games

MDPI

Article

Decentralized, Efficient, and Fair: Mean-Field Predictive Control for Bidirectional EV Coordination Under Uncertainty

Samuel M. Muhindo

Department of Electrical Engineering, Polytechnique Montréal, Montréal, QC H3T 1J4, Canada; samuel.muhindo-mugisho@polymtl.ca

Abstract

We propose a decentralized strategy for coordinating the bidirectional charging and discharging of battery electric vehicles (BEVs) in renewable-powered parking lots. The framework combines mean-field games (MFGs) and model predictive control (MPC) to address the coupled stochastic dynamics induced by uncertain renewable generation and random vehicle arrivals and departures. Solar and wind power fluctuations are modeled using autoregressive moving-average (ARMA) processes, while the time-varying vehicle population is represented through finite Poisson processes. The coordination problem is formulated as a large-scale game, where an aggregator designs individual cost functions to maximize available energy utilization while promoting fairness through near-equal states of charge (SOCs) at departure. Scalability is achieved through MFG theory, ensuring convergence and stability even under highly volatile generation and fluctuating agent populations. Numerical simulations validate the proposed strategy against two straightforward algorithms: capacity-ordered saturation allocation (COSA) and capacity-ordered fair allocation (COFA). These centralized approaches achieve high target fulfillment in static, low-intensity environments, where available energy accommodates a stable fleet without exceeding power limits. However, their efficacy degrades significantly in dynamic, high-intensity environments, where the interplay of volatile generation, continuous fleet turnover, and strict power constraints strains the system. In contrast, the proposed MFG-MPC framework provides a decentralized response that elegantly navigates the trade-offs between energy availability, demand stochasticity, and power limits. Ultimately, this approach ensures robust energy utilization while safeguarding vehicle equity, confirming its strong suitability for real-time deployment.

Keywords: autoregressive moving average; battery electric vehicles; charging; discharging; grid-to-vehicle; mean-field game; model predictive control; Nash equilibrium; Poisson process; renewable-powered parking lot; solar energy; vehicle-to-grid; wind energy

Academic Editors: Subhadip Ghosh and Arka Mukherjee

Received: 4 June 2026
Revised: 26 June 2026
Accepted: 30 June 2026
Published: 9 July 2026

1. Introduction

The daily use of battery electric vehicles (BEVs), particularly when powered by renewable energy sources, is widely recognized as a key lever for reducing greenhouse gas emissions in the transportation sector (Canada Energy Regulator, 2026; International Energy Agency, 2026). However, the large-scale deployment of BEVs also raises significant challenges for power systems, notably due to the additional electricity demand (Energy Institute, 2025). At the same time, the aggregate storage capacity embedded in BEV batteries offers an opportunity to mitigate the inherent variability of intermittent renewable energy sources such as solar and wind power (Nour et al., 2020).

A particularly promising use case arises in parking lots equipped with on-site renewable generation. Most BEVs commuting to workplaces remain parked for extended periods during daytime hours, which naturally coincide with peak solar availability. When connected to charging infrastructure in solar-powered parking facilities, BEV batteries can therefore act as distributed energy storage systems (Aghmadi & Mohammed, 2024), absorbing renewable generation that might otherwise be curtailed. Beyond individual charging services, aggregated BEV fleets can be operated as virtual power plants, delivering ancillary services and peak-shaving support to the grid during demand peaks when vehicles return home (Comi & Elnour, 2024; Wan et al., 2024; Zabetian-Hosseini et al., 2022).

Despite these opportunities, coordinating the bidirectional charging of a large population of BEVs remains a challenging control problem. First, renewable energy availability—solar irradiance and wind speed—is intrinsically stochastic and imperfectly predictable. Second, the population of BEVs connected to parking lots is inherently time-varying, with random arrivals/departures, and has heterogeneous batteries. These features create a high-dimensional, stochastic, and dynamic control problem, in which centralized optimization quickly becomes impractical as the number of vehicles increases (Ma et al., 2013; Tuchnitz et al., 2021; Xu et al., 2024).

Mean-field game (MFG) theory provides a powerful framework for addressing large-population control problems in which individual agents interact primarily through aggregate statistical quantities (Huang et al., 2006; Lasry & Lions, 2007). In the limit of infinitely many agents, MFG theory enables tractable decentralized control strategies with negligible loss of optimality. However, a fundamental difficulty arises (Cecchin & Pelino, 2019; Li et al., 2023) when applying standard MFG theory to BEV charging/discharging: while the population size may be large, it is not constant. Instead, the number of agents fluctuates throughout the day due to random arrivals and departures of vehicles, and these variations directly impact the average charging/discharging trajectory as well as the equitable exchange of renewable energy between the vehicles and the grid (Muhindo, 2024).

Building upon the inverse Nash MFG framework previously developed for deterministic solar-powered parking lots (Muhindo, 2024; Muhindo et al., 2021a), this paper proposes a decentralized charging/discharging algorithm capable of handling both renewable uncertainty and fluctuating vehicle populations. The parking-lot aggregator designs individual BEV cost functions that steer the collective behavior toward maximizing the absorption of locally available renewable energy while ensuring fairness among vehicles. Fairness is enforced by promoting reduced dispersion of SOCs at departure and by preserving the relative ordering of SOCs among vehicles that are simultaneously present.

To address uncertainty in renewable generation and vehicle population dynamics, model predictive control (MPC) is integrated into the MFG framework. Short-term forecasts of solar irradiance and wind speed are generated using autoregressive moving-average (ARMA) models, which are well suited for very short prediction horizons due to the strong temporal correlations in meteorological variables. Random arrivals and departures of BEVs are modeled using finite-population Poisson processes, allowing the prediction of expected vehicle counts over rolling horizons. The resulting MFG-MPC architecture enables the charging/discharging policy to adapt in real time to newly observed data while maintaining decentralized implementation at the BEV level (Cominesi et al., 2015; Novickij & Joós, 2019; Parisio et al., 2014; Qi et al., 2013). Unlike centralized strategies, the MFG-MPC approach allows participants to dynamically opt out of bidirectional energy exchanges. By coordinating the fleet through aggregate statistical quantities rather than top-down commands, users retain full authority to disconnect their vehicles or override policies—for instance, to prioritize battery health (Adegbohun et al., 2024) or accommodate unexpected travel—without disrupting collective grid stability.

A key contribution of this work is the introduction of a heuristic algorithm with provable convergence properties that extends the inverse Nash MFG formulation to systems with time-varying numbers of agents. This algorithm relies on appropriately constructed fictitious initial SOCs for anticipated arrivals and departures, allowing the MFG-based control law to be applied as if all vehicles were connected from the beginning of the prediction horizon. This approach preserves the desirable fairness and decentralization properties of the original MFG solution while significantly broadening its applicability to realistic operations.

The envisioned business model motivating this study involves the daily use of renewable-powered parking facilities by BEV owners—such as individuals working or studying in the surrounding area—through a subscription-based arrangement. Participants benefit from guaranteed access to charging and parking services, while the aggregator leverages the collective storage capacity of the fleet. In the operating regime considered herein, the supply–demand ratio is assumed to be less than one, such that not all vehicles reach full charge by the end of the day. Excess renewable energy, when available, may be stored locally and carried over to subsequent operating periods. With user consent, the aggregated fleet may further participate in vehicle-to-grid (V2G) schemes during evening peak demand, with revenues shared between the aggregator and vehicle owners (Comi & Elnour, 2024; Muhindo, 2024).

The main contributions of this paper are summarized as follows:

(i) An integrated stochastic framework modeling both supply and demand uncertainties. Solar and wind fluctuations are represented via autoregressive moving-average (ARMA) processes, while continuous BEV turnover is captured using finite-population Poisson processes.
(ii) Proof of solution existence for the inverse Nash mean-field game (MFG) formulation in both charging and discharging regimes. This ensures the well-posedness of the decentralized equilibrium for a fully heterogeneous fleet with diverse capacities, efficiencies, and initial states of charge.
(iii) A coupled MFG-MPC architecture that extends classical MFG theory to dynamic finite-population environments. By leveraging successive short rolling horizons, this approach guarantees convergence in fluctuating-agent systems while preserving decentralized operability.
(iv) Extensive numerical benchmarking against two straightforward algorithms, demonstrating the proposed strategy's superior balance of aggregate efficiency, robust energy utilization, and vehicle equity under realistic stochastic conditions.

The remainder of the paper is organized as follows. Section 2 develops the mean-field game formulation of the charging and discharging problem under perfect information and establishes the associated equilibrium properties. Section 3 presents the proposed MFG–MPC strategy, emphasizing its implementation under uncertainty and its convergence in dynamically fluctuating populations. Section 4 summarizes the main findings and outlines directions for future research. Appendix A presents the modeling and short-term forecasting framework for renewable generation, detailing the ARMA-based methodology adopted for solar and wind power. Appendix B introduces the stochastic modeling of BEV arrivals and departures through finite-population Poisson non-mixed and mixed processes.

2. MFG-Based Control for a Large Population of BEVs Under Perfect Information

In what follows, and in order to keep the presentation self-contained, we briefly recall several analytical elements from (Muhindo, 2024). In this section, power generation is assumed deterministic and all BEVs are assumed to be charged or discharged at the same time.

During daytime operation, BEVs interact with parking facilities functioning as local microgrids to absorb on-site renewable generation, such as solar and wind power. On the other hand, during nighttime operation, the vehicles reverse the power flow, discharging their stored energy to support the public electricity network. Despite this bidirectional energy exchange—drawing local renewable power by day and injecting power into the public network by night—the mathematical formulation of the coordination problem remains structurally identical. Consequently, throughout this analysis, the interfacing infrastructure in both operating regimes is collectively referred to as the grid.

The equations introduced below constitute a fundamental building block for the more general framework developed later in the paper, where renewable generation and BEV arrivals/departures are modeled as stochastic processes.

2.1. Battery Model

2.1.1. Charging and Discharging Regimes

We consider a population of n heterogeneous BEVs connected to the grid through bidirectional electric vehicle supply equipment (EVSE). All physical quantities are expressed in standard electrical engineering units.

For each BEV i, we define the following:

- t (in `h`), the time.
- $U_{i,t}$ (in `kW`), the power exchanged between the grid and the bidirectional charger, satisfying the physical constraint

$$0 \leq U_{i,t} \leq U_{\max}, \tag{1}$$

  where $U_{\max}$ (in `kW`) represents the maximum admissible power rate.
- $U^b_{i,t}$ (in `kW`), the effective power transferred to or from the battery. The distinction between $U_{i,t}$ and $U^b_{i,t}$ reflects the power conversion losses occurring in the charger. Specifically, $U_{i,t}$ corresponds to the power measured at the grid–charger interface, whereas $U^b_{i,t}$ represents the effective power flowing into or out of the battery cells.
- $\eta_i \in (0, 1]$, the charger efficiency (dimensionless).
- β_i (in `kWh`), the nominal energy capacity of the battery.
- $W_{i,t}$ (in `kWh`), the instantaneous stored battery energy.
- $x_{i,t}$ denotes the state of charge (SOC), defined as the normalized battery energy, which represents the fraction of available battery capacity.

In the present work, and for ease of interpretation in engineering applications, the state of charge is expressed in percentage units. Accordingly,

$$x_{i,t} = \frac{W_{i,t}}{\beta_i} \times 100, \tag{2}$$

so that $x_{i,t} \in [0, 100]$ represents the percentage of available battery capacity.

(a) Charging regime (grid-to-vehicle mode)

When the vehicle is charging, the battery energy evolves according to

$$\begin{aligned} \frac{dW_{i,t}}{dt} &= U^b_{i,t}, \\ &= \eta_i\, U_{i,t}, \end{aligned} \tag{3}$$

which yields the deterministic SOC dynamics in the charging regime:

$$dx_{i,t} = \frac{\eta_i}{\beta_i} U_{i,t}\, dt. \tag{4}$$

(b) Discharging regime (vehicle-to-grid mode)
When the vehicle injects energy back into the grid, the battery energy decreases as

$$\begin{aligned} \frac{dW_{i,t}}{dt} &= -U^b_{i,t}, \\ &= -\frac{1}{\eta_i} U_{i,t}, \end{aligned} \tag{5}$$

which yields the deterministic SOC dynamics in the discharging regime:

$$dx_{i,t} = -\frac{1}{\eta_i \beta_i} U_{i,t}\, dt. \tag{6}$$

2.1.2. Stochastic Battery Model

To account for modeling inaccuracies and small fluctuations in power delivery (Tamilselvi et al., 2021), we introduce a stochastic formulation of the state of charge (SOC) dynamics. In addition, the initial SOCs of vehicles when they connect to the parking facility are naturally heterogeneous and are therefore modeled as random variables whose distribution characterizes the fleet composition. For convenience of interpretation, both SOC $x_{i,t}$ and the conversion factor α_i, defined below, are expressed in percentage units.

Under these conventions, the stochastic SOC dynamics of BEV i are governed by

$$dx_{i,t} = \frac{\alpha_i}{\beta_i} U_{i,t}\, dt + \nu\, d\omega_{i,t}, \tag{7}$$

where

- α_i is the efficiency-adjusted conversion factor that maps the grid-interface power $U_{i,t}$ to the internal variation of the battery's SOC. To express the SOC as a percentage while rigorously accounting for the bidirectional converter's energy losses, α_i (in %) is defined as
$$\alpha_i = \begin{cases} \eta_i \times 100, & \text{during charging,} \\ -\dfrac{1}{\eta_i} \times 100, & \text{during discharging,} \end{cases} \tag{8}$$
- $\omega_{i,t}$ (in $\sqrt{h}$) denotes a standard Brownian motion, assumed to be independent across BEVs.
- $\nu > 0$ (in $\%/\sqrt{h}$) is the diffusion coefficient controlling the magnitude of stochastic fluctuations. The stochastic term $\nu\, d\omega_{i,t}$ captures small random deviations in battery energy evolution that may arise from power tracking errors, thermal effects, or unmodeled electrochemical dynamics.

The assumption of a large population of batteries becomes essential once stochastic effects are introduced. In this case, the empirical distribution of the states of charge converges toward its mathematical expectation by virtue of the law of large numbers, which enables a deterministic mean-field representation of the aggregate system. Nevertheless, owing to the linear structure of the battery dynamics and, as will be shown later, of the optimal control policy, the analysis remains exact for arbitrarily small populations when the stochastic term is absent. Finally, irrespective of whether diffusion is included in the dynamics, the initial states of charge $x_{i,0}$ are modeled as random variables whose distribution reflects the heterogeneous arrival conditions of vehicles at the parking facility.

2.1.3. Elemental Building Block of 1 kWh Batteries

We consider an elemental building block model to replace the n heterogeneous batteries with capacities β_i by n_β elementary batteries with a capacity of 1 kWh, i.e.,

$$n_\beta = \sum_{i=1}^{n} \beta_i. \tag{9}$$

We can then rewrite (7) as follows:

$$dx_{i,t} = \alpha_i \, u_{i,t} dt + \nu \, d\omega_{i,t}. \tag{10}$$

where $u_{i,t}$ is the capacity-normalized power rate in `1/h` (i.e., $u_{i,t} = \frac{U_{i,t}}{\beta_i}$). In conclusion, we use heterogeneous batteries as integer multiples of the elemental building block batteries of 1 kWh. In effect, once the charging or discharging control policy $u^*_{i,t}$ for the elemental battery building block is computed, the charging or discharging control policy $U^*_{i,t}$ for a BEV i will be a multiple of that associated with the building block, i.e.,

$$U^*_{i,t} = \beta_i \, u^*_{i,t}. \tag{11}$$

This results in one simple control policy $U^*_{i,t}$, in contrast to (Muhindo et al., 2021b), for the n heterogeneous BEVs.

2.2. Computing the Target Average SOC Trajectory

The energy available (or demand) is enforced over a finite horizon $[t_0, T]$. At t_0, the initial states of charge $x_{i,0}$, $i = 1, \dots, n$, of the BEVs are assumed known. To characterize the fleet's arrival distribution, we define the arithmetic mean μ_0 and standard deviation (STD) σ_0:

$$\mu_0 = \frac{1}{n} \sum_{i=1}^{n} x_{i,0}, \quad \sigma_0 = \sqrt{\frac{1}{n} \sum_{i=1}^{n} (x_{i,0} - \mu_0)^2}. \tag{12}$$

However, for a fleet with heterogeneous battery capacities, the global energy balance must be anchored by the true capacity-weighted initial average SOC, denoted $\overline{x}_0$:

$$\overline{x}_0 \equiv \mu_{0,\beta} = \frac{\sum_{i=1}^{n} \beta_i x_{i,0}}{n_\beta}. \tag{13}$$

The target average SOC trajectory $\overline{x}_t^{target}$, $\forall t \in [t_0, T]$, is thus initialized with this capacity-weighted aggregate and calculated as follows:

$$\begin{aligned} \frac{d\overline{x}_t^{target}}{dt} &= \frac{\overline{\alpha}}{n_\beta} U_{W_t}, \quad \overline{x}_{t_0}^{target} = \overline{x}_0. \\ \implies \overline{x}_t^{target} &= \overline{x}_0 + \frac{\overline{\alpha}}{n_\beta} \int_{t_0}^{t} U_{W_\tau} \, d\tau, \end{aligned} \tag{14}$$

where U_{W_t} is the assumed available renewable power.

We explain below the reverse engineering mechanism, where all n heterogeneous BEVs use the target average SOC trajectory $\overline{x}_t^{target}$ to calculate their optimal SOC trajectories $x^*_{i,t}$ such that the empirical per-BEV average optimal SOC $\overline{x}^*_t = \frac{1}{n} \sum_{i=1}^{n} x^*_{i,t}$ is equivalent to $\overline{x}_t^{target}$, $\forall t \in [t_0, T]$.

2.3. Establishing the Individual Battery Cost Function

In a game theory context (Kizilkale et al., 2019; Lenet et al., 2021; Muhindo, 2024; Muhindo et al., 2021b), the expected cumulative cost of charging or discharging a BEV i, $i = 1, \ldots, n$, is defined as

$$J_i(x_{i,0}) = \mathbb{E}\left[\int_{t_0}^{T} e^{-\delta t} \Big(\frac{r}{2} u_{i,t}^2 + q_t^y (x_{i,t} - y)^2 + q^{x_0} (x_{i,t} - x_{i,0})^2 \Big) dt \;\Big|\; x_{i,0} \right], \tag{15}$$

where

(i) $[t_0, T]$ is the control horizon, with $t \in [t_0, T]$, in `h`.
(ii) $\delta \geq 0$, in `1/h`, is a discount rate to guarantee the convergence of J in a context of extension to an infinite horizon ($T \to \infty$). Then, $\delta = 0$, for a finite horizon $[t_0, T]$.
(iii) u, in `1/h`, is the normalized control (knowing that the control $U = \beta\, u$).
(iv) r, in `h`, is the normalized control regularization parameter, common to all BEVs, which penalizes the charging or discharging rate.
(v) x, in `%`, is the state of charge.
(vi) y, in `%`, is the tracking state of charge, serving as a possible final destination for x. Therefore, $y = 100\%$ during charging, and $y = 0\%$ during discharging.
(vii) q^y, in $\mathtt{1/(\%^2h)}$, is the time-varying tracking weight, common to all BEVs, which penalizes any distance from y. Then, q^y is the pressure trajectory that would give priority to the less full BEVs. It would be calculated in a reverse engineering mechanism (called inverse Nash) (Lenet et al., 2021; Muhindo, 2024; Muhindo et al., 2021b) detailed below, the aim of which is to reduce the standard deviation of the states of charge of BEVs at the end of charging or discharging. It is q^y that creates the link between agents (here BEVs) in the game (i.e., as the mean field).
(viii) q^{x_0}, in $\mathtt{1/(\%^2h)}$, is the initial-state weighting parameter, common to all BEVs, which penalizes any distance from $x_{i,0}$. So, q^{x_0} preserves the order of the states of charge upon arrival, and so, combined with the effect of q^y, the emptiest BEVs upon arrival must not exceed those arriving fuller. We consequently obtain a fair and decentralized algorithm.

2.4. Obtaining the Agent's Best Response

Knowing (15), we begin by making the appropriate variable changes:

$$\begin{aligned} \tilde{x}_{i,t} &= (x_{i,t} - y) e^{-\frac{\delta t}{2}}, \; \tilde{u}_{i,t} = u_{i,t}\, e^{-\frac{\delta t}{2}}, \; \tilde{\nu}_t = \nu\, e^{-\frac{\delta t}{2}}. \\ \to d\tilde{x}_{i,t} &= (-\tfrac{\delta}{2} \tilde{x}_{i,t} + \alpha_i\, \tilde{u}_{i,t})\, dt + \tilde{\nu}_t\, d\omega_{i,t}. \end{aligned} \tag{16}$$

Following the standard procedure for analyzing linear quadratic optimal control problems (Anderson & Moore, 2007), we hypothesize the following quadratic value function for the optimal cost-to-go function V_i with arbitrary SOC $\tilde{x}_{i,t}$, under the assumption that the number of BEVs is large enough that one can treat the mean SOC trajectory as deterministic:

$$V_i(\tilde{x}_{i,t}) = \tilde{\pi}_{i,t}\, \tilde{x}_{i,t}^2 + \tilde{s}_{i,t}\, \tilde{x}_{i,t} + \tilde{\gamma}_{i,t}, \tag{17}$$

with $\tilde{\pi}_{i,t}$, $\tilde{s}_{i,t}$, and $\tilde{\gamma}_{i,t}$ being the Riccati, linear, and independent gains, respectively. Recall (15),

$$V_i(\tilde{x}_{i,t}) = \min_{\{\tilde{u}\}} \mathbb{E}\left[\int_t^T \Big(\frac{r}{2} \tilde{u}_{i,\tau}^2 + \frac{q_\tau^y}{2} \tilde{x}_{i,\tau}^2 + \frac{q^{x_0}}{2} (\tilde{x}_{i,\tau} - \tilde{x}_{i,0})^2 \Big) d\tau \;\Big|\; x_{i,t} \right], \tag{18}$$

where $\{\tilde{u}\}$ is a state feedback control strategy. The associated Hamilton–Jacobi–Bellman equation (Bagchi, 1993), knowing the system dynamics in (16), and the optimal cost-to-go function in (18), is as follows:

$$\frac{\partial V_i}{\partial t} + \min_{\{\tilde{u}\}}\Big\{\frac{r}{2}\tilde{u}_{i,t}^2 + \frac{q_t^y}{2}\tilde{x}_{i,t}^2 + \frac{q^{x_0}}{2}\big(\tilde{x}_{i,t} - \tilde{x}_{i,0}\big)^2 + \frac{\partial V_i}{\partial \tilde{x}}\big(\alpha_i\,\tilde{u}_{i,t} - \tfrac{\delta}{2}\tilde{x}_{i,t}\big) + \frac{\partial^2 V_i}{\partial \tilde{x}^2}\frac{\tilde{v}_t^2}{2}\Big\} = 0. \quad (19)$$

The first derivative of the Hamiltonian part of (19) with respect to $\tilde{u}_{i,t}$ is set to zero:

$$r\,\tilde{u}_{i,t} + \alpha_i\frac{\partial V_i}{\partial \tilde{x}} = 0. \quad (20)$$

The second derivative of the Hamiltonian part of (19) with respect to $\tilde{u}_{i,t}$ is r, so the resulting $\tilde{u}_{i,t}$ in (20) is optimal if

$$r > 0. \quad (21)$$

Then, the agent's best response is given by

$$\begin{aligned}\tilde{u}_{i,t}^* &= -\frac{\alpha_i}{r}\big(2\tilde{\pi}_{i,t}\,\tilde{x}_{i,t} + \tilde{s}_{i,t}\big),\\ u_{i,t}^* &= \frac{\alpha_i}{r}\big[\pi_{i,t}(y - x_{i,t}) - s_{i,t}\big],\end{aligned} \quad (22)$$

given that

$$\tilde{\pi}_{i,t} = \tfrac{1}{2}\pi_{i,t},\ \tilde{s}_{i,t} = s_{i,t}\,e^{-\frac{\delta t}{2}},\ \tilde{\gamma}_{i,t} = \tfrac{1}{2}\gamma_{i,t}\,e^{-\delta t}. \quad (23)$$

It only remains to obtain the parameters π and s of (22) by identification using (19):

$$\begin{aligned}&\frac{1}{2}\tilde{x}_{i,t}^2\frac{d\pi_{i,t}}{dt} + \tilde{x}_{i,t}\,e^{-\frac{\delta t}{2}}\big(\frac{ds_{i,t}}{dt} - \frac{\delta}{2}s_{i,t}\big) + \frac{1}{2}e^{-\delta t}\big(\frac{d\gamma_{i,t}}{dt} - \delta\gamma_{i,t}\big)\\ &= -\frac{r}{2}\big(-\frac{\alpha_i}{r}(\pi_{i,t}\tilde{x}_{i,t} + s_{i,t}e^{-\frac{\delta t}{2}})\big)^2 - \frac{q_t^y}{2}\tilde{x}_{i,t}^2 - \frac{q^{x_0}}{2}\big(\tilde{x}_{i,t} + (y - x_{i,0})e^{-\frac{\delta t}{2}}\big)^2\\ &- \alpha_i\big(-\frac{\alpha_i}{r}(\pi_{i,t}\tilde{x}_{i,t} + s_{i,t}e^{-\frac{\delta t}{2}})^2\big) + \frac{\delta}{2}\tilde{x}_{i,t}\big(\pi_{i,t}\tilde{x}_{i,t} + s_{i,t}e^{-\frac{\delta t}{2}}\big) - \frac{\tilde{v}_t^2}{2}\pi_{i,t}\\ &= \frac{\alpha_i^2}{2r}\big(\pi_{i,t}^2\tilde{x}_{i,t}^2 + 2\pi_{i,t}s_{i,t}\tilde{x}_{i,t}e^{-\frac{\delta t}{2}} + s_{i,t}^2e^{-\delta t}\big) + \frac{\delta}{2}\pi_{i,t}\tilde{x}_{i,t}^2 + \frac{\delta}{2}s_{i,t}\tilde{x}_{i,t}e^{-\frac{\delta t}{2}}\\ &- \frac{q_t^y}{2}\tilde{x}_{i,t}^2 - \frac{q^{x_0}}{2}\big(\tilde{x}_{i,t}^2 + 2(y - x_{i,0})\tilde{x}_{i,t}e^{-\frac{\delta t}{2}} + (y - x_{i,0})^2e^{-\delta t}\big) - \frac{(\nu e^{-\frac{\delta t}{2}})^2}{2}\pi_{i,t}\\ &= \frac{1}{2}\tilde{x}_{i,t}^2\big[\frac{\alpha_i^2}{r}\pi_{i,t}^2 + \delta\pi_{i,t} - q_t^y - q^{x_0}\big] + \tilde{x}_{i,t}e^{-\frac{\delta t}{2}}\big[\frac{\alpha_i^2}{r}\pi_{i,t}s_{i,t} + \frac{\delta}{2}s_{i,t}\\ &- q^{x_0}(y - x_{i,0})\big] + \frac{1}{2}e^{-\delta t}\big[\frac{\alpha_i^2}{r}s_{i,t}^2 - \nu^2\pi_{i,t} - q^{x_0}(y - x_{i,0})^2\big].\end{aligned} \quad (24)$$

Which, given that $\tilde{x}_{i,t}$ is arbitrary, yields

$$\begin{aligned}\frac{d\pi_{i,t}}{dt} &= \frac{\alpha_i^2}{r}\,\pi_{i,t}^2 + \delta\,\pi_{i,t} - q_t^y - q^{x_0},\\ \frac{ds_{i,t}}{dt} &= \frac{\alpha_i^2}{r}\,\pi_{i,t}\,s_{i,t} + \delta\,s_{i,t} - q^{x_0}(y - x_{i,0}),\\ \frac{d\gamma_{i,t}}{dt} &= \frac{\alpha_i^2}{r}\,s_{i,t}^2 + \delta\,\gamma_{i,t} - \nu^2\,\pi_{i,t} - q^{x_0}(y - x_{i,0})^2.\end{aligned} \quad (25)$$

with boundary conditions of π and s that are discussed below.

2.5. Obtaining the Pressure Trajectory by Inverse Nash Mechanism

The pressure trajectory q^y in the equation of π in (25) still remains unknown at this stage. According to MFG theory (Huang et al., 2007, 2006; Lasry & Lions, 2007), when all n BEVs apply their optimal control law $u^*_{i,t}$ over the control horizon $[t_0, T]$ for a sufficiently large n, by virtue of the law of large numbers, the empirical averages of trajectories

$$\{\overline{x}^*_t, \overline{u}^*_t, \overline{s}_t\} = \frac{1}{n}\sum_{i=1}^{n}\{x^*_{i,t}, u^*_{i,t}, s_{i,t}\}, \tag{26}$$

are approximated by their deterministic mathematical expectations

$$\mathbb{E}\big[\{x^*_{i,t}, u^*_{i,t}, s_{i,t}\}\big],\ i = 1, \ldots, n, \tag{27}$$

Recalling (10), (22), (25), we have

$$\begin{aligned}
\frac{d\overline{x}^*_t}{dt} &= \frac{\overline{\alpha}^2}{r}\big[\overline{\pi}_t(y - \overline{x}^*_t) - \overline{s}_t\big],\\
\frac{d\overline{\pi}_t}{dt} &= \frac{\overline{\alpha}^2}{r}\overline{\pi}_t^2 + \delta\overline{\pi}_t - q^y_t - q^{x_0},\\
\frac{d\overline{s}_t}{dt} &= \frac{\overline{\alpha}^2}{r}\overline{s}_t\overline{\pi}_t + \delta\overline{s}_t - q^{x_0}(y - \overline{x}_0).
\end{aligned} \tag{28}$$

The pressure trajectory q^y defined in (15) is reverse-engineered so that the target average SOC trajectory $\overline{x}^{target}_t$ corresponds to full absorption of the available assumed deterministic energy, i.e.,

$$\overline{x}^*_t \equiv \overline{x}^{target}_t. \tag{29}$$

2.6. Computing the Boundary Conditions

Combining the first equation of (28) with (29) yields

$$\overline{\pi}_t = \frac{1}{y - \overline{x}^{target}_t}\big(\overline{s}_t + \frac{r}{\overline{\alpha}^2}\frac{d\overline{x}^{target}_t}{dt}\big). \tag{30}$$

With (30) in the last equation of (28), we obtain the following equation:

$$\frac{d\overline{s}_t}{dt} = \frac{\overline{s}_t^2\overline{\alpha}^2}{r(y - \overline{x}^{target}_t)} + \overline{s}_t\big[\delta + \frac{\frac{d\overline{x}^{target}_t}{dt}}{y - \overline{x}^{target}_t}\big] - q^{x_0}(y - \overline{x}_0). \tag{31}$$

To solve this equation, it is necessary to know $\overline{s}_T$ since the solution of (31) is obtained backwards. We consider a control horizon T sufficiently long for energy to completely subside at T, i.e.,

$$U_{W_T} = 0. \tag{32}$$

Recalling (14), (28), (29), and (32), and imposing that $\overline{\pi}_t$ and $\overline{s}_t$ settle at time T, we have

$$\begin{aligned}
0 &= \frac{\overline{\alpha}^2}{r}\big[\overline{\pi}_T(y - \overline{x}^{target}_T) - \overline{s}_T\big],\\
0 &= \frac{\overline{\alpha}^2}{r}\overline{\pi}_T^2 + \delta\,\overline{\pi}_T - q^y_T - q^{x_0},\\
0 &= \frac{\overline{\alpha}^2}{r}\overline{\pi}_T\overline{s}_T + \delta\,\overline{s}_T - q^{x_0}(y - \overline{x}_0).
\end{aligned} \tag{33}$$

With appropriate rearrangements,

$$\begin{aligned} \bar{s}_T &= \overline{\pi}_T(y - \overline{x}_T^{target}), \\ (q^{x_0} + q_T^y)\,\bar{s}_T &= \overline{\pi}_T\, q^{x_0}(y - \overline{x}_0), \\ \frac{\overline{\alpha}^2}{r}\overline{\pi}_T^2 + \delta\,\overline{\pi}_T &= (q^{x_0} + q_T^y), \end{aligned} \tag{34}$$

we obtain the following boundary conditions:

$$\begin{aligned} q_T^y &= q^{x_0}\Big(\frac{\overline{x}_T^{target} - \overline{x}_0}{y - \overline{x}_T^{target}}\Big), \\ \overline{\pi}_T &= \frac{-\delta \pm \sqrt{\delta^2 + 4\big[\frac{\overline{\alpha}^2}{r}(q^{x_0} + q_T^y)\big]}}{2\big(\frac{\overline{\alpha}^2}{r}\big)}, \\ \bar{s}_T &= \overline{\pi}_T(y - \overline{x}_T^{target}),\ \text{or}\ \bar{s}_T = \frac{\overline{\pi}_T\, q^{x_0}(y - \overline{x}_0)}{q^{x_0} + q_T^y}. \end{aligned} \tag{35}$$

2.7. Guaranteeing the Existence of a Solution to the MFG Inverse Nash Problem

There will always be a solution to the inverse Nash problem if we can demonstrate that the solution of (31) exists, $\forall t$. Let us rewrite the equation in (31) as follows:

$$-\frac{d\bar{s}_t}{dt} = q^{x_0}(y - \overline{x}_0) + \bar{s}_t\Big[\frac{\frac{d\overline{x}_t^{target}}{dt}}{\overline{x}_t^{target} - y} - \delta\Big] - \frac{(\overline{\alpha}\,\bar{s}_t)^2}{r(y - \overline{x}_t^{target})}. \tag{36}$$

In the grid-to-vehicle charging problem, where we have $\overline{x}_0 \le \overline{x}_t^{target} < y$, $\forall t$, the following conditions are sufficient to have a bounded solution of $\bar{s}_t$ (which is, from (22), a monotonically decreasing function) (Jacobson, 1970).

$$q^{x_0}(y - \overline{x}_0) > 0,\ \frac{1}{r(y - \overline{x}_t^{target})} > 0,\ \bar{s}_T > 0. \tag{37}$$

In the vehicle-to-grid discharging problem, where we have $y < \overline{x}_t^{target} \le \overline{x}_0$, $\forall t$, the following conditions are sufficient to have a bounded solution of $\bar{s}_t$ (which is, from (22), a monotonically increasing function) (Jacobson, 1970).

$$-q^{x_0}(y - \overline{x}_0) > 0,\ -\frac{1}{r(y - \overline{x}_t^{target})} > 0,\ -\bar{s}_T > 0. \tag{38}$$

Since $r > 0$ to have $u_{i,t}^*$ ((19)–(21)), the conditions of (37) or (38) are respected if we use the following in (35):

$$q^{x_0} > 0,\ \pi_T > 0. \tag{39}$$

2.8. Computing the Pressure Trajectory

Knowing $q^{x_0} > 0, \pi_T > 0$, and $\bar{s}_T$ in (35), we solve $\bar{s}_t$ in (31) backwards, and we calculate the Riccati gain $\overline{\pi}_t$ in (30). The pressure trajectory is then computed in the second equation of (28):

$$q_t^y = \frac{\overline{\alpha}^2}{r}\overline{\pi}_t^2 + \delta\,\overline{\pi}_t - \frac{d\pi_t}{dt} - q^{x_0}, \tag{40}$$

where $\frac{d\overline{\pi}_t}{dt}$ is calculated given $\overline{\pi}_t$ by using the mean-value theorem.

2.9. Computing the Agent's Best Response

We solve $\pi_{i,t}$ and $s_{i,t}$ of (25) backwards, knowing that their boundary conditions, $\pi_{i,T}$ and $s_{i,T}$ for each BEV i, from (35), are given by

$$\begin{aligned} \pi_{i,T} &= \frac{-\delta + \sqrt{\delta^2 + 4\left[\frac{\alpha_i^2}{r}(q^{x_0} + q_T^y)\right]}}{2\left(\frac{\alpha_i^2}{r}\right)}, \\ s_{i,T} &= \frac{\overline{\pi}_T\, q^{x_0}(y - x_{i,0})}{q^{x_0} + q_T^y}. \end{aligned} \tag{41}$$

Finally, we have $\pi_{i,t}$ and $s_{i,t}$ that will give us the optimal control law in (11) and (22).

2.10. Different Algorithms for Charging or Discharging Heterogeneous BEVs

2.10.1. Mean-Field Game (MFG) Inverse Nash Algorithm

The MFG inverse Nash algorithm (Algorithm 1) is a decentralized optimal control framework derived from a mean-field game formulation. The aggregator first computes a target average SOC trajectory $\overline{x}_t^{\text{target}}$ based on the available energy and number of BEVs, then solves a backward Riccati system to determine the aggregate shadow price of energy (trajectories $\overline{s}_t$ and $\overline{\pi}_t$). This induces a time-varying incentive signal (pressure trajectory q^y). Each BEV owner then solves their own backward Riccati system and implements a closed-loop optimal feedback control. The result is a dynamically consistent Nash equilibrium where heterogeneous batteries respond optimally while collectively tracking a prescribed mean trajectory $\overline{x}_t^{\text{target}}$.

Algorithm 1 MFG inverse Nash under perfect information.

Require: Horizon $[t_0, T]$, step dt, available energy, heterogeneous BEVs $(x_{i,0}, \beta_i, \alpha_i)$, maximum rate $U_{\max}$, and parameters $y, \nu, q^{x_0}, r, \delta$

1: Compute the target average SOC trajectory $\overline{x}_t^{\text{target}}$ (14)
2: Compute the boundary conditions q_T^y, $\overline{\pi}_T, \overline{s}_T$ (35)
3: Solve backward aggregate Riccati-type equations for $\overline{s}_t$ (31), and $\overline{\pi}_t$ (30)
4: Compute the pressure trajectory q^y (40)
5: **for** $i = 1$ to n **do**
6: Solve backward individual Riccati type-equations for $\pi_{i,t}$ and $s_{i,t}$ (25)
7: Compute unconstrained optimal control: $U_{i,t}^* = \dfrac{\alpha_i}{r}\big(\pi_{i,t}(y - x_{i,t}) - s_{i,t}\big)\beta_i$
8: Enforce charger constraint: $U_{i,t}^* \leftarrow \min\{U_{i,t}^*, U_{\max}\}$
9: Update SOC: $dx_{i,t} = \dfrac{\alpha_i}{\beta_i} U_{i,t}^*\, dt + \nu\, d\omega_{i,t}$
10: **end for**

2.10.2. Capacity-Ordered Algorithms

We consider the following straightforward algorithms (Muhindo et al., 2021a, 2021b) based on engineering intuition for comparison purposes with the MFG inverse Nash algorithm. One should note that a naive aggregator-based centralized algorithm, where at all times the all BEVs present are charged or discharged in proportion to their SOC difference with respect to saturation, is implicitly carried out in the MFG algorithm, albeit in a decentralized way, thanks to our design of the cost function in (15).

- Capacity-ordered saturation allocation (COSA): Energy is allocated sequentially according to a strict saturation rule. Each vehicle is fully charged (or fully discharged) before proceeding to the next one until the global energy budget is exhausted, while enforcing the per-charger power constraint $U_{\max}$. This policy maximizes instantaneous utilization of the available power but is structurally unfair: some vehicles are prioritized while others may receive no allocation when the energy budget is limited. In charging mode, vehicles are sorted in increasing order of battery capacity so that smaller batteries are saturated first. In discharging mode, vehicles are sorted in decreasing order of capacity so that larger batteries are depleted first.

- Capacity-ordered fair allocation (COFA): COFA can be interpreted as a refinement of COSA based on a capacity-aware equalization principle. Vehicles are ordered consistently with COSA (small-capacity first in charging mode, large-capacity first in discharging mode), but instead of sequential saturation, a common target SOC is determined so that the total allocated energy matches the available budget, subject to the per-charger power constraint $U_{\max}$. The resulting algorithm resembles a water-filling mechanism: it guarantees aggregate feasibility while promoting a proportional form of fairness across heterogeneous batteries.

2.10.3. Implementation Assumptions for the Charging and Discharging Algorithms

We have the following implementation assumptions:

(i) There exists a two-way communication infrastructure to coordinate BEV charging or discharging, where we assume that we know power generation, duration of the control horizon, and number of BEVs present.
(ii) The BEVs' owners communicate their battery energy over their usable capacity upon arrival (e.g., $9\,\texttt{kWh}/60\,\texttt{kWh} = 15\%$ SOC).
(iii) BEVs' owners use bidirectional chargers that are equipped with microprocessors, allowing them to implement and compute Algorithm 1 for charging or discharging their batteries, with a maximal power rate $U_{\max}$.

2.11. Experimental Platform and Data

We consider 10 distinct battery capacity classes based on realistic data (Electric Vehicle Database, 2026). Battery capacity classes are represented by color-coded groups corresponding to small, medium, and large capacities, respectively. Within each group, increasing tone intensity (from light to dark) signifies a progressive increase in battery capacity. Each BEV i, $i = 1, \ldots, n$, is equipped with a bidirectional charger delivering up to $U_{\max} = 19.2\,\text{kW}$, representing a standard AC interface. Each charger efficiency $\eta_i \in [0.8, 0.95]$ is a classic capacity-dependent model through its associated nominal C-rate (defined as $U_{\max}/\beta_i$). At arrival, the initial states of charge (SOCs) and charger efficiencies are sampled from Gaussian distributions with finite means and standard deviations. The other parameters are set as follows: $dt = \frac{1}{60}\,\text{h}$, $\nu = 0.01\%(\sqrt{dt})^{-1}(\sqrt{\text{h}})^{-1}, \delta = 0$, $r = 1\,\text{h}$, and $q^{x_0} = 1\,(\%^2\text{h})^{-1}$. Simulations are conducted in Python 3.12.

2.12. MFG-Based Charging and Discharging Results

2.12.1. Data and Baseline Results in the Deterministic Case

We first present the data for the deterministic charging and discharging regimes, as illustrated in Figures 1 and 2, considering the same population of BEVs, where $\mu_0^{\text{discharging}} \approx \mu_T^{\text{charging}} - 5\%$, and $\sigma_0^{\text{discharging}} \approx \sigma_T^{\text{charging}} + 2\%$ (Muhindo, 2024). The initial SOCs are color-coded by capacity class (red for small, orange for medium, and green for large), with increasing color darkness reflecting a larger battery capacity within each respective class. The baseline results are presented in Figures 3 and 4.

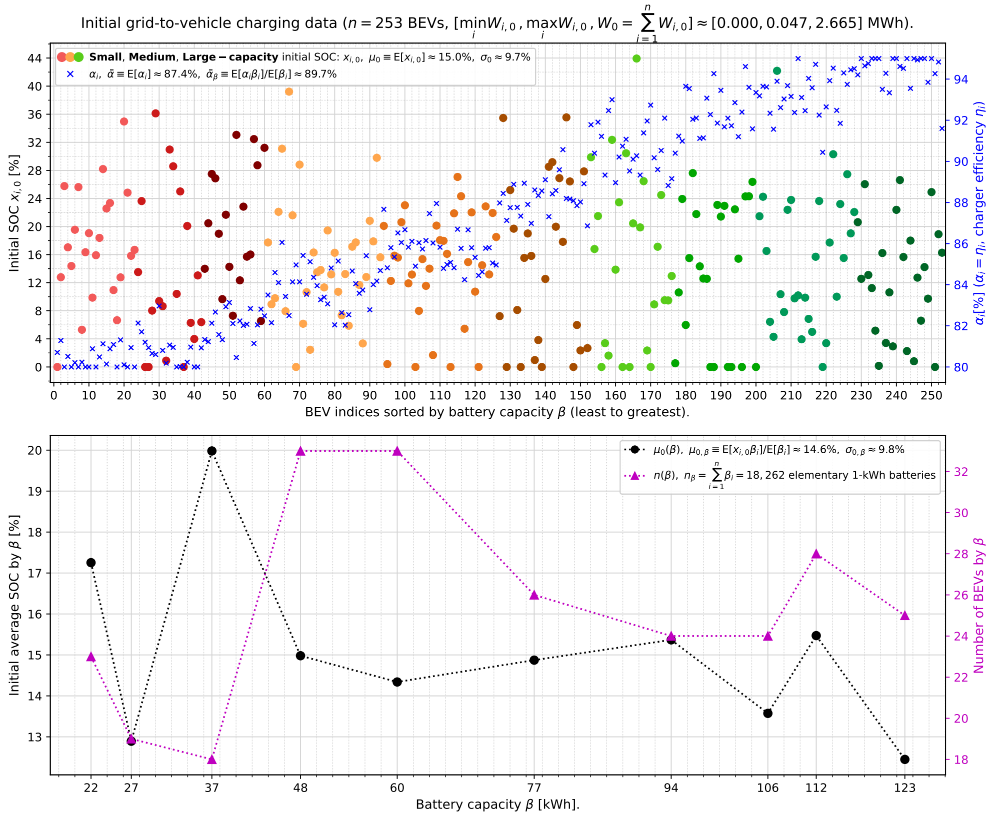

Figure 1. Grid-to-vehicle charging data: (**top**)— 253 BEVs' initial states of charge (SOCs) $x_{i,0}$ and charging efficiencies α_i; (**bottom**)—initial average SOC and number of BEVs, per battery capacity.

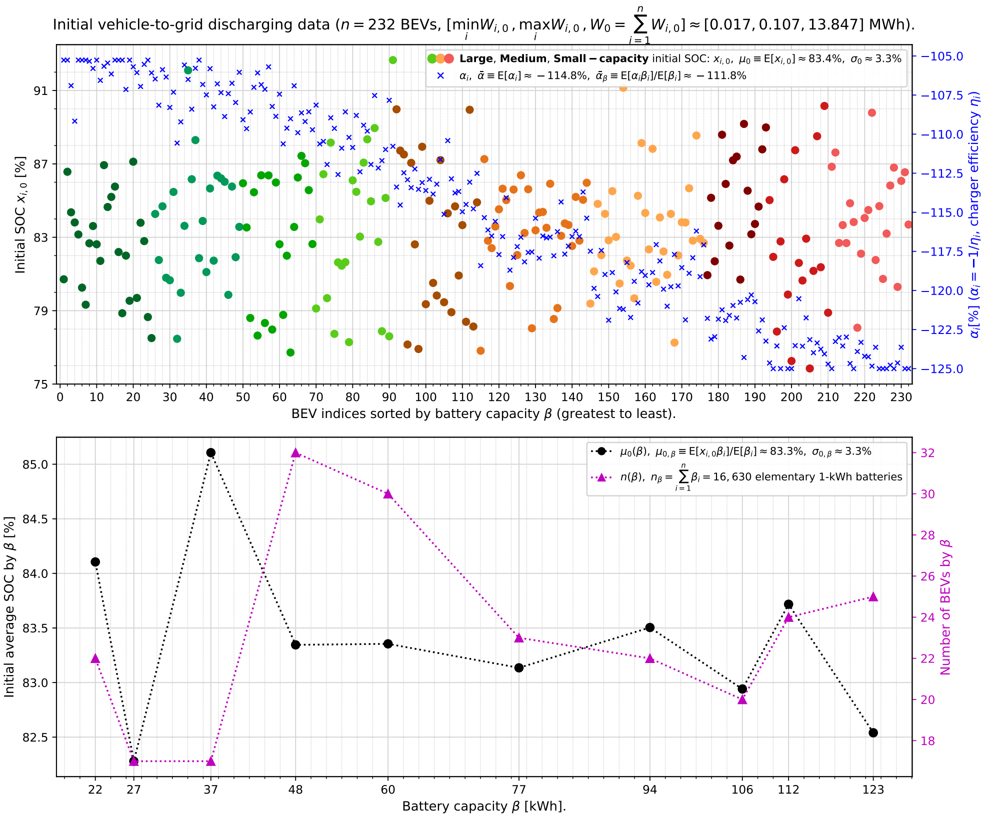

Figure 2. Vehicle-to-grid discharging data: (**top**)—232 BEVs' initial states of charge (SOCs) $x_{i,0}$ and discharging efficiencies α_i; (**bottom**)—initial average SOC and number of BEVs, per battery capacity.

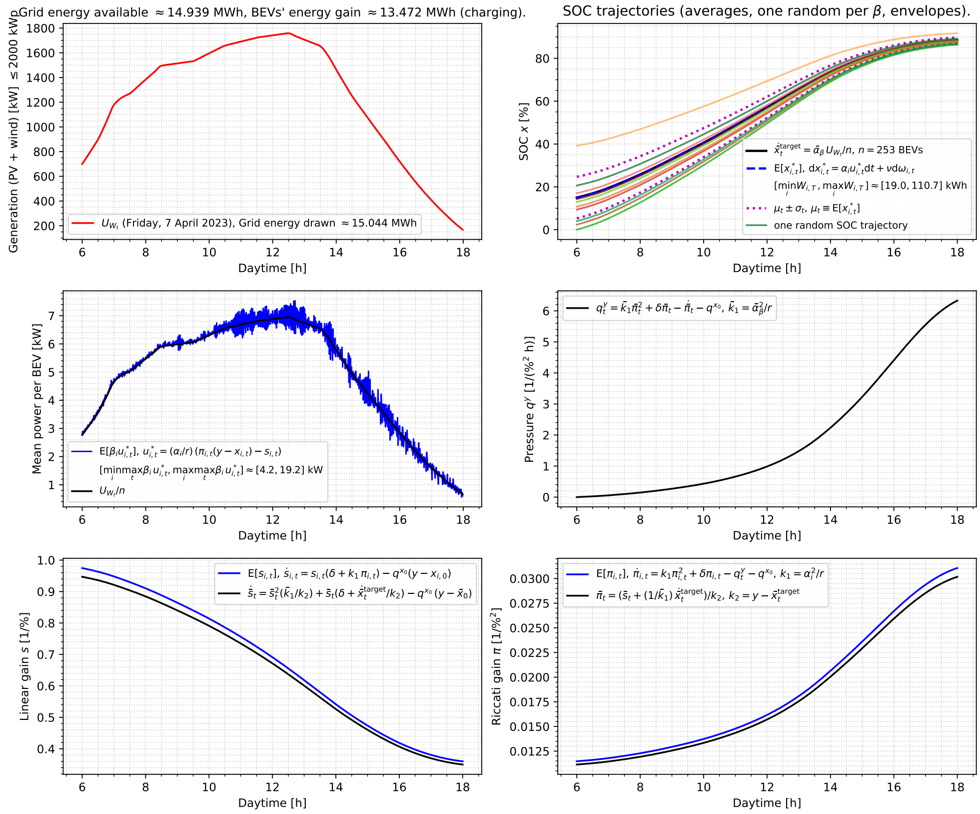

Figure 3. Grid-to-vehicle MFG-based charging results: (**upper-left**)—aggregate power (generation) U_{W_t}; (**upper-right**)—SOC trajectories x_t; (**lower-left**)—mean power per BEV $\overline{U}_t$; (**lower-right**)—pressure trajectory q_t^y; (**bottom-left**)—linear gain s_t; (**bottom right**)—Riccati gain π_t.

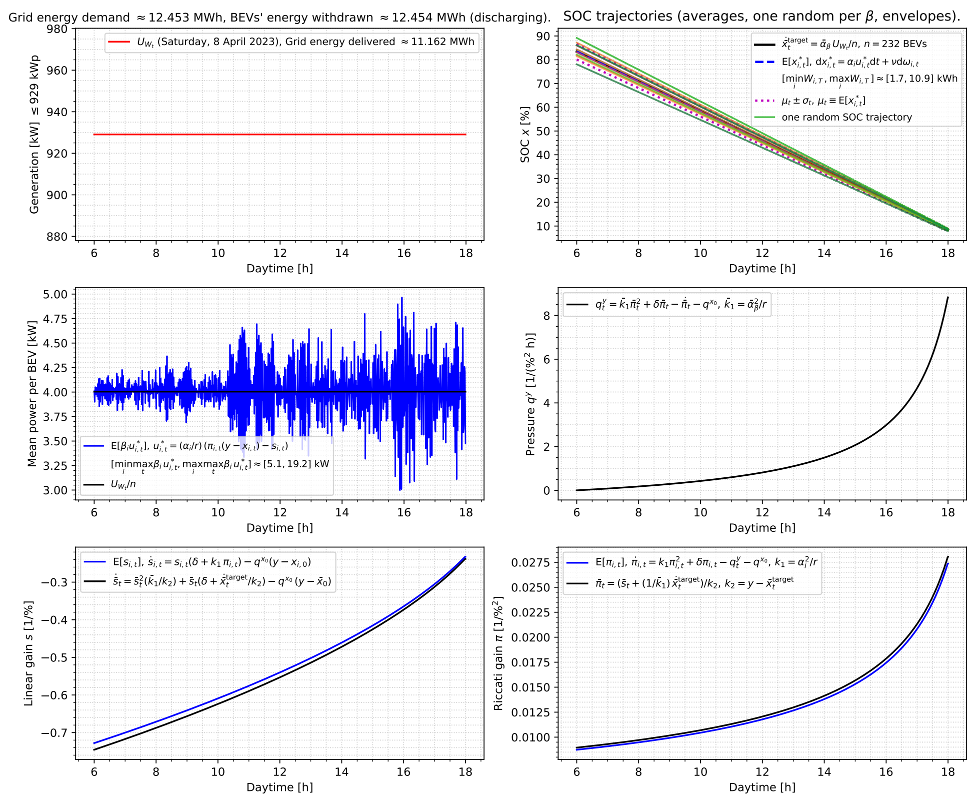

Figure 4. Vehicle-to-grid MFG-based discharging results: (**upper-left**)—aggregate power (generation) U_{W_t}; (**upper-right**)—SOC trajectories x_t; (**lower-left**)—mean power per BEV $\overline{U}_t$; (**lower-right**)—pressure trajectory q_t^y; (**bottom-left**)—linear gain s_t; (**bottom right**)—Riccati gain π_t.

Figure 3 illustrates the grid-to-vehicle (G2V) charging regime. First, the 253 BEVs exhibit a net energy gain of 13,472 kWh, whereas the total energy drawn from the grid amounts to 15,044 kWh. The resulting gap reflects the cumulative conversion losses induced by the non-ideal charger efficiencies of the vehicles. Interestingly, the stochastic diffusion inherent to the individual BEV dynamics can cause the total energy drawn to slightly overshoot the available grid energy (14,939 kWh); nevertheless, this theoretical excess is systematically capped and corrected by the physical limits of the chargers in practice. The optimal mean SOC trajectory clearly converges toward the target mean SOC trajectory, while the dispersion of SOC across BEVs narrows significantly by departure time. We also illustrate, in colored curves, 10 random sample SOC trajectories corresponding to vehicles of different battery capacities but only one representative trajectory is listed in the legend for visual clarity. The optimal mean power dynamically tracks the highly variable, uncoordinated grid power profile—which represents the combined solar and wind generation acquired at the site—while strictly respecting the maximum charging rate of 19.2 kW. It is crucial here to distinguish the actual optimal mean trajectory from the theoretical baseline, which simply represents the uniform per-vehicle allocation of the grid's power. The observed fluctuations in the mean trajectories stem from diffusion noise and a structural discrepancy: aggregate trajectories are computed using the capacity-weighted mean of the charger efficiencies, whereas local trajectories rely on their unweighted expectation. The mathematical structure of the MFG inverse Nash elegantly dictates the evolution of the control gains. The pressure q^y and the Riccati gain π exhibit monotonic growth, reflecting the intensifying urgency to strictly meet the terminal SOC constraints as the departure time approaches. In contrast, the feedforward gain s decreases monotonically in this charging mode; as the fleet naturally accumulates energy and approaches the reference trajectory y, the need for predictive forward guidance diminishes.

Figure 4 illustrates the vehicle-to-grid (V2G) discharging regime. Here, the grid receives a net energy injection of 11,162 kWh to help meet a demand of 12,453 kWh, which represents 90% of the fleet's total available energy. Meanwhile, the energy withdrawn from the 232 BEVs reaches 12,454 kWh. Furthermore, any marginal excess in the aggregate energy extraction—arising from the controller's need to overcome stochastic diffusion noise—is strictly capped by the physical hardware limits of the chargers. The fleet's optimal mean SOC trajectory is steered steadily downward toward the target average SOC, exhibiting a marked contraction of dispersion as vehicles approach their departure times. The optimal mean power dynamically tracks the per-vehicle share of the constant 929 kW aggregate discharging demand, while strictly respecting the maximum rate of 19.2 kW. As observed in the charging regime, the actual optimal trajectory fluctuates around the theoretical baseline—which, in this case, is perfectly flat due to the constant nature of the demand. The fundamental differences between the G2V and V2G regimes are elegantly captured by the endogenous adjustment of the control gains, despite sharing the exact same backward algebraic structure. In the discharging regime, both the Riccati gain π and the feedforward gain s grow monotonically. While π continues to penalize terminal deviations, the magnitude of s must increase to sustain the massive, continuous power extraction required by the grid right up to the final moments. Furthermore, the algebraic sign of s inherently adapts to the energy flow: it remains positive during charging to pull the SOC upward, and turns negative during discharging to force the SOC downward. Finally, the pressure q^y and the Riccati gain π share the same structural positivity and monotonic growth across both regimes, underscoring the universality of the MFG Inverse Nash formulation.

2.12.2. Comparison of Charging and Discharging Algorithms in the Deterministic Case

Figures 5 and 6 compare the three algorithms—mean-field game (MFG), capacity-ordered fair allocation (COFA), and capacity-ordered saturation allocation (COSA)—under heterogeneous BEVs in charging and discharging regimes, with a maximum rate of 19.2 kW.

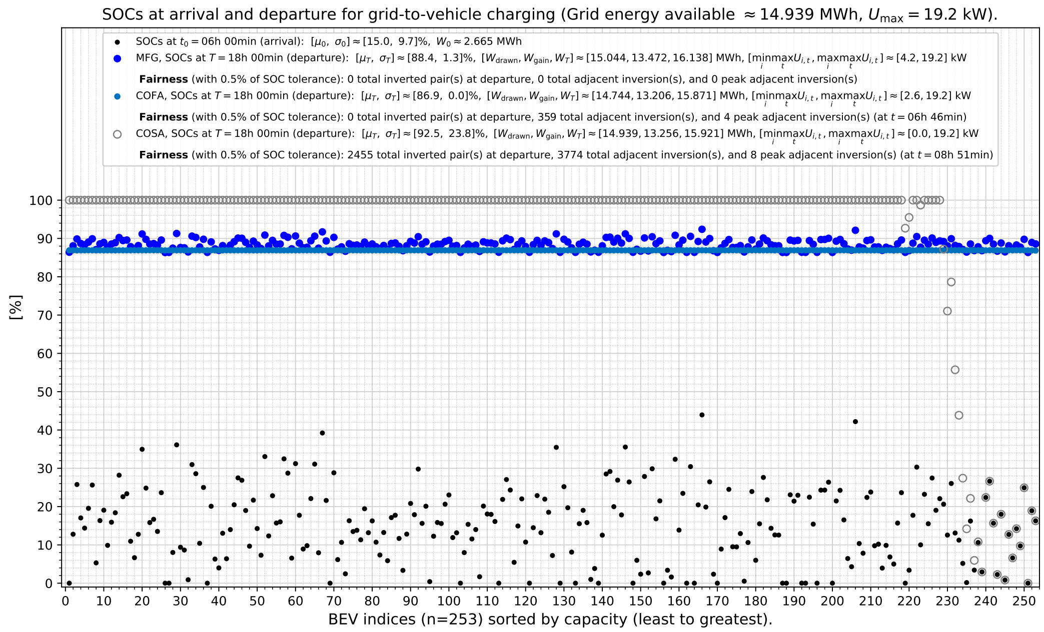

Figure 5. Grid-to-vehicle static charging algorithms: SOCs at arrival and departure.

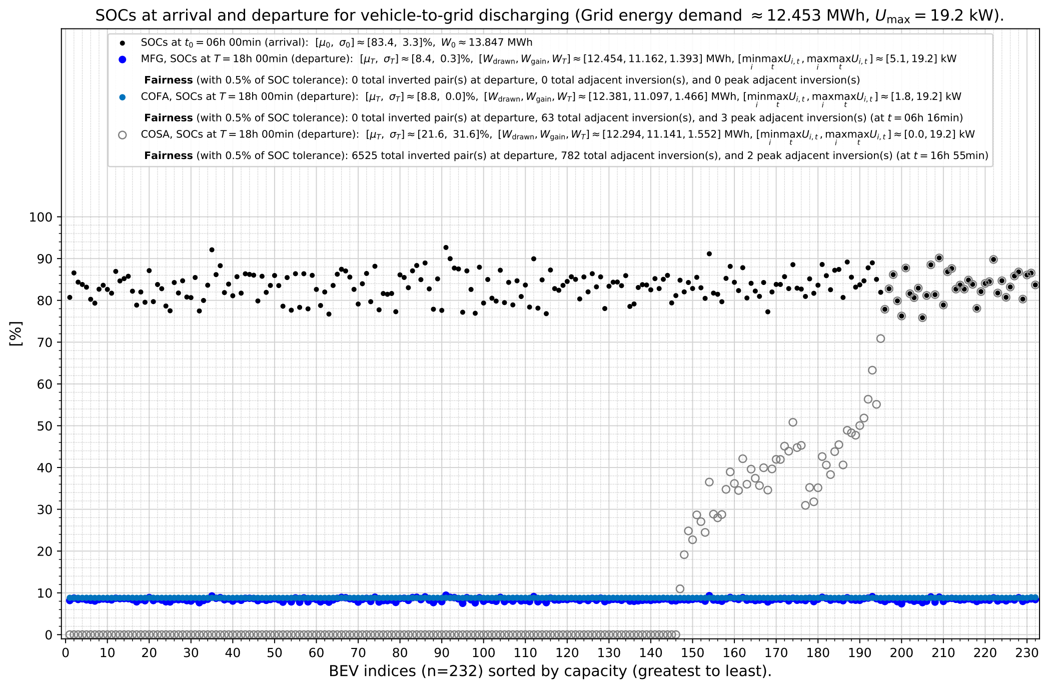

Figure 6. Vehicle-to-grid static discharging algorithms: SOCs at arrival and departure.

In the charging regime (Figure 5), the fundamental differences in allocation strategies become highly apparent. By design, the rule-based COSA algorithm allocates the available grid energy sequentially by strictly prioritizing small-capacity batteries first. While it maximizes total energy transfer, it leaves many higher-capacity BEVs completely uncharged, resulting in a severely stratified and unfair departure structure ($\mu_T \approx 92.5\%$, $\sigma_T \approx 23.8\%$). The COFA algorithm shares this sequential allocation philosophy but modifies its objective: it attempts to lift all vehicles toward a common, aggregate mean SOC dictated by the total available energy. Although egalitarian in principle—achieving zero variance ($\mu_T \approx 86.9\%$, $\sigma_T \approx 0.0\%$)—this strict uniformity proves mechanically inefficient. When the aggregate energy requirement is large, attempting to simultaneously drive all batteries to the exact same relative level creates severe bottlenecks under the maximum charging rate of 19.2 kW. Consequently, COFA becomes power-constrained and fails to extract the maximum available energy (drawing only 14,744 kWh). In contrast, the MFG approach distributes energy through a decentralized mean-field equilibrium that naturally internalizes the heterogeneity in both battery capacities and charger efficiencies. Drawing the grid energy available, it elegantly achieves a tight, dynamically consistent clustering of departure SOCs ($\mu_T \approx 88.4\%$, $\sigma_T \approx 1.3\%$) without hitting the artificial power bottlenecks induced by absolute uniformity. Importantly, MFG guarantees relative fairness by inherently preserving the initial SOC ordering of the fleet at departure.

In the discharging regime (Figure 6), these structural behaviors are fundamentally mirrored. To meet the grid's energy demand, COSA reverses its saturation logic, prioritizing the depletion of large-capacity BEVs first. This again leads to a highly polarized departure distribution ($\mu_T \approx 21.6\%$, $\sigma_T \approx 31.6\%$), disproportionately penalizing vehicles with larger batteries while shielding others. COFA, pursuing a uniform target mean SOC ($\mu_T \approx 8.8\%$, $\sigma_T \approx 0.0\%$), performs slightly better in this scenario (drawing 12,381 kWh) since the required energy exchange does not trigger the hardware rate limits as severely. Nevertheless, MFG remains superior: it ensures controlled, robust energy extraction to perfectly meet the grid's demand (12,454 kWh drawn), while maintaining an extraordinarily tight clustering of departure SOCs ($\mu_T \approx 8.4\%$, $\sigma_T \approx 0.3\%$) that, once again, perfectly preserves the initial SOC hierarchy of the fleet.

Conceptually, this comparative analysis reveals the profound advantages of the MFG framework over deterministic, rule-based algorithms. Both COSA and COFA rely on rigid, centralized coordination that requires explicit sorting of vehicle capacities. COSA maximizes energy utilization but entirely sacrifices fairness, whereas COFA pursues absolute fairness (zero variance) at the cost of operational efficiency, creating unnecessary power bottlenecks. MFG operates at the optimal intersection: relying on a decentralized architecture where each BEV responds optimally to a broadcast mass trajectory. In this way, MFG maximizes energy utilization like COSA, dynamically reduces variance like COFA, yet fundamentally outperforms both. It endogenously adapts to physical heterogeneities without centralized sorting, ensures massive scalability, and uniquely preserves the initial SOC ordering—offering a mathematically rigorous guarantee of fairness that neither centralized algorithm can naturally provide.

3. MFG-MPC Algorithm with Fluctuating BEVs and Stochastic Power Generation

A fully rigorous treatment of stochastic BEV arrivals/departures and renewable generation would require embedding a common-noise structure into the mean-field formulation. Renewable production acts as a common disturbance and therefore induces correlation across all SOC trajectories (Ahuja, 2015). Such an extension, although theoretically appealing, would substantially increase the analytical and computational complexity.

Instead, we adopt a model predictive control (MPC) approximation. Over short prediction windows, all stochastic inputs are replaced by conditional forecasts and treated as deterministic. The MFG equilibrium is then solved repeatedly in a receding-horizon fashion under quasi-perfect information.

MPC is a standard receding-horizon methodology in control engineering (Novickij & Joós, 2019). Coupled with the MFG load-allocation structure, it yields a computationally tractable and operationally implementable framework for BEV charging and discharging under uncertainty. Renewable forecasts rely on ARMA predictors, while BEV arrivals and departures are modeled via finite-population Poisson increments.

3.1. Adapting Mean-Field Games to Fleet Fluctuations and Stochastic Power Generation: A Fairness-Preserving Approach

Classical mean-field game (MFG) theory inherently relies on an asymptotic regime that strictly assumes an infinite and constant number of agents. In practical parking-lot operations, however, the population of connected BEVs fluctuates continuously, which directly distorts the aggregate per-vehicle target trajectory. To overcome this theoretical limitation, the proposed MFG-MPC architecture operates over successive 15-min receding horizons. At the beginning of each prediction window, all stochastic forecasts are refreshed, and a short-horizon MFG problem is formulated and solved utilizing the latest predicted trajectories for both the fleet composition and available power.

3.1.1. Stochastic Charging Regime

In charging mode, the aggregator anticipates the average SOC trajectory of the expected fleet using one-hour-ahead renewable forecasts, updated every 15 min. At each time, a one-hour prediction window is constructed from the currently observed fleet size. To satisfy the steady-state transversality conditions in (32) and (33), the predicted renewable power profile is extended by an exponentially decaying fictitious continuation over an additional hour. This artificial tail ensures smooth convergence of the Riccati system and produces a two-hour aggregate horizon for computing the target trajectory. Although this extension modifies the far-horizon dynamics, the inverse Nash equilibrium structure—namely, instantaneous utilization of available renewable energy—remains preserved. The distortion affects only the intertemporal allocation across vehicles, not the aggregate energy balance. Since near-term forecasts dominate the optimization, the induced approximation error remains limited.

3.1.2. Stochastic Discharging Regime

In discharging mode, the predicted power delivered to the grid relies directly on the expected fleet composition, rendering any fictitious continuation unnecessary. A strict one-hour predictive horizon is therefore applied and advanced in 15-min receding steps. This mathematical asymmetry highlights the fundamental operational distinction between the two regimes: charging is primarily renewable-driven, whereas discharging is strictly constrained by fleet availability. Accordingly, the anticipative target average SOC trajectory is continuously evaluated over this rolling one-hour window.

3.1.3. Anticipative Target Average SOC Trajectory

BEVs expected to connect within the forthcoming hour are treated as virtually present at the beginning of the horizon. Local control actions are nevertheless implemented only over the first 15-min step. This construction yields an anticipative average SOC trajectory used by currently connected BEVs to compute their feedback controls. Forecast errors in renewable production create temporary imbalances, absorbed either by a stationary parking battery (charging regime) or by remaining stored energy in BEVs (discharging regime). The overarching objective is to ensure efficient real-time utilization of available energy. The cost function in (15) inherently prioritizes vehicles with lower initial SOCs. Since instantaneous charging power is identical across BEVs, fairness must be enforced at the cumulative energy level rather than at the power level. The key idea is to artificially inflate the initial SOC of vehicles expected to connect later within the hour. By doing so, the MFG solution naturally prioritizes actually connected BEVs while still reserving energy for upcoming arrivals.

Let P_0 denote the set of BEVs initially present at t_0, with initial average SOC $\overline{x}_0^{P_0}$. Denote by W^{P_0} the share of forecast equivalent energy allocated to that population, and $\overline{x}_{60}^{P_0}$ its desired mean SOC at the end of the hour. The problem reduces to determining a fictitious weighted initial mean SOC $\overline{x}_{0,w}^{P_0}$—including inflated mean SOC values for future arrivals—such that applying the MFG solution with $\overline{x}_{0,w}^{P_0}$ produces exactly $\overline{x}_{60}^{P_0}$ for the P_0 BEV population after one hour. The mapping $\overline{x}_{0,w}^{P_0} \mapsto \overline{x}_{60}^{P_0}$ is monotone. Hence, $\overline{x}_{0,w}^{P_0}$ can be computed efficiently via a bisection procedure. Once obtained, the anticipative average SOC trajectory is computed and the associated optimal feedback optimal law is applied during the first 15-min interval only, after which forecasts are updated.

3.1.4. Predictor-Based Substitutions in the Aggregate Dynamics

On each interval $[t_0, t_0 + 60\,\text{min}]$, uncertain quantities are replaced by conditional expectations.

(i) Renewable power U_{W_t} is replaced by

$$\hat{U}_{W_t} = \mathbb{E}[U_{W_t} \mid U_{W_{t_0}}]. \tag{42}$$

In charging mode, an exponentially decaying one-hour extension is appended.

(ii) Fleet size n is replaced by its Poisson-based predictor:

$$\hat{n}(t) = n(t_0) + \sum_{k=1}^{3} \mathbb{E}\big[n(t_0 + 15k) - n(t_0 + 15(k-1)) \mid n(t_0)\big]. \tag{43}$$

This yields predicted quantities $\hat{\overline{\alpha}}$ (8) and $\hat{n}_\beta$ (9). This produces the predicted target SOC trajectory $\hat{\overline{x}}_t^{target}$ (14), where $\hat{\overline{x}}_0$ (13) is replaced by $\overline{x}_{0,w}^{P_0}$. We detail below the heuristic procedure used to compute the fictitious capacity-weighted initial SOC $\overline{x}_{0,w}^{P_0}$.

(iii) The linear gain, the Riccati gain, and the pressure trajectory are computed using the predicted target SOC trajectory, yielding $\hat{\overline{s}}_t$ (31), $\hat{\overline{\pi}}_t$ (30), and $\hat{q}_t^y$ (40).

3.1.5. Observer-Based Substitutions in the Optimal Control Equations

On each interval $[t_0, t_0 + 15\,\text{min}]$, each connected BEV:

(i) Computes its gain trajectories $\pi_{i,t}$ and $s_{i,t}$ (25) using $\hat{q}_t^y$.
(ii) Applies the optimal control law (22) and updates its SOC (10).

3.2. Determining the Fictitious Weighted Initial Mean SOC

3.2.1. Heuristic Mechanism

Fix an update time t_0 and consider the one-hour predictive horizon $[t_0, t_0 + 60\,\text{min}]$. Let P_0 denote the subset of BEVs that are connected at time t_0 and remain connected until $t_0 + 60\,\text{min}$. These vehicles are the only ones whose terminal SOC at $t_0 + 60\,\text{min}$ is effectively constrained by the one-hour optimization. Define the capacity mass of the P_0 BEV population:

$$C^{P_0} := \sum_{i \in P_0} \beta_i. \tag{44}$$

Let $\mathcal{P}_0^{\text{virtual}}$ denote the virtually augmented BEV population at t_0, defined as the union of the BEVs actually present (P_0) and those forecast to connect during the hour and remain connected until $t_0 + 60\,\text{min}$. Its total capacity mass is given by

$$C^{\mathcal{P}_0^{\text{virtual}}} := \sum_{i \in \mathcal{P}_0^{\text{virtual}}} \beta_i. \tag{45}$$

Let W_{60} be the forecast equivalent energy available over $[t_0, t_0 + 60\,\text{min}]$. Fairness requires allocating this energy proportionally to the capacity-time exposure of each subgroup. Under proportional allocation, the share of forecast energy assigned to the P_0 BEV population is

$$W^{P_0} = W_{60} \frac{C^{P_0}}{C^{\mathcal{P}_0^{\text{virtual}}}}. \tag{46}$$

Let

$$\overline{x}_0^{P_0} = \frac{\sum_{i \in P_0} \beta_i x_{i,0}}{C^{P_0}} \tag{47}$$

be the initial capacity-weighted mean SOC of the P_0 BEV population. Its desired terminal mean SOC at $t_0 + 60\,\text{min}$ is therefore

$$\overline{x}_{60}^{P_0} = \begin{cases} \overline{x}_0^{P_0} + \dfrac{W^{P_0}}{C^{P_0}}, & \text{during charging,} \\ \overline{x}_0^{P_0} - \dfrac{W^{P_0}}{C^{P_0}}, & \text{during discharging.} \end{cases} \tag{48}$$

The heuristic mechanism consists in introducing a fictitious weighted initial mean SOC for the virtually augmented BEV population at t_0, denoted by $\overline{x}_{0,w}^{P_0}$, obtained by artificially inflating the initial mean SOC of vehicles expected to connect later. The MFG equilibrium is then solved with $\overline{x}_{0,w}^{P_0}$. Define the mapping

$$\Phi(\bullet) = \text{capacity-weighted mean SOC of } P_0 \text{ at } t_0 + 60\,\text{min}, \tag{49}$$

when the MFG solution is computed using the initial aggregate mean SOC and the resulting feedback optimal law is applied to the P_0 BEV population.

3.2.2. Monotonicity, Existence, Uniqueness, and Convergence

The optimal MFG feedback law depends linearly on the deviation from the aggregate mean SOC trajectory. Within this framework, adjusting the fictitious initial aggregate mean SOC $\overline{x}_{0,w}^{P_0}$ artificially shifts this reference trajectory, thereby directly modulating the collective power exchange of the P_0 BEV population.

Importantly, while this relationship establishes a strictly monotone mapping $\Phi(\bullet)$ between $\overline{x}_{0,w}^{P_0}$ and $\overline{x}_{60}^{P_0}$, the direction of this monotonicity fundamentally depends on the active bidirectional regime. In the charging regime, decreasing the fictitious initial SOC creates a larger perceived aggregate energy deficit. This stimulates a more aggressive collective charging effort from the algorithm to reach equilibrium, which paradoxically elevates the actual terminal SOC of the physical fleet. In contrast, in the discharging regime, this dynamic is inherently inverted, as the algorithm allocates the discharging effort based on the perceived energy surplus relative to the reference trajectory. Despite this directional duality, the mapping $\Phi(\bullet)$ remains strictly monotonic in both operational modes. Given that terminal SOC values are inherently bounded by physical battery capacities, the existence of a unique $\overline{x}_{0,w}^{P_0}$ is guaranteed, such that

$$\Phi(\overline{x}_{0,w}^{P_0}) = \overline{x}_{60}^{P_0}. \tag{50}$$

Let $[x_{\min}, x_{\max}]$ be a valid SOC interval containing $\overline{x}^{P_0}_{0,w}$. The classical bisection algorithm—dynamically adapted to the correct gradient direction of the active regime—generates a sequence $\{x_k\}$ satisfying

$$|x_k - \overline{x}^{P_0}_{0,w}| \leq \frac{x_{\max} - x_{\min}}{2^k}. \tag{51}$$

Hence, convergence is linear and mathematically guaranteed independently of initialization, ensuring robust real-time solving across both supply and demand cycles.

3.2.3. Fairness and Relevance

The construction of $\overline{x}^{P_0}_{0,w}$ ensures that vehicles present during the entire hour receive exactly their proportional share of forecast energy, while reserving capacity for future arrivals. Since the MFG solution is recomputed every 15 min, any forecast error is corrected at the next MPC update. The heuristic therefore preserves:

- Proportional fairness with respect to heterogeneous capacities;
- Consistency with the mean-field equilibrium structure;
- Numerical stability and guaranteed convergence.

3.3. Mean-Field Predictive Control (MFPC) Algorithm

Building upon the perfect-information framework of the MFG inverse Nash algorithm (Algorithm 1), the MFPC algorithm (Algorithm 2) adapts this decentralized strategy for real-time operations under uncertainty. Operating on a receding horizon, the aggregator continually updates short-term forecasts of renewable generation ($\hat{U}_{W_t}$) and fleet dynamics ($\hat{n}(t)$) to construct an anticipative target average SOC trajectory ($\hat{\overline{x}}^{\text{target}}_t$), with a fictitious weighted initial mean SOC ($\overline{x}^{P_0}_{0,w}$). By solving the backward Riccati system over this limited forward-looking window—often extended by a fictitious decaying tail during charging regimes to ensure stability—the aggregator derives an estimated, time-varying pressure signal ($\hat{q}^y_t$). This signal is broadcast to the currently connected BEVs, which apply the optimal control policy over a brief 15-minute interval before the system updates actual battery realizations and shifts the optimization window forward.

Algorithm 2 Mean-Field Predictive Control (MFPC).

Initialize t, with T as length of the horizon.
while $t < T$ **do**
    Update renewable one-hour forecast power $\hat{U}_{W_t}$.
    Update BEV arrival/departure one-hour forecast fleet $\hat{n}(t)$.
    Compute $\overline{x}^{P_0}_{0,w}$, and set $\hat{\overline{x}}^{target}_{t_0} \leftarrow \overline{x}^{P_0}_{0,w}$.
    if charging regime **then**
        Extend $\hat{U}_{W_t}$ with one-hour exponentially decaying fictitious tail.
        Compute the two-hour horizon anticipative target SOC trajectory $\hat{\overline{x}}^{target}_t$.
    else
        Compute the one-hour horizon anticipative target SOC trajectory $\hat{\overline{x}}^{target}_t$.
    end if
    Solve a backward Riccati system $(\hat{\overline{s}}_t, \hat{\overline{\pi}}_t)$, and compute $\hat{q}^y_t$.
    Apply optimal control (using $\hat{q}^y_t$) over 15 min to currently connected BEVs.
    Update realizations (parking battery), and set $t \leftarrow t + 15$ min.
end while

3.4. MFPC-Based Charging and Discharging Results

3.4.1. Baseline Results in the Stochastic Case

Using the same BEV characteristics as in the deterministic case, we present the baseline results for the stochastic charging and discharging regimes in Figures 7–9.

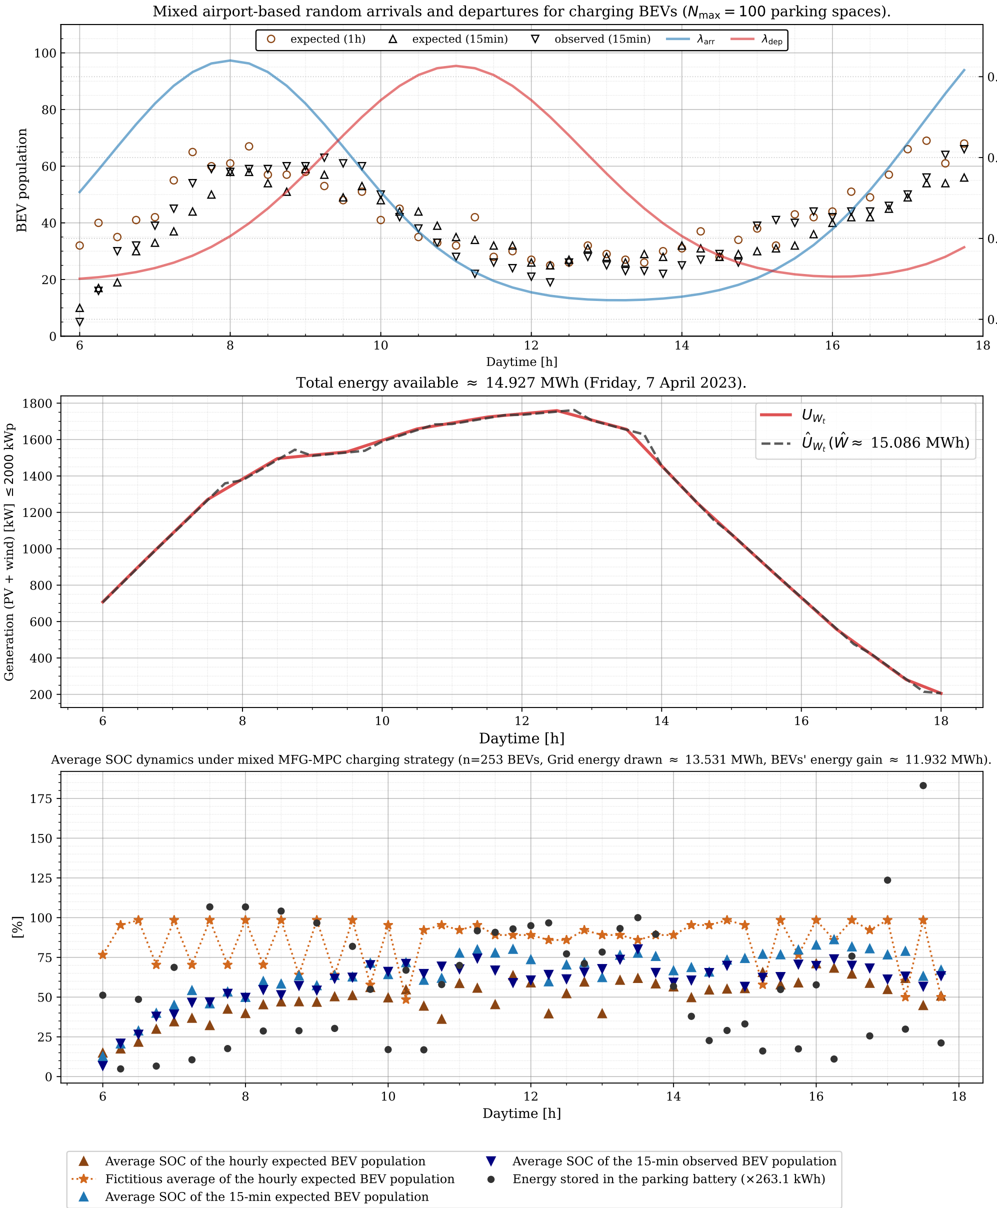

Figure 7. Grid-to-vehicle MFPC mixed-charging results: (**top**)—expected and observed BEV population using mixed airport-based Poisson intensities; (**middle**)—aggregate power (generation); (**bottom**)—average SOC dynamics and parking battery energy.

In the mixed-charging regime (Figure 7), the MFG-MPC controller elegantly overcomes the inherent challenges of intense fleet turnover, where stochastic mixity is at its strongest. To accommodate the extreme intensity of concurrent power flows inherent to such dynamic environments, the physical power threshold is deliberately set to $U_{\max} = 50$ kW, corresponding to standard DC fast infrastructure. Although the MFG-MPC controller operates with a receding aggregate horizon and anticipates expected arrivals, it leverages this expanded bandwidth to maintain impressive tracking accuracy without relying on real-time verification of available power. By relying exclusively on power curve forecasts and fleet predictions, the algorithm preserves a fully decentralized architecture—distinguishing itself from centralized benchmarks like COSA and COFA. Consequently, the observed 15-min average state of charge (SOC) rigorously and smoothly tracks the expected mean SOC, demonstrating the algorithm's robust resilience to high-frequency composition changes. The empirical statistics strongly validate this performance: across a dynamic fleet of 253 BEVs, the algorithm flawlessly orchestrates an energy gain of approximately

11,932 kWh from 13,531 kWh drawn from the grid. This closely aligns with the available renewable energy of 14,927 kWh, seamlessly absorbing the intermittent resources without the tracking deviations previously hypothesized. The tracking problem is thus resolved into an elegantly balanced dynamic strategy that maximizes renewable utilization while effectively buffering forecast uncertainties.

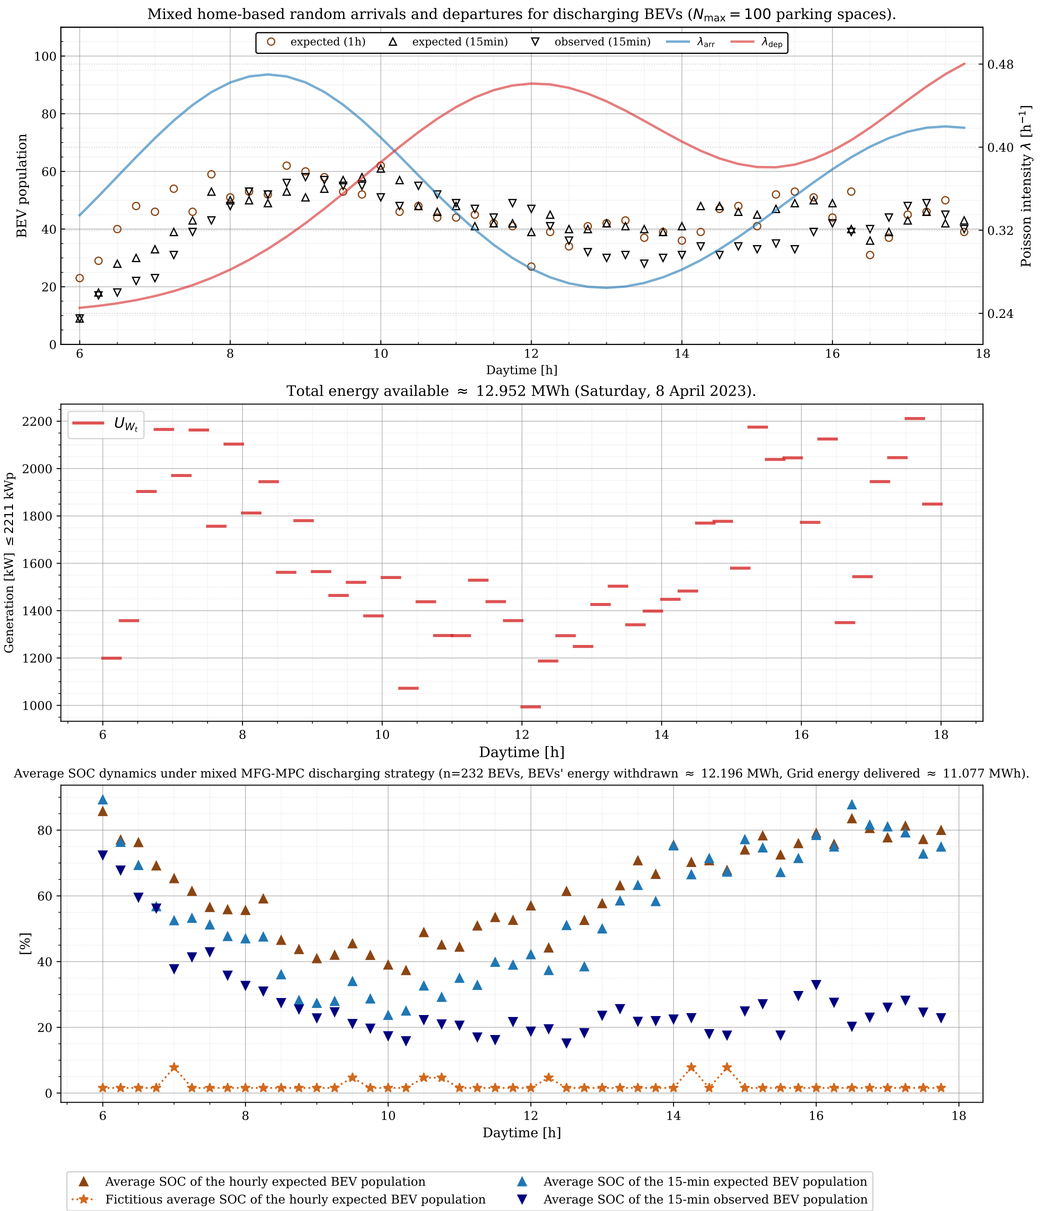

Figure 8. Vehicle-to-grid MFPC mixed-discharging results: (**top**)—expected and observed BEV population using mixed home-based Poisson intensities; (**middle**)—aggregate power (generation); (**bottom**)—average SOC dynamics.

Similarly, in the mixed-discharging regime (Figure 8) the resulting power profile stands as a testament to the efficacy of this purely predictive, receding-horizon strategy. Operating under the same $U_{\max} = 50$ kW DC fast-charging capacity to absorb the high-frequency composition changes, the MFG-MPC controller recurrently updates the predicted aggregate based solely on forecasts. This naturally induces a piecewise adaptive, rather than temporally rigid, discharge profile. For instance, rather than attempting to perfectly extract an arbitrary proportion at the exact moment of connection, the algorithm applies its fractional policy to the rolling one-hour expected aggregates. Operating on a diverse population of 232 BEVs, the control strategy withdraws approximately 12,196 kWh from the vehicles, delivering roughly 11,077 kWh to the grid against an available capacity target of 12,952 kWh. Notably, the portion of energy not extracted by the grid simply remains securely stored within the BEV batteries. This guarantees that no energy is inherently wasted, thereby safeguarding onboard flexibility for vehicle owners and striking an optimal balance between substantial grid support and user convenience in a highly dynamic mobility environment.

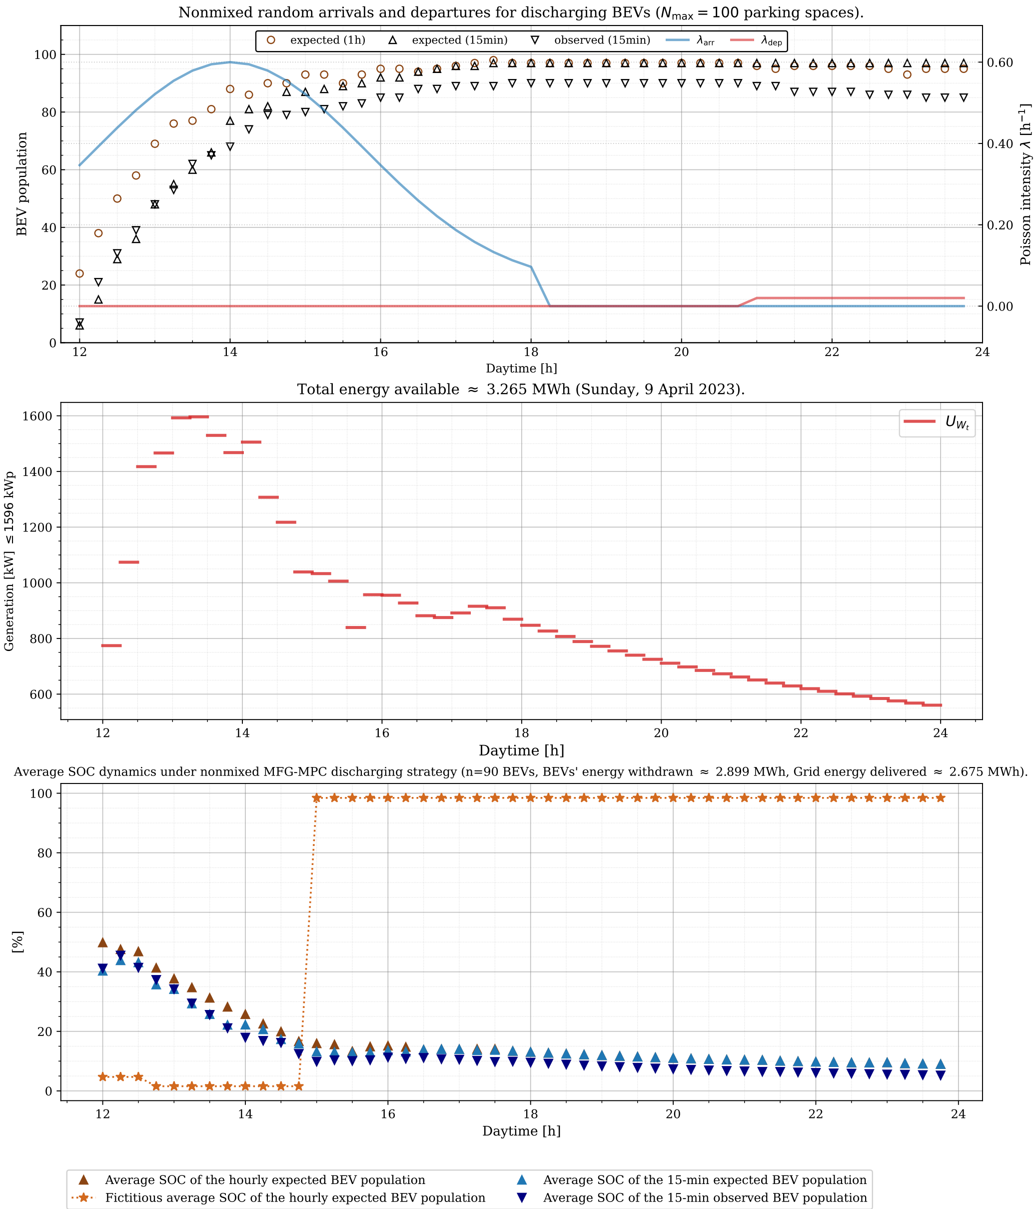

Figure 9. Vehicle-to-grid MFPC non-mixed-discharging results: (**top**)—expected and observed BEV population using non-mixed Poisson intensities; (**middle**)—aggregate power (generation); (**bottom**)—average SOC dynamics.

In contrast, the non-mixed discharging regime (Figure 9) presents a distinct operational paradigm characterized by moderate battery power flows, which directly correlate to a lower overall volume of available energy. Under these milder dynamic conditions, a standard upper power limit of $U_{\max} = 19.2$ kW proves entirely sufficient for the algorithm to smoothly extract the required energy. This stands in sharp contrast to the highly demanding mixed regimes discussed previously, which necessitated a much higher limit of $U_{\max} = 50$ kW.

3.4.2. Comparison of Charging and Discharging Algorithms in the Stochastic Case

Figures 10–12 compare the three algorithms—mean-field predictive control (MFPC), capacity-ordered fair allocation (COFA), and capacity-ordered saturation allocation (COSA)—under heterogeneous BEVs in mixed-charging and -discharging regimes. The SOCs at arrival are color-coded by capacity class (red for small, orange for medium, and green for large), with increasing color darkness reflecting a larger battery capacity within each respective class.

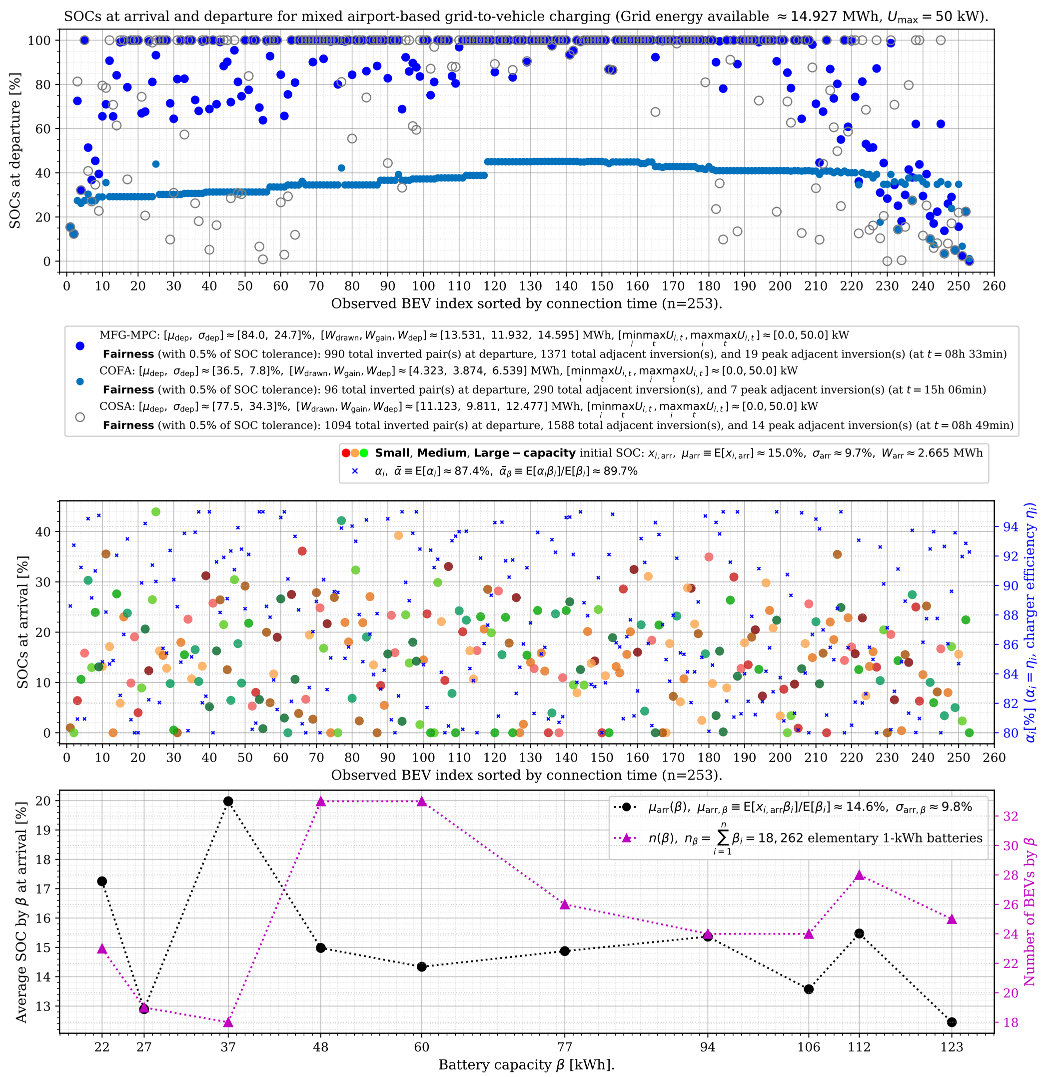

Figure 10. Grid-to-vehicle dynamic mixed-charging algorithms: (**top**)—SOCs at departure; (**middle**)—SOCs at arrival and charging efficiencies; (**bottom**)—initial average SOC and number of BEVs, per battery capacity.

In the mixed-charging regime (Figure 10), operating with an available grid energy of 14,927 kWh and an expanded physical capacity limit of $U_{\max} = 50$ kW to accommodate extreme power flows, the three algorithms display sharply differentiated energy utilization profiles. The MFG-MPC algorithm achieves the highest energy absorption, drawing roughly 13,531 kWh and yielding a net vehicle energy gain of 11,932 kWh. This robust performance demonstrates its capacity to elegantly capture the bulk of the renewable supply while remaining fully decentralized and preserving the relative SOC ordering induced by arrival heterogeneity. In contrast, the centralized COSA algorithm, despite its aggressive sequential saturation logic, falls short of the MFG-MPC's aggregate absorption, extracting only 11,123 kWh. Furthermore, COSA's mechanical efficiency comes at a severe expense to equity: departure SOCs are highly dispersed, reflecting its structural bias toward fully charging smaller-capacity batteries first. COFA performs exceptionally poorly in this intense dynamic setting, extracting a mere 4323 kWh. This severe limitation is not due to the concept of fairness per se, but stems from the restrictive interaction between its strict equalization objective and the per-charger power constraint across multiple 15-min control horizons. Because COFA continuously attempts to drive all connected SOCs toward a common uniform level, it becomes chronically power-constrained at each receding step, completely failing to absorb the abundant renewable supply despite the 50 kW bandwidth.

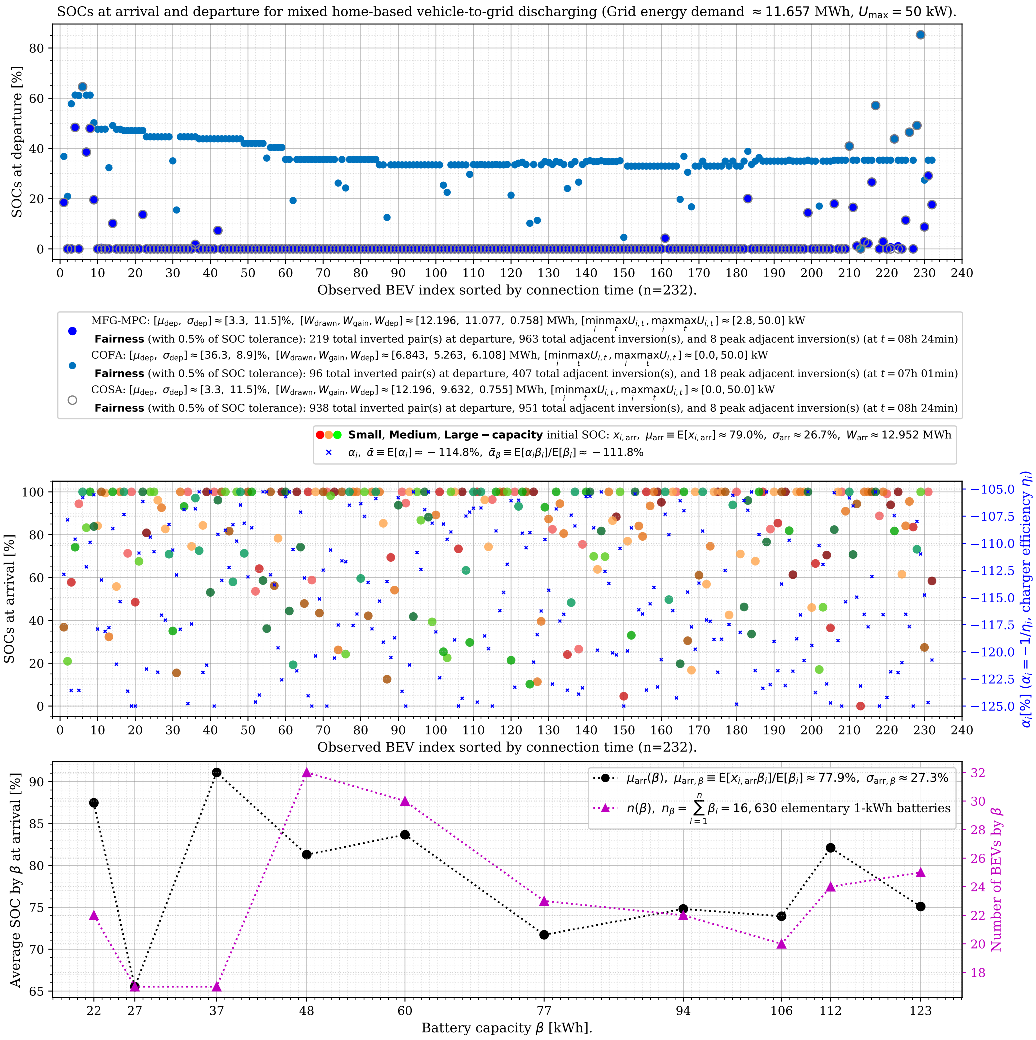

Figure 11. Vehicle-to-grid dynamic mixed-charging algorithms: (**top**)—SOCs at departure; (**middle**)—SOCs at arrival and discharging efficiencies; (**bottom**)—initial average SOC and number of BEVs, per battery capacity.

In the mixed-discharging regime (Figure 11), the aggregate grid demand is set to approximately 11,657 kWh against a total initial BEV fleet energy of roughly 12,952 kWh. Operating under the same $U_{max} = 50$ kW threshold to manage the intense fleet turnover, the MFG-MPC algorithm again demonstrates superior efficacy. It extracts 12,196 kWh from the vehicles to successfully deliver roughly 11,077 kWh to the grid. Importantly, by computing each vehicle's contribution locally and enforcing a consistent fractional policy, it effectively maintains fairness and largely preserves the initial SOC hierarchy, preventing extreme discrepancies at departure. Interestingly, the COSA algorithm draws the exact same aggregate energy from the fleet (12,196 kWh) but delivers slightly less to the grid (10,889 kWh). Moreover, COSA achieves its target by greedily depleting the largest batteries first; this structural front-loading drastically penalizes equity, resulting in severe SOC dispersion for vehicles departing early. COFA, on the other hand, struggles profoundly, drawing only 6843 kWh. The same structural limitation observed during charging reappears here: enforcing strict SOC equalization under rolling horizons heavily restricts the aggregate discharge potential when facing intense stochastic mixity.

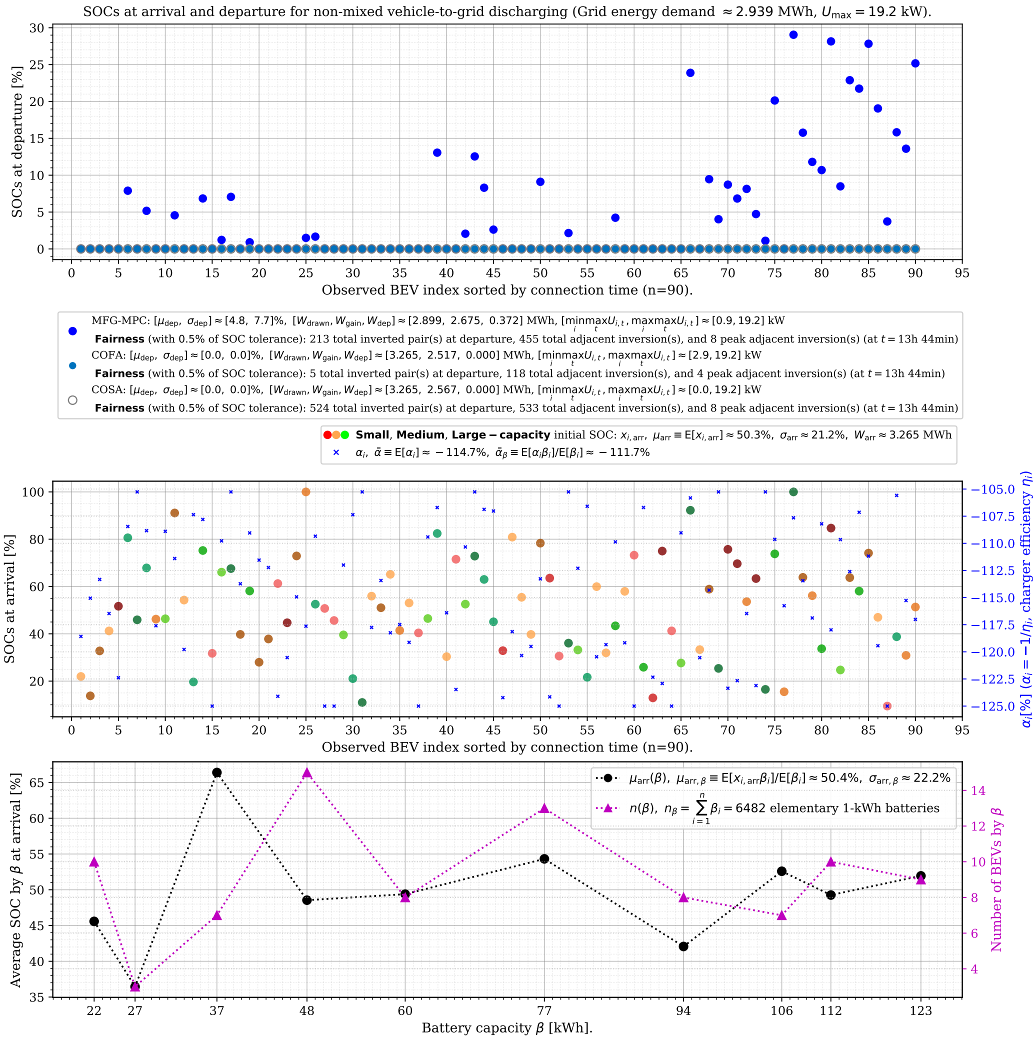

Figure 12. Vehicle-to-grid dynamic non-mixed-charging algorithms: (**top**)—SOCs at departure; (**middle**)—SOCs at arrival and discharging efficiencies; (**bottom**)—initial average SOC and number of BEVs, per battery capacity.

In constrast, the non-mixed-discharging scenario (Figure 12) presents a distinct operational paradigm where the COSA and COFA algorithms exhibit demonstrably better tracking of the aggregate target. In this statically stable environment, the grid demand is a modest 2939 kWh. Here, both COSA and COFA successfully withdraw 3265 kWh to deliver exactly 2928 kWh, nearly perfectly fulfilling the grid requirement. This relative success is primarily driven by the moderate intensity of power flows within the batteries, which correlates directly to a lower overall volume of available energy and the complete absence of high-frequency fleet turnover. Under these milder dynamic conditions, a standard upper power limit of $U_{\max} = 19.2$ kW proves more than sufficient, allowing the centralized algorithms to smoothly extract the required energy without encountering the severe saturation constraints that paralyze them in mixed regimes. Meanwhile, the MFG-MPC algorithm extracts a slightly lower total of 2899 kWh (delivering 2675 kWh). This conservative behavior is a natural consequence of its design: by rigorously maintaining fairness and enforcing its proportional policy across the fleet, it purposefully avoids over-discharging high-SOC vehicles. While its cumulative extraction falls moderately short of the ideal target in this specific static context, MFG-MPC guarantees fleet stability.

4. Conclusions

In this paper, we have engineered and validated a mean-field game–model predictive control (MFG-MPC) framework for the charging and discharging of heterogeneous battery electric vehicles (BEVs). Representing a significant original contribution to large-scale fleet management, this algorithm seamlessly orchestrates operations under both deterministic and intensely stochastic conditions. By natively embedding uncertainties in fleet dynamics and available power—rigorously modeled via Poisson and ARMA processes, respectively—the proposed framework delivers a purely predictive, decentralized, and scalable strategy. It elegantly equalizes state-of-charge (SOC) levels at departure while maximizing the utilization of available renewable energy. This architecture not only incentivizes BEV owners to utilize daytime parking for cost-effective green charging, but also unlocks massive, dynamic flexibility for vehicle-to-grid (V2G) support during peak hours. Importantly, the decentralized design ensures that each vehicle autonomously computes its optimal contribution without requiring burdensome centralized coordination. Through rolling 15-min updates, the controller dynamically reconciles predicted macroscopic aggregates with actual high-frequency fleet turnover. By handling a fully heterogeneous environment—encompassing diverse battery capacities, efficiencies, initial SOCs, and stochastic mobility patterns—this framework emerges as a robust, practically deployable engineering solution for next-generation urban mobility and grid integration.

The structural superiority of our algorithm is highlighted when benchmarked against two straightforward algorithms: capacity-ordered saturation allocation (COSA) and capacity-ordered fair allocation (COFA). While COSA employs an aggressive sequential saturation logic that can nearly exhaust energy budgets in certain regimes, this mechanical efficiency comes at the severe expense of user equity, inducing pronounced SOC dispersion by structurally prioritizing specific battery sizes. On the other hand, COFA attempts to enforce strict SOC equalization at departure; however, this extreme rigidity severely bottlenecks the system under physical power limits, paralyzing aggregate energy utilization across successive short control horizons. In contrast, the MFG-MPC framework strikes an optimal, elegant compromise between aggregate energy efficiency, algorithmic fairness, and dynamic adaptability. It achieves remarkable target compliance while strictly preserving the relative SOC ordering induced by arrival heterogeneity. Unlike COSA, it inherently avoids the inequitable charging/discharging of specific batteries, thereby safeguarding onboard flexibility for vehicle owners. Unlike COFA, it smoothly leverages macroscopic aggregate flexibility without sacrificing the responsiveness required in highly volatile mixed regimes. Ultimately, the MFG-MPC approach provides a physically consistent, highly resilient, and equitable strategy that effectively domesticates the intense stochastic mixity of modern BEV fleets.

Funding: This research received no external funding.

Data Availability Statement: All data were reported in the paper. For solar irradiance and wind speed data, the PVGIS database is on https://re.jrc.ec.europa.eu/pvg_tools/en/tools.html (accessed on 29 June 2026).

Conflicts of Interest: The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Renewable Power Modeling and Short-Term ARMA Forecasting

The purpose of this section is not to introduce novel forecasting techniques, but to provide reliable short-term renewable power predictions suitable for integration into the proposed MFG-MPC framework for coordinated BEV charging. The focus is therefore placed on methodological consistency, robustness, and compatibility with rolling-horizon control, rather than on maximizing pure forecasting accuracy (Alsharif et al., 2019; Chatfield, 2000; Zhang et al., 2019).

Appendix A.1. Site Description and Generation Data

The meteorological and renewable (solar and wind) generation data are extracted from the Photovoltaic Geographical Information System (PVGIS), an online platform developed by the Joint Research Centre of the European Commission. The geographical coordinates correspond to Montréal-Trudeau International Airport. This specific location serves as a representative empirical testbed to validate the robustness of the proposed control architecture under realistic solar irradiance profiles, rather than functioning as a site-specific sizing study. Consequently, the analytical results and operational conclusions derived herein are primarily methodological; they are fully generalizable and transposable to other geographical areas exhibiting varying meteorological characteristics.

The PVGIS configuration parameters for hourly data extraction are as follows:

- Geographical coordinates (latitude, longitude): 45.457°, −73.751°;
- Solar radiation database: PVGIS–ERA5;
- Year: 2023 (representing the most recent fully validated historical dataset currently available in the PVGIS repository);
- Photovoltaic (PV) mounting type and technology: fixed, crystalline silicon;
- PV module tilt angle (slope) and azimuth: 39°, −6° (both heuristically optimized for maximum annual energy yield);
- Installed PV peak power and system losses: 1000 kWp, 14% (accounting for thermal degradation, inverter inefficiencies, and ohmic losses).

The raw hourly meteorological measurements feature timestamps centered at HH:30 in Coordinated Universal Time (UTC). To synchronize the renewable generation profiles with local grid demand and vehicle mobility patterns, all timestamps are shifted to the local standard time utilizing a constant offset (e.g., UTC − 5). Notably, daylight-saving time adjustments are strictly excluded. This intentional omission prevents the introduction of artificial temporal discontinuities or phase shifts, thereby ensuring the strict temporal stationarity required for robust ARMA forecasting and predictive control.

Appendix A.1.1. Solar Power Reconstruction from Irradiance

Rather than directly using the PVGIS-provided PV power (P_{pvgis}) as ground truth, PV power is reconstructed from irradiance data in order to maintain physical consistency with the forecasting layer. Let $G_{\mathrm{POA}i}$ denote the global plane-of-array irradiance (Alsharif et al., 2019):

$$G_{\mathrm{POA}i} = G_{bi} + G_{di} + G_{ri}, \tag{A1}$$

where G_{bi}, G_{di}, and G_{ri} are, respectively, the global beam, diffuse, and ground-reflected irradiance components provided by PVGIS. The reconstructed PV power is defined as

$$P_{\text{rec}}(t) = P_{\text{nom}}\, \tilde{G}_{\text{POAi}}(t)\, \text{PR}, \tag{A2}$$

where P_{nom} denotes the installed peak capacity under standard test conditions (STC), $\tilde{G}_{\text{POA}i}$ is the global plane-of-array irradiance normalized with respect to STC irradiance, and PR is the performance ratio of the site. The performance ratio accounts for all system-level losses that prevent the conversion of incident irradiance into usable power, including temperature effects, inverter inefficiencies, wiring losses, and mismatch effects. It therefore represents an aggregate efficiency factor linking theoretical irradiance-based production to actual delivered power. In this study, PR is assumed constant over the year, which is appropriate for long-term energy consistency analysis. The value of PR is calibrated using an energy-consistency criterion,

$$\text{PR} = \frac{\sum_t P_{\text{pvgis}}(t)\, \Delta t}{\sum_t P_{\text{nom}}\, \tilde{G}_{\text{POAi}}(t)\, \Delta t}, \tag{A3}$$

thereby ensuring that the total reconstructed annual energy matches the reference energy obtained from PVGIS data. This calibration guarantees that the irradiance-driven reconstruction preserves the correct long-term energy balance while retaining high temporal resolution for control and forecasting purposes. To quantify the power reconstruction accuracy, the signed relative energy reconstruction error is defined as

$$\epsilon_W^{\text{rec}} = \frac{W_{\text{rec}} - W_{\text{pvgis}}}{W_{\text{pvgis}}}, \qquad W = \sum_t P(t)\, \Delta t. \tag{A4}$$

Appendix A.1.2. Wind Power Modeling

Wind speed data provided by PVGIS at a reference height ($h = 10$ m) are extrapolated to the turbine hub height H using a power-law profile:

$$v_H(t) = v_h(t) \left(\frac{H}{h}\right)^{\Upsilon}, \tag{A5}$$

where Υ is the surface roughness exponent. Wind power is then computed using a standard piecewise cubic turbine power curve:

$$P_{\text{wind}}(t) = \begin{cases} 0, & v_H < v_{\text{ci}} \text{ or } v_H \geq v_{\text{co}}, \\ P_{\text{rated}} \left(\dfrac{v_H - v_{\text{ci}}}{v_{\text{r}} - v_{\text{ci}}}\right)^3, & v_{\text{ci}} \leq v_H < v_{\text{r}}, \\ P_{\text{rated}}, & v_{\text{r}} \leq v_H < v_{\text{co}}, \end{cases} \tag{A6}$$

where v_{ci}, v_{r}, and v_{co} are, respectively, the cut-in, rated, and cut-out wind speeds. The adopted parameters are as follows:

- Wind speeds (rated, cut-out, cut-in): 11.5 m/s, 25 m/s, 3 m/s;
- Turbine hub height: 100 m;
- Surface roughness exponent: 0.14;
- Installed wind rated power: 1000 kW.

Appendix A.1.3. Hourly Meteorological Data in a Selected Site for One Year from PVGIS

Figure A1 presents the hourly air temperature at 2 m, global plane-of-array irradiance (G_{POAi}), reconstructed photovoltaic (PV) power from G_{POAi}, wind speed at 10 m, and reconstructed wind power data for the year 2023.

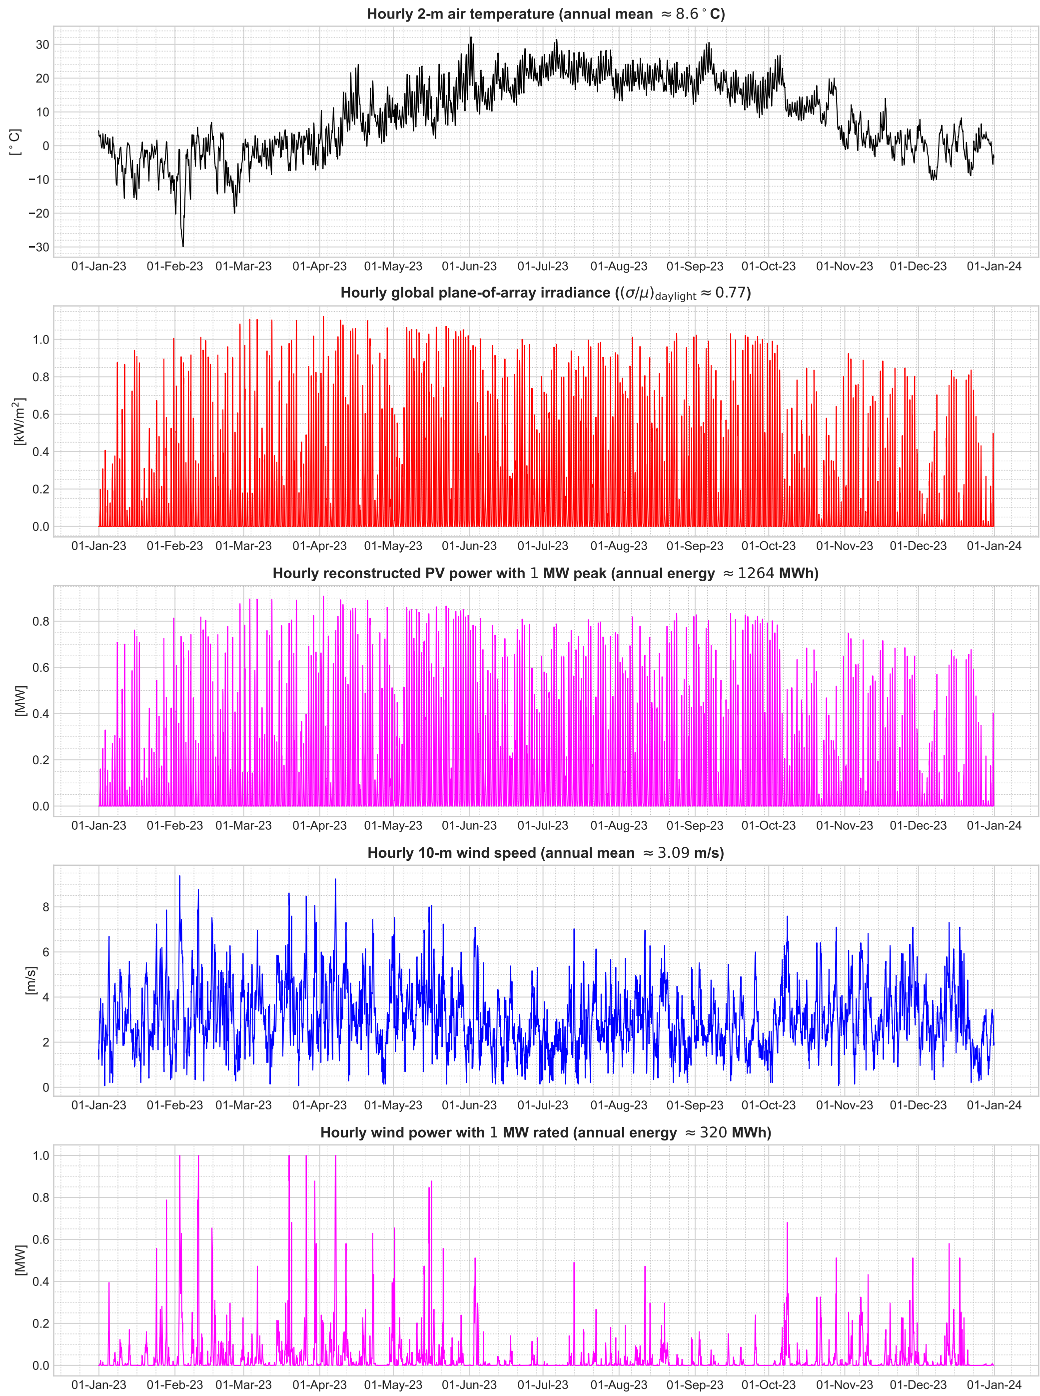

Figure A1. Hourly meteorological data at a selected site for 2023 from PVGIS.

The meteorological data reflects the pronounced seasonal amplitude characteristic of the site, with an annual mean air temperature of approximately 8.6°C and rigorous winter conditions spanning from November to March. Despite this northern climate, the G_{POAi} time series demonstrates strong diurnal regularity and highly sustained irradiance levels, particularly from spring to autumn. The daylight variability index, defined as $(\sigma/\mu)_{daylight} \approx 0.77$, indicates moderate intraday variability, which is well suited to short-horizon statistical forecasting methods such as ARMA models. Using a fixed performance ratio calibrated from PVGIS data, the reconstructed PV power yields a substantial annual energy production of 1264 MWh per megawatt installed. This translates to a good solar capacity factor of 14.4%.

On the other hand, the 10 m wind speed distribution exhibits a low annual mean (≈3.09 m/s) interspersed with highly intermittent peaks. As a result, the reconstructed wind power generates only 320 MWh annually per megawatt rated, yielding a marginal wind capacity factor of 3.65%. While these results clearly establish the location as a solar-dominant site rather than a high-yield wind site, relying exclusively on photovoltaics is insufficient for a year-round BEV charging infrastructure. Wind peaks frequently occur during winter storms and nocturnal hours, directly offsetting the solar deficits. Therefore, hybridizing solar and wind generation constitutes a mathematically sound and physically necessary strategy to construct a more resilient, complementary power supply curve for the charging hub across all seasons.

Appendix A.1.4. Daily PV and Wind Power Profiles

Figure A2 compares daily photovoltaic (PV) and wind power profiles for the period from 1 April to 7 April 2023.

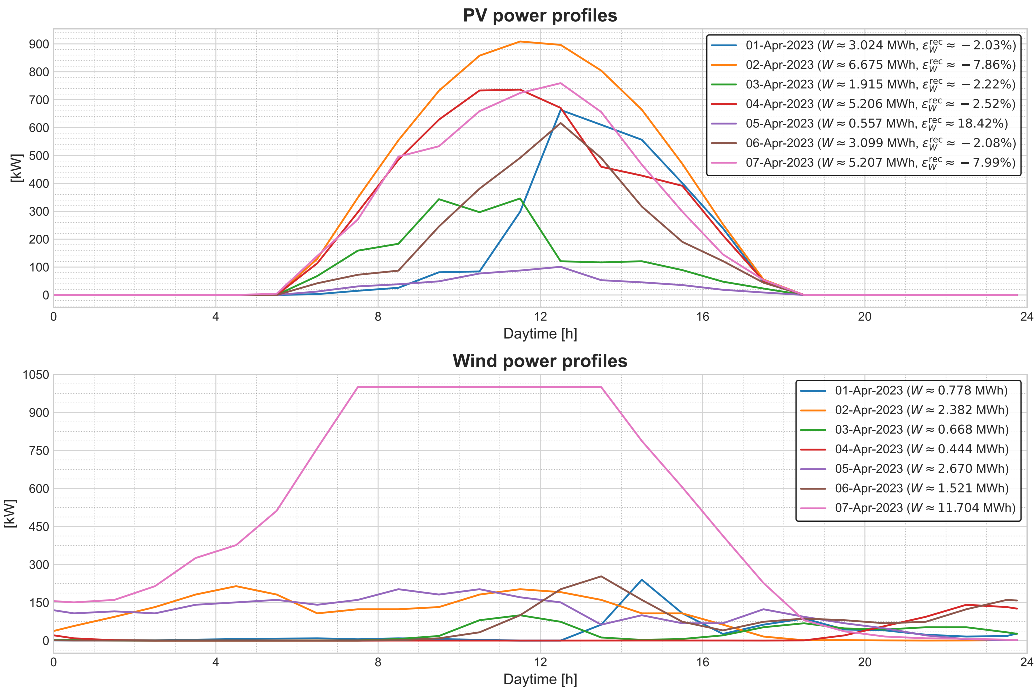

Figure A2. PV and wind power curves in a parking lot reconstructed from PVGIS reference data.

The daily PV profiles, illustrated over the first week of April 2023, exhibit a consistent bell-shaped structure, with generation naturally peaking around solar noon. On clear-sky days (e.g., 2 April), the PV system performs near its installed capacity, yielding a substantial daily energy of approximately 6.6 MWh. In contrast, heavily overcast conditions (e.g., 5 April) severely distort this ideal curve, suppressing the daily yield to a mere 0.55 MWh. Despite this weather-induced variability, the reconstruction error (ϵ_W^{rec}) remains tightly bounded—typically between −2% and −8% on productive days—confirming the high reliability of the irradiance-to-power conversion model, even if relative errors mathematically inflate during extreme low-irradiance events (18.4%).

In contrast, the daily wind power profiles exhibit profound irregularity and a distinct lack of diurnal correlation. While baseline days (1 to 4 April) show marginal and sporadic generation fluctuating between 0.4 MWh and 2.3 MWh, the data reveals the highly volatile and episodic nature of wind resources. During significant meteorological events—such as the sustained wind pattern captured on 7 April—the turbine reaches its 1 MW rated capacity and remains saturated for over eight hours. This single event generated a massive 11.7 MWh, nearly doubling the best solar yield of that same week.

Overall, these daily profiles elegantly validate the necessity of a hybrid generation architecture. Solar energy provides a predictable and dominant diurnal baseload, highly suitable for standard daytime BEV charging. Wind energy, while generally sporadic and unreliable on a daily basis, acts as a powerful, episodic complement. Notably, the meteorological systems that bring heavy cloud cover and diminish solar output are often accompanied by strong winds; thus, wind generation effectively offsets PV deficits during stormy periods, reinforcing the energetic resilience of the charging hub.

Appendix A.2. ARMA Modeling of Meteorological Time Series

Short-term forecasts of $G_{\mathrm{POA}i}$ and wind speed are generated using autoregressive moving-average (ARMA) models (Alsharif et al., 2019; Chatfield, 2000). The deterministic trajectories introduced previously serve as reference benchmarks for evaluating forecast accuracy.

Appendix A.2.1. Preprocessing

Let $\{z_t\}$ denote the observed time series corresponding to solar irradiance or wind speed measurements. For solar irradiance, a logarithmic variance-stabilizing transformation is applied to the raw data in order to mitigate right-skewness and heteroskedasticity, whereas wind speed is modeled without transformation. Nighttime observations associated with near-zero irradiance levels are excluded from the identification procedure to prevent numerical ill-conditioning during parameter estimation.

Appendix A.2.2. Centering and Stationarity

Let μ denote the empirical mean of $\{z_t\}$ computed over the identification window. The centered process is defined as

$$y_t = z_t - \mu. \tag{A7}$$

Approximate covariance stationarity of the centered series $\{y_t\}$ over the estimation window is verified using standard unit-root tests such as Dickey–Fuller or Phillips–Perron. All ARMA modeling is performed on $\{y_t\}$.

Appendix A.2.3. ARMA Formulation

For prescribed orders (p, q), the centered process satisfies

$$y_t = \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} + \varepsilon_t, \tag{A8}$$

where $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$ is a white-noise innovation.

Appendix A.2.4. Parameter Estimation and Order Selection

Model identification proceeds in two stages: order selection and parameter estimation. For each candidate (p, q) within a predefined search set, parameters are estimated using the conditional sum-of-squares (CSS) method. Innovations are computed recursively as

$$\varepsilon_t = y_t - \sum_{i=1}^{p} \phi_i y_{t-i} - \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}. \tag{A9}$$

The CSS estimator minimizes the empirical mean squared innovation:

$$(\hat{\phi}, \hat{\theta}) = \arg\min_{\phi,\theta} \frac{1}{N_s} \sum_{t=1}^{N_s} \varepsilon_t^2, \tag{A10}$$

where N_s denotes the number of samples in the identification window. The innovation variance is estimated as

$$\hat{\sigma}^2 = \frac{1}{N_s} \sum_{t=1}^{N_s} \hat{\varepsilon}_t^2. \tag{A11}$$

Model-order selection is performed using the Bayesian information criterion (BIC):

$$\mathrm{BIC}(p, q) = N_s \ln(\hat{\sigma}^2) + (p+q) \ln(N_s). \tag{A12}$$

The optimal order (p^*, q^*) minimizes $\mathrm{BIC}(p, q)$. The ARMA parameters are then re-estimated using CSS with the selected order.

Appendix A.2.5. Rolling Multi-Step Forecasting

Given observations $\{z_1, \ldots, z_T\}$, the forecasting procedure is carried out as follows. If required, the raw measurements are first transformed (e.g., logarithmic transformation for solar irradiance). The series is then centered according to $y_t = z_t - \mu$, where μ denotes the empirical mean over the identification window. Historical innovations $\hat{\varepsilon}_t$ are reconstructed recursively from the fitted ARMA model. Multi-step forecasts are subsequently computed via

$$\hat{y}_{T+k} = \sum_{i=1}^{p^*} \hat{\phi}_i \hat{y}_{T+k-i} + \sum_{j=1}^{q^*} \hat{\theta}_j \hat{\varepsilon}_{T+k-j}, \tag{A13}$$

where future innovations are set to their conditional expectation, $\hat{\varepsilon}_{T+k} = 0$, $k \geq 1$. The data forecasts are then recovered as $\hat{z}_{T+k} = \hat{y}_{T+k} + \mu$, and the inverse transformation is applied whenever required to return to physical units.

Appendix A.3. ARMA Forecasting Performance

Appendix A.3.1. Forecasting Performance Metrics

Forecast accuracy is evaluated using both instantaneous and energy-based indicators. Classical pointwise metrics such as the mean absolute error (MAE) and the root mean squared error (RMSE) quantify short-term deviations. However, for charging control applications, the most relevant measure is the signed relative energy prediction error:

$$\epsilon_W^{\mathrm{pred}} = \frac{W_{\mathrm{pred}} - W_{\mathrm{obs}}}{W_{\mathrm{obs}}}, \tag{A14}$$

which directly captures the mismatch between predicted and available renewable energy over the control horizon. This metric is particularly appropriate for MPC-based charging strategies, which rely primarily on anticipated energy availability rather than instantaneous power precision.

Appendix A.3.2. Numerical Results

Figures A3 and A4 present rolling ARMA forecasts with a one-hour prediction horizon and updates every 15 min. Results are shown for PV and wind power profiles calculated from forecasts of $G_{\mathrm{POA}i}$ and wind speed.

The ARMA model deployed for PV forecasting (Figure A3) accurately captures the smooth diurnal evolution of solar generation. As illustrated by the predictive profiles, forecast errors are predominantly concentrated around local maxima and during rapid, weather-induced transitions—a characteristic limitation of purely statistical models when facing abrupt irradiance variations. Nevertheless, the reconstructed PV power forecasts inherit a high degree of predictive accuracy. For instance, on highly productive days (e.g., 6 and 7 April), the root mean square error (RMSE) remains well contained below 8.5 kW, translating to an excellent normalized RMSE (NRMSE) of roughly 1%. Even on heavily overcast days (e.g., 5 April), the mean absolute error (MAE) is restricted to merely 1.8 kW. Notably, the relative energy prediction error (ϵ_W^{pred}) consistently remains below the 5% threshold. This specific metric is of paramount importance for BEV charging applications, as it directly and reliably quantifies the aggregate energy available to the fleet over the entire charging window.

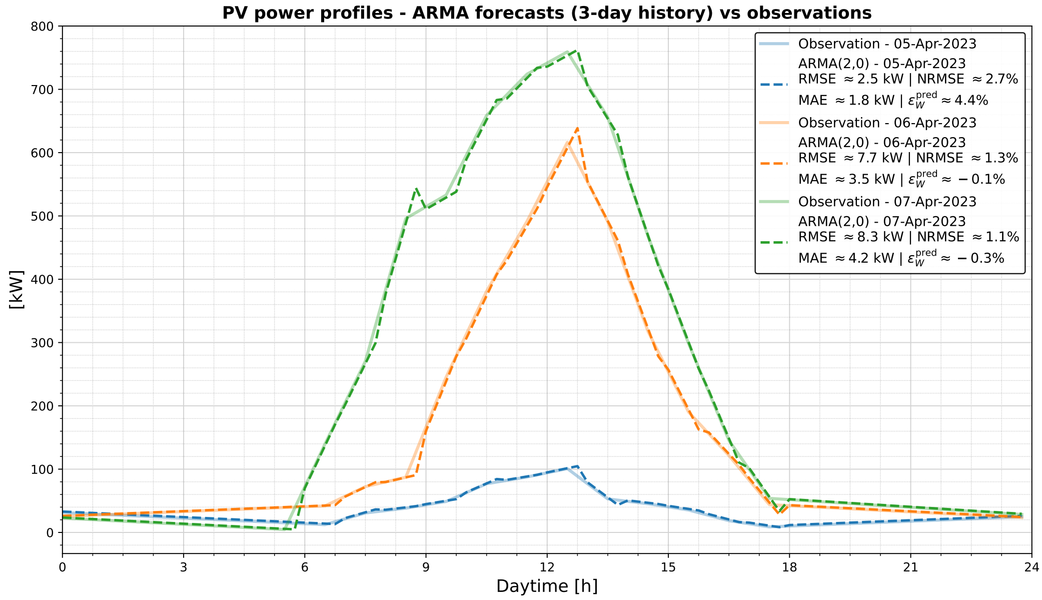

Figure A3. ARMA forecasts and observations for PV power profiles.

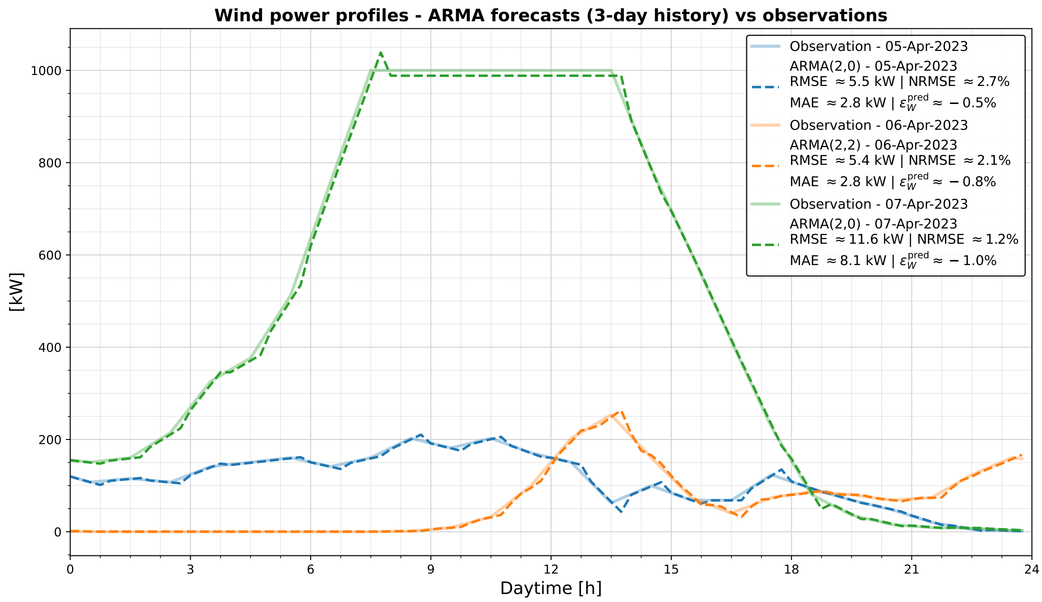

Figure A4. ARMA forecasts and observations for wind power profiles.

The ARMA model applied to wind forecasting (Figure A4) also exhibits strong structural agreement with the observations, confirming the suitability of autoregressive methods for short-term wind prediction under moderately varying conditions. However, the statistical metrics highlight an inherent physical challenge: due to the highly nonlinear nature of the turbine's power curve (which scales with the cube of the wind speed prior to saturation), even minor deviations in wind speed forecasting translate into amplified power prediction errors. This nonlinear amplification clearly explains the comparatively higher absolute RMSE values observed for wind power (reaching up to 11.6 kW during the 7 April wind event), despite the underlying wind speed forecasts being statistically accurate. Nonetheless, the relative predictive metrics remain robust, with the wind NRMSE staying below 3% and the relative energy error (ϵ_W^{pred}) bounded within an impressive $\pm 1\%$ across all tested days.

Overall, the combination of statistically consistent ARMA forecasting with physically rigorous power reconstruction yields highly accurate short-term renewable energy predictions. This performance validates the integration of these models into the proposed MFG-MPC framework. The forecasting layer thus provides a robust, reliable, and mathematically sound interface between the stochastic nature of renewable generation and the decentralized coordination of the BEV charging control.

Appendix B. Predictive Modeling of Random Arrivals and Departures of BEVs

The objective of this section is not to develop a statistically optimal parking-lot demand model, but to generate short-term, MPC-compatible predictions of the number of battery electric vehicles (BEVs) connected to a charging/discharging facility under realistic operational constraints. These predictions are subsequently used as inputs to the proposed MFG-MPC charging/discharging framework.

Appendix B.1. State-Dependent Poisson Modeling with Finite Population

Arrivals and departures of vehicles are modeled using state-dependent Poisson counting processes under finite population constraints. This framework effectively captures both the stochastic nature of parking demand and the physical limitations of the infrastructure. Notably, should empirical observations suggest overdispersion in the data, this formulation can be seamlessly extended to a negative binomial model to better account for increased variance. The operating day is discretized into uniform time intervals of duration $\Delta = 15$ min. Let t_k denote the beginning of the k-th interval and $n_k \in \{0, 1, \ldots, N_{\max}\}$ the number of BEVs connected to the infrastructure at time t_k, where $N_{\max}$ is the nominal parking capacity.

Appendix B.2. Modeling of Random Arrivals

The number of new connections at time t_{k+1} is modeled as

$$A_{k+1} \sim \text{Poisson}\big((N_{\max} - n_k)\,\lambda_{\text{arr}}(t_{k+1})\,\Delta\big), \tag{A15}$$

where λ_{arr} is the arrival intensity, assumed piecewise constant over each 15-min interval. The factor $(N_{\max} - n_k)$ explicitly captures the finite-capacity effect: as the parking lot fills up, the expected number of arrivals naturally decreases. Physical arrivals may occur continuously within an interval, but a BEV is assumed to become connected and available for charging/discharging only at the next sampling instant.

Appendix B.3. Modeling of Random Departures

Departures are modeled differently from arrivals in order to account for the empirically correlation between arrival time and parking duration (Benenson et al., 2008; Small &

Verhoef, 2007). Vehicles arriving later tend to remain parked longer. To capture this behavior without introducing a full duration-dependent hazard model, each BEV is associated with its connection time considering a first-in-first-out (FIFO) queue. The number of departure requests at time t_{k+1} is drawn as

$$D_{k+1} \sim \text{Poisson}\Big(n_k\,\lambda_{\text{dep}}(t_{k+1})\,\Delta\Big), \tag{A16}$$

where λ_{dep} is the departure intensity, also assumed piecewise constant.

Appendix B.4. Population Dynamics and Rolling Prediction

The 15-min connected BEV population, bounded by $N_{\max}$, evolves according to

$$n_{k+1} = n_k + A_{k+1} - D_{k+1}. \tag{A17}$$

For each 15-min interval, the realized numbers of arrivals and departures are then observed and used to update the conditional distributions for the next intervals. This rolling update yields one-hour-ahead predictions of the expected BEV population, refreshed every 15 min. This conditional prediction mechanism naturally integrates with a receding-horizon MPC scheme and provides the MFG-MPC charging/discharging algorithm with short-term forecasts of the number of participating vehicles. At each time step, the model produces both the expected numbers of BEV arrivals and departures over the next four 15-min intervals.

Appendix B.5. Numerical Results for Non-Mixed and Mixed Poisson Random Arrivals and Departures

Figures A5–A7 illustrate non-mixed and mixed Poisson models for random arrivals and departures.

A non-mixed Poisson model means:

- arrival and departure processes independently sampled;
- departures not conditioned on congestion or parking occupancy;
- each vehicle assigned to a parking space.

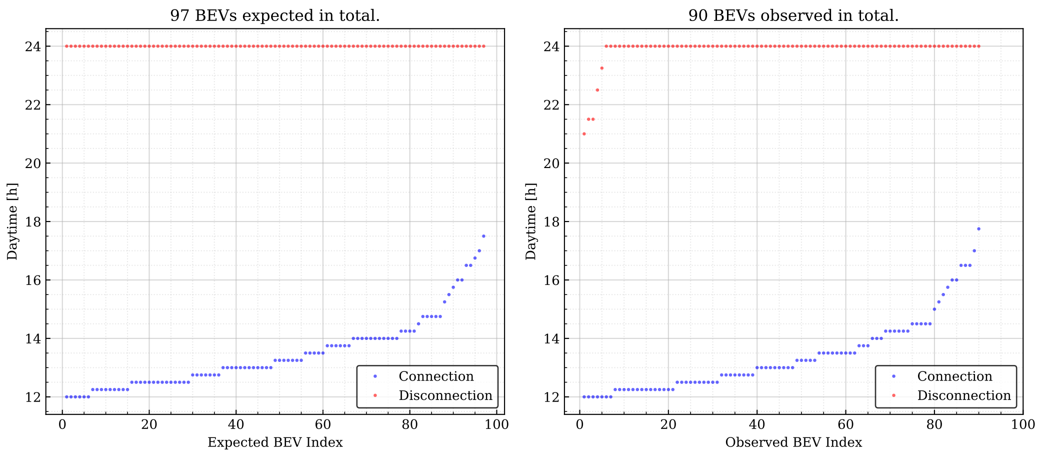

Figure A5. Expected and observed BEV population for non-mixed-discharging BEVs.

A mixed Poisson model means:

- conditional arrival probabilities;
- possible queuing or congestion effects;
- a parking space shared among vehicles.

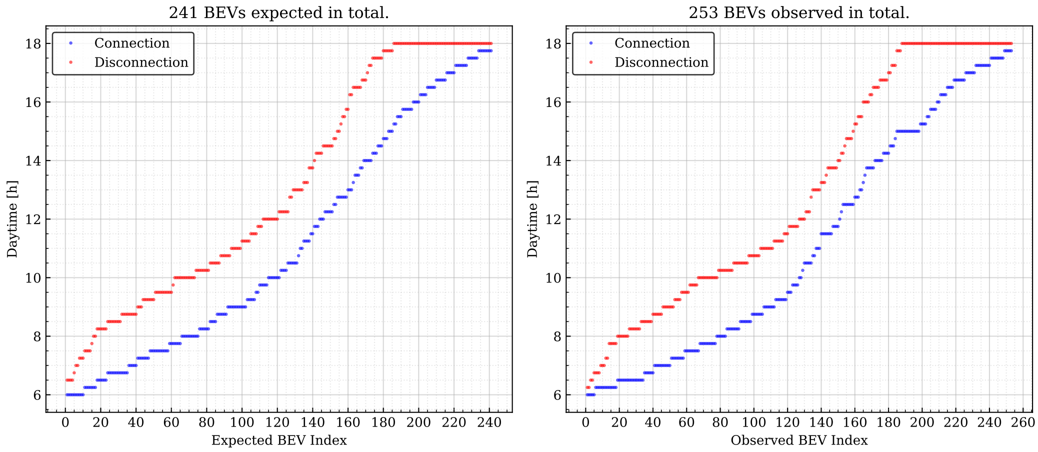

Figure A6. Expected and observed BEV population for mixed airport-based charging BEVs.

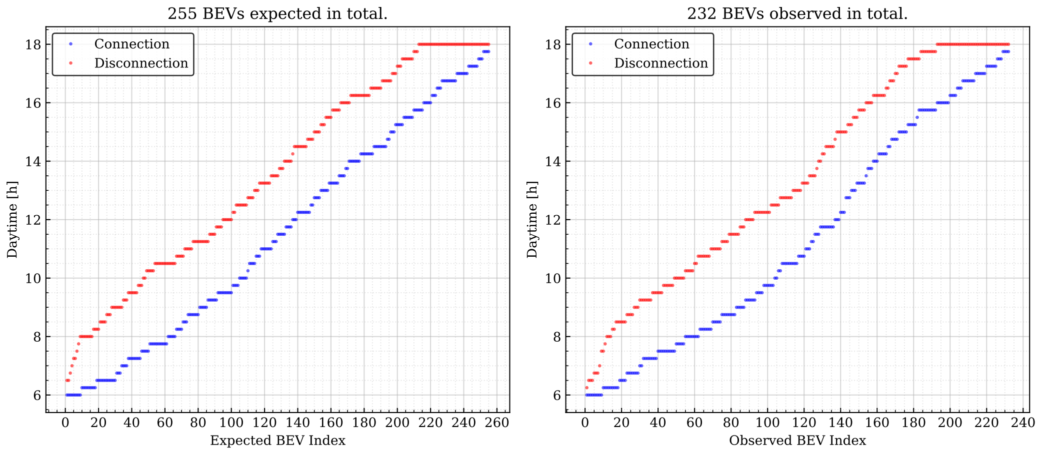

Figure A7. Expected and observed BEV population for mixed home-based discharging BEVs.

References

Adegbohun, F., von Jouanne, A., Agamloh, E., & Yokochi, A. (2024). A review of bidirectional charging grid support applications and battery degradation considerations. *Energies*, *17*(6), 1320. [CrossRef]

Aghmadi, A., & Mohammed, O. A. (2024). Energy storage systems: Technologies and high-power applications. *Batteries*, *10*(4), 141. [CrossRef]

Ahuja, S. (2015). *Mean field games with common noise* [Doctoral's thesis, Stanford University].

Alsharif, M. H., Younes, M. K., & Kim, J. (2019). Time series ARIMA model for prediction of daily and monthly average global solar radiation: The case study of Seoul, South Korea. *Symmetry*, *11*(2), 240. [CrossRef]

Anderson, B. D. O., & Moore, J. B. (2007). *Optimal control, linear quadratic methods* (12th ed.). Dover Publications.

Bagchi, A. (1993). *Optimal control of stochastic systems*. Prentice Hall.

Benenson, I., Martens, K., & Birfir, S. (2008). PARKAGENT: An agent-based model of parking in the city. *Computers, Environment and Urban Systems*, *32*(6), 431–439. [CrossRef]

Canada Energy Regulator. (2026). *Canada's energy future 2026: Energy supply and demand projections to 2050* (Technical report). Canada Energy Regulator.

Cecchin, A., & Pelino, G. (2019). Convergence, fluctuations and large deviations for finite state mean field games via the master equation. *Stochastic Processes and their Applications*, *129*(11), 4510–4555. [CrossRef]

Chatfield, C. (2000). *Time-series forecasting*. Chapman.

Comi, A., & Elnour, E. (2024). Challenges for implementing vehicle-to-grid services in parking lots: A state of the art. *Energies*, *17*(24), 6240. [CrossRef]

Cominesi, S. R., Farina, M., Giulioni, L., Picasso, B., & Scattolini, R. (2015, July 1–3). *Two-layer predictive control of a micro-grid including stochastic energy sources* [Conference presentation]. 2015 American Control Conference (ACC), Chicago, IL, USA.

Electric Vehicle Database. (2026). Available online: https://ev-database.org/ (accessed on 29 June 2026).

Energy Institute. (2026). *Statistical review of world energy 2026, 75th edition* (Technical report). Energy Institute.

Huang, M., Caines, P. E., & Malhamé, R. P. (2007). Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε-Nash equilibria. *IEEE Transactions on Automatic Control*, *52*(9), 1560–1571. [CrossRef]

Huang, M., Malhamé, R. P., & Caines, P. E. (2006). Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, *6*(3), 221–252. [CrossRef]

International Energy Agency. (2026). *Global EV outlook 2026: Growing sales amid an energy crisis* (Technical report). International Energy Agency.

Jacobson, D. H. (1970). New conditions for boundedness of the solution of a matrix Riccati differential equation. *Journal of Differential Equations*, *8*(2), 258–263. [CrossRef]

Kizilkale, A. C., Salhab, R., & Malhamé, R. P. (2019). An integral control formulation of mean field game based large scale coordination of loads in smart grids. *Automatica*, *100*, 312–322. [CrossRef]

Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, *2*, 229–260. [CrossRef]

Lenet, Q., Nazir, M. S., & Malhamé, R. P. (2021, December 13–17). *An inverse Nash mean field game-based strategy for the decentralized control of thermostatic loads* [Conference presentation]. 60th IEEE Conference on Decision and Control (CDC 2021), Austin, TX, USA.

Li, P., Wang, X., Li, S., Chan, H., & An, B. (2023, May 1–5). *Population-size-aware policy optimization for mean-field games* [Conference presentation]. The Eleventh International Conference on Learning Representations, Kigali, Rwanda.

Ma, Z., Callaway, D. S., & Hiskens, I. A. (2013). Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Transactions on Control Systems Technology*, *21*(1), 67–78. [CrossRef]

Muhindo, S. M. (2024). Mean field game-based algorithms for charging in solar-powered parking lots and discharging into homes a large population of heterogeneous electric vehicles. *Energies*, *17*(9), 2118. [CrossRef]

Muhindo, S. M., Malhamé, R. P., & Joos, G. (2021a). A novel mean field game-based strategy for charging electric vehicles in solar powered parking lots. *Energies*, *14*(24), 8517. [CrossRef]

Muhindo, S. M., Malhamé, R. P., & Joós, G. (2021b). *Mean field game-based control of sharing daily solar energy between electric vehicles in a parking lot* (Les Cahiers du GERAD No. G-2021-29). Groupe d'études et de recherche en analyse des décisions.

Nour, M., Chaves-Ávila, J. P., Magdy, G., & Sánchez-Miralles, Á. (2020). Review of positive and negative impacts of electric vehicles charging on electric power systems. *Energies*, *13*(18), 4675. [CrossRef]

Novickij, I., & Joós, G. (2019, May 5–8). *Model predictive control based approach for microgrid energy management* [Conference presentation]. 2019 IEEE Canadian Conference of Electrical and Computer Engineering (CCECE), Edmonton, AB, Canada.

Parisio, A., Rikos, E., & Glielmo, L. (2014). A model predictive control approach to microgrid operation optimization. *IEEE Transactions on Control Systems Technology*, *22*(5), 1813–1827. [CrossRef]

Qi, W., Liu, J., & Christofides, P. D. (2013). Distributed supervisory predictive control of distributed wind and solar energy systems. *IEEE Transactions on Control Systems Technology*, *21*(2), 504–512. [CrossRef]

Small, K. A., & Verhoef, E. T. (2007). *The economics of urban transportation* (2nd ed.). Taylor and Francis Group.

Tamilselvi, S., Gunasundari, S., Karuppiah, N., Razak, R. K. A., Madhusudan, S., Nagarajan, V. M., Sathish, T., Shamim, M. Z. M., Saleel, C. A., & Afzal, A. (2021). A review on battery modelling techniques. *Sustainability*, *13*(18), 10042. [CrossRef]

Tuchnitz, F., Ebell, N., Schlund, J., & Pruckner, M. (2021). Development and evaluation of a smart charging strategy for an electric vehicle fleet based on reinforcement learning. *Applied Energy*, *285*, 116382. [CrossRef]

Wan, M., Yu, H., Huo, Y., Yu, K., Jiang, Q., & Geng, G. (2024). Feasibility and challenges for vehicle-to-grid in electricity market: A review. *Energies*, *17*(3), 679. [CrossRef]

Xu, F., Zhao, K., & Fu, Q. (2024). Mobility-aware optimal trade planning and mean field game-based decentralized charging control for large-scale electric vehicles. *Asian Journal of Control*, *26*(2), 598–616.

Zabetian-Hosseini, A., Joos, G., & Boulet, B. (2022, April 25–28). *Centralized control design for bidirectional DC charging stations to enable V2G* [Conference presentation]. 2022 IEEE/PES Transmission and Distribution Conference and Exposition (T&D), New Orleans, LA, USA.

Zhang, S., Zhou, T., Zhou, T., Sun, L., & Liu, C. (2019). Kernel ridge regression model based on beta-noise and its application in short-term wind speed forecasting. *Symmetry*, *11*(2), 282. [CrossRef]