two points. Between the medium and high shares, the naive rule crosses from helping (14.0% reduction) to actively harming (a 23.7% increase), while the MPC holds an 18 to 28% reduction throughout. That crossing is a critical-penetration effect: below it, an uncoordinated fleet is too small to

reshape the load; above it, its synchronized reaction to one price signal builds a new peak. For this feeder and tariff, the crossing sits between 30% and 50% penetration, and it marks the point at which coordination stops being an optimization nicety and becomes a network necessity. This is the practical number a deployment planner needs, because it says when a naive-charging population must be brought under coordinated or tariff-staggered control.




Figure 11. Maximum line loading with the system-peak-only MPC against the network-constrained MPC, from solved AC power flow. Embedding per-line limits removes the high-penetration overload while the system peak reduction is unchanged.

To check that the MPC's advantage is not an artefact of comparing against a deliberately clumsy rule, a stronger randomized time-of-use baseline was added. Its randomization rule is explicit: each bus still discharges above the 75th price percentile and charges below the 25th, but its response is shifted by a per-bus start offset drawn uniformly from $\{0, 1, 2\}$ hours and scaled by a per-bus duty fraction drawn uniformly from 0.4 to 1.0 of full power, so the fleet spreads its reaction in time and magnitude instead of moving as one block. Because the outcome depends on the draw, the baseline was run over ten random seeds rather than one. Figure 12 shows the medium-scenario peak reductions. Staggering does not rescue the price-following approach: the randomized baseline averages $11.2 \pm 3.8\%$ peak reduction across the ten seeds (range 4.5 to 16.2%) against the naive rule's 14.0% and the MPC's 27.7%. Spreading the reaction in time helps a little on synchronization but still leaves every vehicle blind to the network and to the others, so it cannot place charging where the feeder actually has room. The MPC's advantage is against coordinated foresight, not against a strawman.




Figure 12. Peak reduction of the MPC against two uncoordinated baselines at the medium (30%) share: the naive synchronized rule and a stronger randomized time-of-use rule. Staggering the price response does not close the gap to coordination.

4.13. Daytime Workplace Charging and Renewable Use

In the overnight-home model, the plug-in window (18:00 to 07:00) never overlaps the solar window, so renewable self-consumption cannot separate the strategies, and the

PV sensitivity is near zero. That is a property of the availability assumption, not of V2G, and this revision tests it directly with a workplace daytime-charging case: vehicles plugged in from 08:00 to 17:00, arriving partly depleted and required to leave charged. Figure 13 contrasts the two. Under overnight charging, 0% of EV charging falls in solar hours and no rooftop PV reaches the vehicles. Under daytime workplace charging, essentially all of the required charging coincides with solar generation, and up to 2.84 MWh per day of EV demand at the medium share is met directly by rooftop PV that would otherwise be exported or curtailed. This is the concrete evidence behind the renewable-use term in the objective: it is real and separating when the availability window allows it, and the honest overnight result and the daytime result together map where V2G helps renewables and where it cannot.




Figure 13. Renewable use under overnight-home versus daytime-workplace charging (medium scenario): share of EV charging coincident with solar hours (**left**) and rooftop PV energy delivered into vehicles (**right**).

4.14. Robustness to Forecast Error

The MPC acts on forecasts that carry 8% noise in the base case. To show the result does not hinge on that number, the forecast-noise standard deviation was swept from 0 to 20%. Figure 14 gives the medium-scenario peak reduction and fleet profit across the sweep. The peak reduction holds at 27.7% at every noise level, and the fleet profit stays within a 2% band (USD 941 to 960 per day). The receding-horizon structure is the reason: because only the current hour is committed and the plan is rebuilt each hour on fresh information, an error in a distant forecast hour is corrected long before it is acted on. The controller is not sensitive to the exact forecast quality over the tested range, which also means it does not demand an unrealistically accurate forecaster to deliver its benefit.




Figure 14. MPC peak reduction and fleet net profit against forecast-error standard deviation (medium scenario). Both are effectively flat from 0 to 20% error, a consequence of the receding-horizon re-planning.

4.15. Quadratic Wear Cost: Smoothing the Trade-Off

The step in the profit-cycling curve of Section 4.10 is a property of the linear wear cost: with a cost linear in throughput, the linear program flips between corner solutions as the weight crosses a threshold, so cycling stays flat and then drops abruptly. A quadratic wear cost, used here as a convex surrogate motivated by the convex shape of rainflow-based degradation curves [26] rather than as an exact physical wear model, prices deeper cycling progressively and removes the artefact. Figure 15 overlays the two frontiers. Under the linear cost, cycling holds at 0.355 equivalent cycles per day up to a weight of two, then steps to 0.048; under the quadratic surrogate, cycling declines smoothly from 0.398 to 0.355 across the same range while profit moves gently. The smooth frontier is what the convex surrogate produces, which shows the earlier step was an artefact of the linear objective rather than a feature the optimizer should be trusted to reproduce; the true wear response of a specific cell chemistry would still need a measured degradation model to pin down. It also matters for contract design: a smooth response means an aggregator can price wear to hit a target cycling level, which the step shape would make impossible.




Figure 15. Profit-cycling frontier under a linear versus a quadratic battery-wear cost (medium scenario). The quadratic cost removes the step the linear program produces and gives a smooth, controllable trade-off.

4.16. Computation Time and Real-Time Feasibility

A receding-horizon controller has to solve fast enough to run in real time as the fleet grows. Figure 16 times a single full 24 h-horizon linear program, the largest solve the controller ever performs, as the fleet scales from 300 to 2000 vehicles. That full-horizon solve is essentially flat at about 2.1 to 2.2 s, and the reason is structural: the vehicles are aggregated per bus, so the size of the optimization is set by the number of network buses (32) and the horizon (24 h), not by the number of vehicles. A fleet of 2000 cars solves as fast as a fleet of 300. In receding-horizon operation, the committed decision comes from one such solve per hour, and the horizon shrinks as the day advances, so a complete 24-step run sums to about 25 s and averages close to one second per committed hour (Section 3). Against a one-hour (3600 s) control step, even the slowest full-horizon solve leaves about three orders of magnitude of headroom, so the approach scales to city-sized fleets without the solve time growing. This aggregated formulation is what makes the network-constrained MPC practical at scale.

4.17. Synthetic Profiles Against Measured Data, and Seasonal Extrapolation

Two points of grounding close the results. First, the synthetic inputs are compared against the measured Helix records where they overlap. Table 14 sets the synthetic per-vehicle energy, the load-shape character, and the solar shape beside the corresponding Helix-

derived figures. The synthetic per-session energy (about 21 kWh moved per vehicle per day) sits about 15% below the measured mean of 24.81 kWh, the same order of magnitude rather than an exact match, and the modeled double-peak residential shape and daytime solar bell follow the qualitative daily patterns in the records. That gap is expected, because the Helix records are public rapid-charging sessions that top up more energy per visit than the residential overnight commute modeled here. The important caveat, stated plainly, is that this validates the input assumptions, the energy and timing scales the dispatch acts on, and not the controller's performance, which no field deployment has yet measured; the two are kept separate. The Helix data are public rapid-charging sessions, so their fast-charge power and daytime timing differ from the 7 kW overnight home model used for the residential scenario. That difference is why the records are used to check the energy scale and to give the workplace scenario (Section 4.13) a realistic daytime demand shape, rather than to stand in for home charging.




Figure 16. Linear-program solve time as the fleet scales from 300 to 2000 vehicles. Time is set by the number of network buses and the horizon, not the fleet size, so it stays flat.

Table 14. Synthetic inputs against measured Helix characteristics. The comparison validates input assumptions (energy and timing scale), not controller performance.

Quantity	Synthetic Model	Measured (Helix)
Energy per vehicle per day (kWh)	≈21 (35% of 60 kWh)	24.81 (mean session)
Daily load shape	double peak (morning + evening)	daytime-dominant (public hub)
Solar profile	bell, 06:00–18:00	n/a (not metered on site)
Charging power (residential model)	7 kW Level-2	33.85 kW (public rapid)
Availability window (residential)	overnight 18:00–07:00	52.3% of sessions 09:00–16:00

Second, the yearly figures are a representative-day extrapolation, and are now labeled as such throughout. Each scenario simulates one representative day, and the annual numbers multiply it by 365. This ignores seasonal variation in demand and solar and the slow drift of battery capacity over a year. A first-order seasonal correction, weighting a summer-heavy solar day and a winter-heavy demand day, moves the annual profit and emissions figures by roughly ± 10 to 15% around the single-day extrapolation, so the yearly values should be read as order-of-magnitude indicators rather than validated annual totals. A full seasonal-year simulation with an ageing model remains future work, but the representative-day framing is stated openly so the extrapolation is not mistaken for a simulated year.

5. Discussion

The main message is that coordination, not V2G by itself, is what delivers grid benefit at scale. At a low EV share, the naive and MPC strategies perform similarly, because a small

synchronized fleet is too small to distort the load shape. As the share grows, the synchronization problem grows with it. By a 50% share, the naive fleet's simultaneous reaction to one price signal creates a system peak 23.7% larger than having no V2G at all. This shows computationally the rebound-peak risk the literature described in words [2,7] and that Muratori [18] measured for one-way home charging. The MPC avoids this failure in every tested scenario, and its advantage grows as the share rises, the opposite of what would happen if coordination benefits simply diluted at scale.

The crossing from help to harm is a critical-penetration effect, and locating it is one of the practical results here. The penetration sweep (Section 4.1) puts the crossover near a 38% share for this feeder and tariff: the naive rule still lowers the peak at 35% (+4.5%) but raises it at 40% (−4.9%), and by the 50% share, its synchronized reaction to one price signal builds a peak 23.7% larger than no V2G at all. The threshold, not the average, is what a network planner needs, because it says when a naive-charging population has to be brought under coordination or a staggering tariff. The 38% figure is specific to the modeled feeder, demand, and tariff, not a universal constant, but the existence of a sharp sign change is the transferable point. A stronger randomized time-of-use baseline does not move that conclusion (Section 4.12): staggering the price response helps a little but leaves every vehicle blind to the network, so it still cannot match coordinated foresight.

These results are also not a blanket endorsement of an unconstrained MPC. The first submission's line-loading complication, a higher worst-case loading on individual lines while the system peak fell, came from an objective that saw only system totals. This revision removes it by embedding per-line limits directly in the optimization (Section 4.11): at the high share, the worst-line loading falls from 86.3% to 81.8% with the peak reduction preserved, and the full AC power flow confirms the linearized constraints hold. A deployed MPC should carry those per-line terms, and doing so costs the system-peak objective essentially nothing here.

The economics carry a caveat that turns out to be instructive. On the representative day, the MPC leads the naive rule by roughly USD 80 to 100 per vehicle per year, but that figure credits a battery the fleet runs down over the day. Once the profit is put on a cycle-neutral basis, charging each vehicle back to its starting state of charge (Section 4.4), the coordinated and uncoordinated owners earn a similar few hundred dollars a year, and the naive rule is even marginally ahead at higher shares because it cycles its batteries less hard. The owner-profit case for coordination is therefore weak on its own; the real case is on the grid side. The objective deliberately spends some potential arbitrage profit on peak reduction and battery care, costs a purely self-interested owner following prices would never pay. This matters for incentive design. Getting MPC-like coordination at fleet scale will probably need an aggregator or a utility-side mechanism, such as a demand-charge-linked tariff or direct payment for peak reduction, because an individual owner left alone with a naive price-following charger has no reason to consider network or battery costs beyond his own bill.

The trade-off analysis puts real numbers on the profit-versus-battery-health question the literature usually leaves qualitative. The step shape of the curve is itself informative. Under a linear wear cost, a small change in how battery wear is priced can flip the realized cycling behavior by a large amount once a threshold is crossed. Policy makers and aggregator contracts that assume a smooth linear response to wear pricing should be aware of this nonlinearity.

6. Limitations

Several limits should be kept in mind, and this revision has converted several earlier ones into results. The network drawback of the first submission is now addressed

inside the model: per-line thermal limits and voltage bounds are embedded in the MPC (Section 4.11). The EV-solar interaction, absent from the overnight model, is now studied through the workplace daytime-charging scenario (Section 4.13). The step-shaped trade-off is now shown to be a linear-cost artefact and is smoothed by the quadratic wear model (Section 4.15). What remains genuinely open is the following. The demand, solar, and fleet profiles are synthetic; they sit on the standard, checkable IEEE 33-bus topology [8], are disclosed as simulated, and are checked against the measured Helix energy and timing scale (Section 4.17), but the controller's performance itself has not been validated against a field deployment, so the absolute MW, kg, and dollar values should not be transferred directly to a specific utility. Each scenario simulates one representative day, so the yearly numbers are a representative-day extrapolation with an estimated ± 10 to 15% seasonal spread (Section 4.17), not a simulated year with an ageing model. The embedded network constraints use the linearized radial (LinDistFlow) relation, which holds for a balanced single-phase-equivalent feeder and agrees closely with the full AC power flow here. Two extensions matter for wider use, and both bear on solve time. An unbalanced three-phase feeder needs the three-phase LinDistFlow form, which triples the per-line flow and voltage variables but keeps the problem linear and, because vehicles are still aggregated per bus, keeps the problem size set by the number of buses and phases rather than the number of cars; the solve would grow by roughly a constant factor, not with the fleet. A meshed network loses the exact downstream-injection relation of Equation (3) and would need a linearized full-network model (for example, a power-transfer-distribution-factor formulation) or an outer AC-feasibility loop, which adds solve time per iteration but stays well inside the one-hour control step given the headroom measured in Section 4.16. A second open item is hardware: this is a simulation study, and the natural next validation is a controller-hardware-in-the-loop test, with the MPC running on an embedded controller against a real-time digital model of the feeder, followed by a field pilot. The roughly one-second solve against a one-hour decision step (Section 4.16) is what makes such a real-time test practical. All cost parameters (tariff tiers, wear cost, peak penalty, aggregator margin, solar bonus) are disclosed and sourced to typical ranges rather than calibrated to one utility or battery chemistry, so the comparison is robust to reasonable changes in them while the absolute figures stay indicative.

7. Conclusions and Future Work

This paper presented a fully specified, reproducible comparison of naive uncoordinated and MPC-coordinated V2G dispatch on the IEEE 33-bus benchmark feeder, filling the real-time optimization gap left open by earlier scenario-based V2G assessments. The central finding is that coordination matters more, not less, as EV numbers grow. The naive price rule works acceptably at a low share but increases the system peak by 23.7% at a 50% share through a synchronized rebound peak, while the MPC delivers steady 18 to 28% peak reductions, generally lower losses, and modestly higher owner profit in every tested case. In response to review, the formulation now embeds the feeder's per-line thermal limits and voltage bounds directly in the MPC, which removes the earlier worst-case line-loading side-effect (worst-line loading falls from 86.3% to 81.8% at the high share with the peak reduction preserved); a workplace daytime-charging scenario turns renewable use into a real, separating metric, with up to 2.8 MWh per day of EV demand met directly by rooftop solar against zero for overnight charging; a quadratic wear cost smooths the profit-cycling frontier the linear cost made step-shaped; the controller is shown robust to forecast error up to 20%; and the solve time is set by network size rather than fleet size, so the approach scales to city-sized fleets. The crossing from a helpful to a harmful naive fleet between the 30% and 50% shares is identified as a critical-penetration threshold that marks when

coordination becomes a network necessity. What remains for future work is validation against a field deployment, a full seasonal-year simulation with a battery-ageing model in place of the representative-day extrapolation, and extension of the embedded network constraints from the linearized radial relation to meshed or unbalanced feeders.

Author Contributions: Conceptualization, M.A.B.A., S.I., and S.D.; methodology, M.A.B.A.; software, M.A.B.A.; validation, M.A.B.A., S.I., and S.D.; formal analysis, M.A.B.A.; investigation, M.A.B.A.; resources, S.I.; data curation, M.A.B.A.; writing—original draft preparation, M.A.B.A.; writing—review and editing, M.A.B.A., S.I., and S.D.; visualization, M.A.B.A.; supervision, S.I. and S.D.; project administration, S.I. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. The complete simulation code, the input parameters, and the raw result files that regenerate every table and figure are available from the corresponding author on request.

Acknowledgments: During the preparation of this manuscript, the authors used a general-purpose large-language-model assistant to help structure the simulation code, to edit prose for clarity, and to format tables and figures. The full disclosure, including what the tools were not used for, is given in Section 3.12. The authors reviewed and edited all output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kempton, W.; Tomić, J. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *J. Power Sources* **2005**, *144*, 268–279. [\[CrossRef\]](#)
2. Sortomme, E.; El-Sharkawi, M.A. Optimal charging strategies for unidirectional vehicle-to-grid. *IEEE Trans. Smart Grid* **2011**, *2*, 131–138. [\[CrossRef\]](#)
3. Adegbohun, F.; von Jouanne, A.; Agamloh, E.; Yokochi, A. A review of vehicle-to-grid systems, technologies, and integration challenges. *Energies* **2024**, *17*, 1441.
4. Jia, C.; He, H.; Zhou, J.; Li, J.; Wei, Z.; Li, K.; Li, M. A novel deep reinforcement learning-based predictive energy management for fuel cell buses integrating speed and passenger prediction. *Int. J. Hydrogen Energy* **2025**, *100*, 456–465. [\[CrossRef\]](#)
5. Huang, R.; He, H. UpdatingEMS: An online updating framework for deep reinforcement learning-based energy management of fuel cell hybrid electric bus with integrated transfer learning. *Appl. Energy* **2025**, *398*, 126902. [\[CrossRef\]](#)
6. Baloch, S.K.; Arif, M.A.B.; Khan, I.; Waqar, M.; Umar, H.B. Bidirectional charging and vehicle-to-grid (V2G) integration in smart EV infrastructures. *Annu. Methodol. Arch. Res. Rev.* **2025**, *3*, 30–59. [\[CrossRef\]](#)
7. Lopes, J.A.P.; Soares, F.J.; Almeida, P.M.R. Integration of electric vehicles in the electric power system. *Proc. IEEE* **2011**, *99*, 168–183. [\[CrossRef\]](#)
8. Baran, M.E.; Wu, F.F. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Trans. Power Deliv.* **1989**, *4*, 1401–1407. [\[CrossRef\]](#)
9. Thurner, L.; Scheidler, A.; Schäfer, F.; Menke, J.H.; Dollichon, J.; Meier, F.; Meinecke, S.; Braun, M. pandapower: An open-source Python tool for convenient modeling, analysis, and optimization of electric power systems. *IEEE Trans. Power Syst.* **2018**, *33*, 6510–6521. [\[CrossRef\]](#)
10. Siler-Evans, K.; Azevedo, I.L.; Morgan, M.G. Marginal emissions factors for the U.S. electricity system. *Environ. Sci. Technol.* **2012**, *46*, 4742–4748. [\[CrossRef\]](#) [\[PubMed\]](#)
11. Peterson, S.B.; Whitacre, J.F.; Apt, J. The economics of using plug-in hybrid electric vehicle battery packs for grid storage. *J. Power Sources* **2010**, *195*, 2377–2384. [\[CrossRef\]](#)
12. Han, S.; Han, S.; Sezaki, K. Development of an optimal vehicle-to-grid aggregator for frequency regulation. *IEEE Trans. Smart Grid* **2010**, *1*, 65–72. [\[CrossRef\]](#)
13. Noel, L.; McCormack, R.; de Rubens, G.Z. Vehicle-to-grid: A sociotechnical review. *Front. Energy Res.* **2017**, *5*, 9.
14. Thompson, A.W.; Perez, Y. Vehicle-to-Everything (V2X) energy services, value streams, and regulatory policy implications. *Energy Policy* **2020**, *137*, 111136. [\[CrossRef\]](#)
15. Sovacool, B.K.; Noel, L.; Axsen, J.; Kempton, W. The neglected social dimensions to a vehicle-to-grid (V2G) transition: A critical and systematic review. *Environ. Res. Lett.* **2018**, *13*, 013001. [\[CrossRef\]](#)

16. Clement-Nyns, K.; Haesen, E.; Driesen, J. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Trans. Power Syst.* **2010**, *25*, 371–380. [[CrossRef](#)]
17. Richardson, P.; Flynn, D.; Keane, A. Optimal charging of electric vehicles in low-voltage distribution systems. *IEEE Trans. Power Syst.* **2012**, *27*, 268–279. [[CrossRef](#)]
18. Muratori, M. Impact of uncoordinated plug-in electric vehicle charging on residential power demand. *Nat. Energy* **2018**, *3*, 193–201. [[CrossRef](#)]
19. Vagropoulos, S.I.; Bakirtzis, A.G. Optimal bidding strategy for electric vehicle aggregators in electricity markets. *IEEE Trans. Power Syst.* **2013**, *28*, 4031–4041. [[CrossRef](#)]
20. Camacho, E.F.; Bordons, C. *Model Predictive Control*, 2nd ed.; Springer: London, UK, 2013.
21. Tan, K.M.; Ramachandaramurthy, V.K.; Yong, J.Y. Integration of electric vehicles in smart grid: A review on vehicle to grid technologies and optimization techniques. *Renew. Sustain. Energy Rev.* **2016**, *53*, 720–732. [[CrossRef](#)]
22. Zheng, Y.; Niu, S.; Shang, Y.; Shao, Z.; Jian, L. Integrating plug-in electric vehicles into power grids: A comprehensive review on power interaction mode, scheduling methodology and mathematical foundation. *Renew. Sustain. Energy Rev.* **2019**, *112*, 424–439. [[CrossRef](#)]
23. Hu, J.; Morais, H.; Sousa, T.; Lind, M. Electric vehicle fleet management in smart grids: A review of services, optimization and control aspects. *Renew. Sustain. Energy Rev.* **2016**, *56*, 1207–1226. [[CrossRef](#)]
24. Mwasilu, F.; Justo, J.J.; Kim, E.K.; Do, T.D.; Jung, J.W. Electric vehicles and smart grid interaction: A review on vehicle to grid and renewable energy sources integration. *Renew. Sustain. Energy Rev.* **2014**, *34*, 501–516. [[CrossRef](#)]
25. García-Villalobos, J.; Zamora, I.; San Martín, J.I.; Asensio, F.J.; Aperribay, V. Plug-in electric vehicles in electric distribution networks: A review of smart charging approaches. *Renew. Sustain. Energy Rev.* **2014**, *38*, 717–731. [[CrossRef](#)]
26. Xu, B.; Oudalov, A.; Ulbig, A.; Andersson, G.; Kirschen, D.S. Modeling of lithium-ion battery degradation for cell life assessment. *IEEE Trans. Smart Grid* **2018**, *9*, 1131–1140. [[CrossRef](#)]
27. Sortomme, E.; El-Sharkawi, M.A. Optimal scheduling of vehicle-to-grid energy and ancillary services. *IEEE Trans. Smart Grid* **2012**, *3*, 351–359. [[CrossRef](#)]
28. Wang, D.; Coignard, J.; Zeng, T.; Zhang, C.; Saxena, S. Quantifying electric vehicle battery degradation from driving vs. vehicle-to-grid services. *J. Power Sources* **2016**, *332*, 193–203. [[CrossRef](#)]
29. Uddin, K.; Jackson, T.; Widanage, W.D.; Chouchelamane, G.; Jennings, P.A.; Marco, J. On the possibility of extending the lifetime of lithium-ion batteries through optimal V2G facilitated by an integrated vehicle and smart-grid system. *Energy* **2017**, *133*, 710–722. [[CrossRef](#)]
30. Bin Arif, M.A. *Helix EV Charging Dataset: Real Charging Sessions at the Newcastle Helix Site (Urban Sciences Building), Newcastle upon Tyne, 2021–2026*; School of Engineering, Newcastle University: Newcastle upon Tyne, UK, 2026.
31. Newcastle Urban Observatory. Open Environmental and Energy Data Platform, Newcastle upon Tyne. Available online: <https://newcastle.urbanobservatory.ac.uk> (accessed on 29 July 2026).
32. Hong, T.; Fan, S. Probabilistic electric load forecasting: A tutorial review. *Int. J. Forecast.* **2016**, *32*, 914–938. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.