Spatial Equilibrium of V2G-Mobility Markets: A Mean Field Game Approach

Zhiyang Chen[1] and Fangni Zhang[1]

[1]Department of Data and Systems Engineering, The University of Hong Kong, Hong Kong, China

Short summary

The integration of Vehicle-to-Grid (V2G) services into electric autonomous ride-sourcing fleets creates a fundamental trade-off between spatial displacement for mobility and temporal retention for energy arbitrage. This study investigates the stationary equilibrium of such large-scale fleets to analytically characterize the structural properties of the optimal fleet management strategy. We formulate a mean field game (MFG) framework under continuous energy constraints. We prove that the optimal policy exhibits a unique, spatially dependent threshold structure characterized by forced charging and V2G participation boundaries. To compute the equilibrium, we develop a Structure-Preserving Threshold Fictitious Play (SP-TFP) algorithm that ensures robust convergence by operating directly on the threshold manifold. Numerical results reveal that the equilibrium spontaneously organizes into Voronoi-like service cells. We further demonstrate that V2G infrastructure can either smooth spatial value disparities or exacerbate market polarization, depending on its deployment relative to demand hotspots. Ultimately, this study offers platform operators strategic insights into energy-spatial sorting and provides policymakers with actionable levers to enhance service equity and network robustness through targeted infrastructure planning.
Keywords: Electric autonomous vehicles, equilibrium, mean field, shared mobility, vehicle-to-grid.

1 Introduction

The rapid adoption of electric autonomous vehicles (EAVs) in on-demand ride-sourcing platforms is fundamentally altering the nature of transportation assets (Yan et al., 2024). Unlike traditional vehicles, EAVs possess a dual identity as both transportation tools and mobile energy storage units. Vehicle-to-Grid (V2G) technology unlocks this duality by allowing EAVs to discharge stored energy back to the power grid (Guille & Gross, 2009), enabling platforms to expand their operational boundaries toward integrated mobility and energy synergy.

However, optimizing operations for a large-scale, V2G-enabled fleet presents significant computational challenges, as traditional centralized control suffers from the curse of dimensionality. Therefore, a distributed control mechanism is highly practical, wherein each EAV operates as an independent, rational agent making decentralized decisions to maximize its own virtual profit. At the microscopic level, each EAV must continuously resolve a fundamental conflict between spatial displacement for mobility revenue and temporal retention for V2G arbitrage. Because an individual EAV's profitability is heavily dependent on the macroscopic market state, where aggregate fleet behavior endogenously shapes local mobility demand and V2G dynamics, analyzing this system requires an equilibrium framework that couples microscopic vehicle strategies with macroscopic fleet distributions.

Existing literature on electric ride-sourcing primarily treats energy strictly as a constraint that must be replenished, rather than a tradable commodity (Ke et al., 2019; Cai et al., 2023; Wei et al., 2023; Yang et al., 2024). Conversely, V2G economics research largely assumes vehicles are parked and available for grid interaction for the majority of the day, ignoring the stochastic mobility demands of ride-sourcing fleets where the opportunity cost of V2G includes the lost income from potential ride requests (Gao et al., 2025; Lai et al., 2025; Zhang et al., 2025). Furthermore, while mean field game (MFG) theory provides a powerful framework for modeling decentralized interactions and has been applied to transportation, existing models predominantly rely on discrete state spaces or simplified variables (Zhu et al., 2021; Ameli et al., 2022; Wu et al., 2025). They rarely incorporate complex internal states, like continuous battery energy levels, coupled with spatial dynamics.

To bridge these gaps, this study proposes an MFG framework to model the spatial equilibrium of a decentralized, V2G-enabled EAV fleet on a continuous spatial-energy state space. The primary objective is to analytically characterize the structural properties of the independent EAV's optimal

strategy at a stationary equilibrium. Specifically, this study aims to prove that the optimal policy exhibits a unique, spatially dependent threshold structure and to demonstrate how strategic V2G infrastructure deployment can act as a topological lever to either smooth spatial value disparities or exacerbate market polarization.

2 Methodology

We consider a continuous-time, infinite-horizon MFG to model the decentralized operations of a large-scale EAV ride-sourcing fleet. In this distributed framework, each EAV acts as a rational, independent agent seeking to maximize its own virtual profit.

Micro-level single vehicle problem

The state of an individual EAV is defined by a tuple $s = (x, e) \in \mathcal{S}$, where $x \in \mathcal{X} \subset \mathbb{R}^2$ represents the geographical location, and $e \in [0, e_{\max}]$ denotes the battery energy level. The EAV continuously monitors its state and environmental signals to choose an action from the set $\mathcal{A} = \{\text{Charge, Wait}\}$.

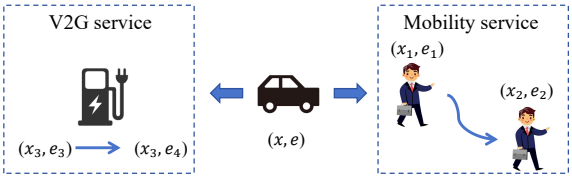


Figure 1: V2G-mobility trade-off

Let $V(x, e)$ denote the optimal value function (expected discounted cumulative virtual profit) for an EAV in state (x, e). If the vehicle chooses to stop waiting and proceed to a charging station $c \in \mathcal{C}$ at location x_c, the value of this charging action, denoted as $\mathcal{M}V(x, e)$, is given by:

$$\mathcal{M}V(x, e) = \max_{c \in \mathcal{C}} \{-C_{\text{charge}} + e^{-\rho(\tau(x,x_c)+\tau_c)} V(x_c, e_{\max})\}, \tag{1}$$

where C_{charge} includes the travel, energy purchase, and service time costs, $\tau(x, x_c)$ is travel time, τ_c is charging time, and ρ is the discount rate.
Conversely, if the vehicle chooses to wait, its value evolution is governed by the infinitesimal generator $\mathcal{L}$ associated with the stochastic arrival of mobility requests (at rate $\lambda_m(x)$) and V2G opportunities (at rate $\lambda_v(x)$):

$$\mathcal{L}V(x, e) = \lambda_m(x)\mathbb{E}_j[\Delta V_m^*(x, e, j)] + \lambda_v(x)\mathbb{E}_k[\Delta V_v^*(x, e, k)], \tag{2}$$

where ΔV_m^* and ΔV_v^* represent the optimal value increments upon optimally pricing and accepting a mobility trip j or determining the optimal discharge amount for a V2G session k, respectively.
To explicitly quantify these value increments, we construct a unified cost structure comprising distance-based operating costs (c_o), time-based opportunity costs (c_t), energy replacement prices (c_e), and battery degradation costs (c_d).
Upon the arrival of a mobility request j with origin x_j, destination $x_{j'}$, and trip distance d_j, the EAV evaluates the immediate net revenue R_m. This revenue accounts for the proposed fare $F_{m,j}$ minus the total operational costs for both the pickup and the service phases:

$$R_m(F_{m,j}, x, j) = F_{m,j} - (c_o + c_e\eta)(d(x, x_j) + d_j) - c_t(\tau(x, x_j) + \tau_j), \tag{3}$$

where η is the energy consumption rate and τ_j is the service duration. The optimal value increment for a mobility trip is then formulated as a pricing problem:

$$\Delta V_m^*(x, e, j) = \max_{F_{m,j}} \left\{ D(F_{m,j}) \left[R_m(F_{m,j}, x, j) + e^{-\rho(\tau(x,x_j)+\tau_j)} V(x_{j'}, e_{\text{post}}) - V(x, e) \right] \right\}, \tag{4}$$

where e_{post} is the residual energy after the trip, and $D(\cdot) \in [0, 1]$ is the strictly decreasing probability function of a customer accepting the price $F_{m,j}$.
Similarly, for a V2G opportunity k at location x_k, the immediate net revenue R_v includes the V2G payment minus travel and explicit degradation costs:

$$R_v(e_{v,k}, x, k) = F_v e_{v,k} - (c_o + c_e\eta)d(x, x_k) - c_t(\tau(x, x_k) + \tau_v) - (c_e + c_d)e_{v,k}, \tag{5}$$

where F_v is the grid purchasing price and $e_{v,k}$ is the chosen discharge amount. The corresponding optimal value increment is:

$$\Delta V_v^*(x,e,k) = \max_{0\le e_{v,k}\le e_{\text{avail}}} \left\{ R_v(e_{v,k},x,k) + e^{-\rho(\tau(x,x_k)+\tau_v)} V(x_k, e_{\text{post}}) - V(x,e) \right\}, \tag{6}$$

where e_{avail} ensures the battery does not drop below a safety threshold.
Combining the options of charging and waiting, $V(x,e)$ must satisfy the following continuous-time Quasi-Variational Inequality (QVI) of a Hamilton-Jacobi-Bellman (HJB) equation:

$$\min\{\rho V - \mathcal{L}V, V - \mathcal{M}V\} = 0. \tag{7}$$

This QVI naturally partitions the state space $\mathcal{S}$ into distinct regions indicating the optimal policy π^*: a charging region Ω_C where immediate intervention is optimal ($V = \mathcal{M}V$), a mobility acceptance region Ω_M, and a V2G acceptance region Ω_V.

Macro-level population evolution

While the QVI characterizes the optimal strategy of a single EAV, the macroscopic collective behavior is described by the evolution of the population distribution. Let $m(t,x,e)$ denote the probability density function of available vehicles at time t, location x, and energy level e. The continuous evolution of $m(t,x,e)$ is governed by a Forward Kolmogorov (FPK) Master Equation, which captures the net result of probability inflows $\mathcal{J}_{\text{in}}$ and outflows $\mathcal{J}_{\text{out}}$ across the state space:

$$\frac{\partial m(t,x,e)}{\partial t} = \mathcal{J}_{\text{in}}(t,x,e) - \mathcal{J}_{\text{out}}(t,x,e). \tag{8}$$

The local outflow rate is strictly determined by the active policy in the respective functional regions:

$$\mathcal{J}_{\text{out}}(t,x,e) = m(t,x,e) \cdot [\delta_{\text{charge}} \cdot 1_{\Omega_C}(x,e) + \lambda_m(x) \cdot 1_{\Omega_M}(x,e) + \lambda_v(x) \cdot 1_{\Omega_V}(x,e)], \tag{9}$$

where $1_\Omega(\cdot)$ is the indicator function and δ_{charge} represents the departure rate to a charging station. Conversely, $\mathcal{J}_{\text{in}}(t,x,e)$ integrates vehicles entering state (x,e) from three channels: arriving after completing a mobility trip, completing a V2G discharge, or returning to the fleet with a full battery $e_{\max}$ after charging.

Stationary V2G-mobility equilibrium

The coupling between the microscopic EAV agent and the macroscopic population occurs through a market matching mechanism. The effective supply of vehicles for mobility $N_m(x,t)$ and V2G $N_v(x,t)$ at location x is determined by integrating the density $m(t,x,e)$ over the respective optimal decision regions Ω_M and Ω_V.
The opportunity arrival rates perceived by individual EAVs are endogenous variables derived from a frictional matching function $\psi(\cdot)$:

$$\lambda_m(x,t) = \frac{\Lambda_m^{\text{dem}}(x) \cdot \psi\left(\frac{N_m(x,t)}{\Lambda_m^{\text{dem}}(x)}\right)}{N_m(x,t)}, \tag{10}$$

$$\lambda_v(x,t) = \frac{\Lambda_v^{\text{dem}}(x) \cdot \psi\left(\frac{N_v(x,t)}{\Lambda_v^{\text{dem}}(x)}\right)}{N_v(x,t)}, \tag{11}$$

where $\Lambda_m^{\text{dem}}(x)$ and $\Lambda_v^{\text{dem}}(x)$ are exogenous demand rates for mobility services and V2G services, respectively. This creates a closed-loop system: the HJB policy regions determine the effective supply, which determines the congestion-aware matching rates, which in turn feed back into Eq. (7).
A stationary mean field equilibrium (SMFE) is achieved when the system reaches a time-invariant state ($\partial m/\partial t = 0$) where the optimal policy π^* derived from the HJB equation and the stationary distribution m^* derived from the FPK equation are mutually consistent with the generated market rates λ_m^* and λ_v^*.

3 Results and discussion

Structural properties of optimal policy

Theoretical analysis of the continuous-time QVI reveals that the interaction between the EAV's internal energy dynamics and the external market environment naturally induces a unique optimal control structure.

Theorem 1 (Existence and Uniqueness of Control Thresholds) *For any location $x \in \mathcal{X}$, there exist a charging threshold $e_c(x)$ and a V2G threshold $e_v(x)$ satisfying $e_c(x) < e_v(x)$, which partition the energy domain into three continuous regimes:*

- *Forced charging ($e \le e_c(x)$): The vehicle rejects all orders and proceeds to a charging station, satisfying $V(x,e) = \mathcal{M}V(x,e)$.*
- *Mobility-only ($e_c(x) < e < \min\{e_v(x), e_{\max}\}$): The vehicle remains available for mobility services but rejects V2G opportunities because the marginal value of energy is still higher than the net V2G revenue.*
- *Mixed service ($e \ge e_v(x)$ [1]): The vehicle accepts both mobility and V2G opportunities as the shadow price of energy drops below the market-clearing price.*

Furthermore, in the forced charging region, the optimal choice of charging station induces an additively weighted Voronoi diagram based on economic distance rather than purely physical distance. Having established the operational regimes via threshold functions, we further formalize how the independent EAV's optimal decisions correspond to distinct mechanisms of economic arbitrage.

Proposition 1 *For a mobility request j from x_j to $x_{j'}$ with distance d_j and duration τ_j, the optimal reservation price $F^*_{m,j}$ proposed by the independent EAV is given by:*

$$F^*_{m,j} = C_{\text{ops}}(j) + \left[V(x,e) - e^{-\rho(\tau(x,x_j)+\tau_j)}V(x_{j'}, e_{\text{post}})\right] - \mathcal{H}(F^*_{m,j}), \tag{12}$$

where $C_{\text{ops}}(j)$ represents the total operational cost (including energy and time) and $\mathcal{H}(\cdot)$ is the monopoly markup.

Eq. (12) reveals a sophisticated mechanism of spatial arbitrage driven autonomously by the EAV. The bracketed term indicates that the EAV explicitly prices the change in its own internal state value. If the destination $x_{j'}$ promises high future expected rewards (e.g., a high-demand core), the term $e^{-\rho(\tau(x,x_j)+\tau_j)}V(x_{j'}, e_{\text{post}})$ becomes large. This dynamically reduces the effective opportunity cost, leading to a lower fare that essentially subsidizes the EAV's own relocation to a more profitable area. Conversely, for destinations with low expected future value, the EAV autonomously demands a higher fare to compensate for the penalty of being relocated to a suboptimal zone. Thus, decentralized pricing acts as a self-regulating dispatching signal, leveraging spatial gradients to guide fleet redistribution without explicit central commands.

Proposition 2 *Within the mixed service region ($e \ge e_v(x)$), the optimal discharge amount $e^*_{v,k}$ for a V2G opportunity at x_k satisfies the following condition:*

$$e^{\rho(\tau(x,x_k)+\tau_v)}\frac{\partial V}{\partial e}(x_k, e_{\text{post}}) = F_v - c_e - c_d. \tag{13}$$

This condition encapsulates the temporal arbitrage behavior. The partial derivative $\partial V/\partial e$ serves as the dynamic reservation price for the vehicle's energy inventory. Crucially, the EAV does not blindly discharge all available energy down to the threshold $e_v(x)$. Because the value function is strictly concave with respect to energy, $\partial V/\partial e$ increases as energy decreases. This creates a negative feedback loop: as the vehicle sells power to the grid, its remaining energy becomes more scarce and valuable, preventing further discharge to preserve a core reserve for future mobility tasks.

Theorem 2 (Infrastructure Proximity Effect) *Let $d_{\text{charge}}(x)$ represent the generalized distance from location x to the nearest charging station. The control thresholds satisfy the following spatial properties:*

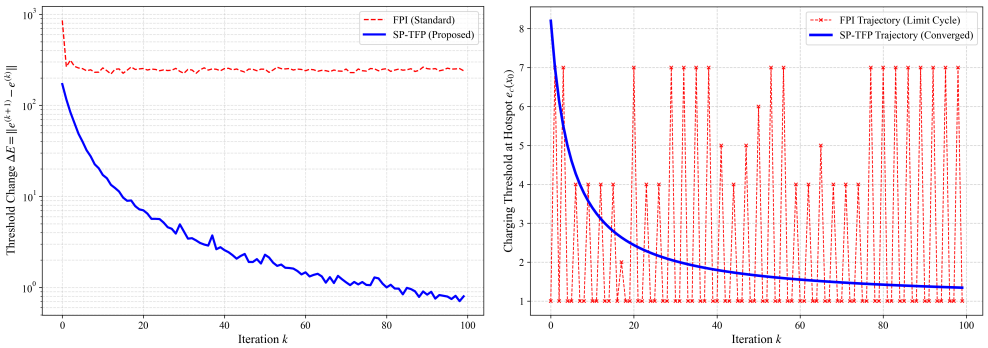
[1] This regime exists physically if and only if $e_v(x) < e_{\max}$.

- *The charging threshold $e_c(x)$ increases as $d_{\text{charge}}(x)$ decreases.*
- *The V2G threshold $e_v(x)$ decreases as $d_{\text{charge}}(x)$ decreases.*

Theorem 2 uncovers a fundamental spatial mechanism driven by the cross-derivative $\partial^2 V/\partial x \partial e$. When EAVs are near charging stations, the generalized cost of replenishment is minimal, maximizing the value of the charging option. To equilibrate this, the vehicle requires a higher energy level to justify waiting, pushing $e_c(x)$ upwards. Concurrently, easily replaceable energy yields a lower shadow price ($\partial V/\partial e$), causing it to intersect the market V2G price at a lower level, thus pushing $e_v(x)$ downwards. This spatial variation of $\Delta e(x) = e_v(x) - e_c(x)$ mathematically defines two distinct operational regimes: a narrow band inducing an opportunity charging regime with frequent, shallow cycles near stations, and a wide band inducing a range preservation regime in remote areas to avoid the high penalty of deadheading to a station.

Algorithm convergence and spatial heterogeneity

To compute the stationary equilibrium without the numerical diffusion typical of standard MFG solvers, we developed the Structure-Preserving Threshold Fictitious Play (SP-TFP) algorithm. SP-TFP operates directly on the threshold manifolds $e_c(x)$ and $e_v(x)$, significantly reducing the dimensionality of the learning process.



(a) Convergence analysis of the Bellman error (log scale)

(b) Policy evolution trajectory

Figure 2: Performance of the SP-TFP

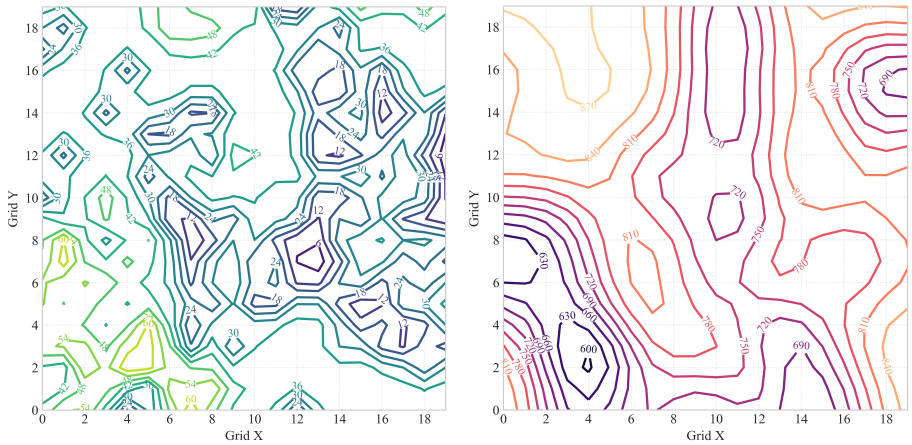
Numerical experiments (illustrated in Figure 2) validate that standard fixed point iteration (FPI) fails to converge, trapping the system in persistent limit cycles. This instability is rooted in the strong strategic substitutability of the V2G-mobility market. When fleet density in a region is high, local matching rates drop, prompting the HJB solver to output a greedy response that raises the charging threshold (instructing vehicles to leave and charge). In the subsequent iteration, the sudden drop in fleet density causes a spike in perceived matching rates, causing the threshold to collapse again.

To break this congestion-induced negative feedback loop, the proposed SP-TFP algorithm operates directly on the threshold manifolds. Instead of adopting the greedy thresholds $\hat{e}_c(x)$ and $\hat{e}_v(x)$ immediately, SP-TFP applies a convex smoothing update in the function space:

$$e_c^{(n+1)}(x) = (1 - \alpha_n) e_c^{(n)}(x) + \alpha_n \hat{e}_c(x). \tag{14}$$

By dampening the update with a decaying step size α_n, SP-TFP preserves the three-regime geometric structure proven in Theorem 1 while eliminating high-frequency oscillations. This structure-preserving approach achieves robust, exponential convergence, strictly maintaining the physical interpretability of the policy throughout the computation process.

According to Figure 3, spatially, the computed SMFE confirms a highly heterogeneous policy. The charging threshold $e_c(x)$ exhibits a cellular structure peaking significantly near charging stations where the generalized cost of access is low. This induces an opportunity charging regime characterized by frequent, shallow cycling. A strong negative spatial correlation exists between $e_c(x)$ and $e_v(x)$; in energy-scarce remote areas, V2G operations are effectively shut down to preserve the battery for high-value mobility tasks.



(a) Spatial distribution of the charging threshold

(b) Spatial distribution of the V2G threshold

Figure 3: Spatial heterogeneity of control policy

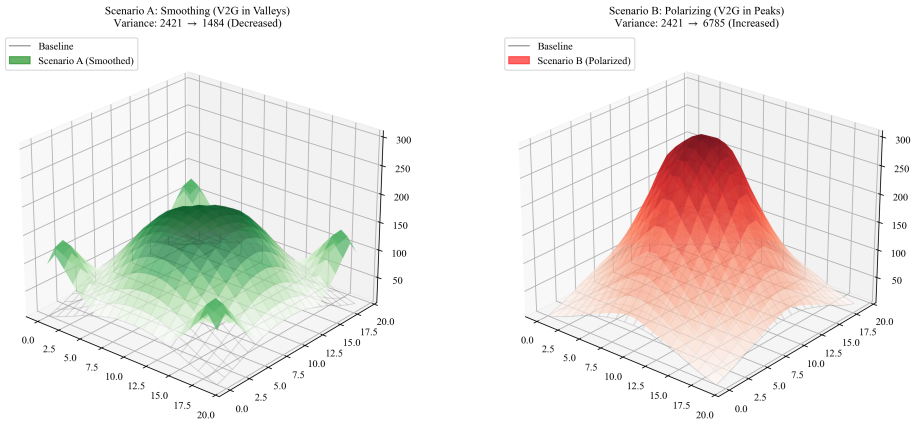


Figure 4: Impact of V2G deployment location

Macroscopic impact of V2G infrastructure

Finally, we evaluate the macroscopic impact of V2G infrastructure deployment on the spatial value landscape, quantified by the spatial variance of the value function. Strategic placement of V2G stations acts as a topological lever to reshape spatial equity.
According to Figure 4, we can find two different impacts of V2G infrastructure:

- Scenario A (smoothing strategy): Deploying V2G stations exclusively in low-demand peripheral areas provides a supplementary revenue stream that lifts the value floor of these valleys. This significantly reduces spatial variance (from 2421 to 1484), fostering a balanced fleet distribution and inducing a decentralized self-rebalancing effect.
- Scenario B (polarizing strategy): Co-locating V2G stations with high-demand mobility hotspots exacerbates spatial inequality. Vehicles in the city center enjoy dual revenue streams, drastically increasing spatial variance (from 2421 to 6785). This creates a walled city effect where the optimal pricing strategy imposes heavy penalties on trips leaving the center, economically locking the fleet in the core and degrading service in underserved regions.

4 Conclusions

This study establishes a comprehensive theoretical framework for the spatial equilibrium of decentralized, V2G-enabled electric autonomous vehicle (EAV) ride-sourcing fleets. By formulating the problem as a mean field game (MFG) with a continuous spatial-energy state space, we internalize the fundamental operational trade-off between spatial displacement for mobility services and

temporal retention for energy arbitrage. Theoretically, we proved that the optimal decentralized policy for each independent EAV follows a unique, spatially dependent threshold structure. This threshold structure governs the emergent self-regulation of the entire fleet, thereby eliminating the need for computationally prohibitive centralized dispatching

Methodologically, the development of the Structure-Preserving Threshold Fictitious Play (SP-TFP) algorithm successfully exploits these geometric properties to achieve robust convergence, overcoming the instability and limit cycles often observed in complex congestion games. Crucially, our numerical experiments reveal that V2G infrastructure is not merely a passive economic value-add, but an active instrument for spatial reshaping. For platform operators and urban policymakers, this demonstrates that the strategic placement of V2G stations can serve as a topological lever to enhance service equity and induce fleet self-rebalancing in underserved areas.

References

Ameli, M., Faradonbeh, M. S. S., Lebacque, J.-P., Abouee-Mehrizi, H., & Leclercq, L. (2022). Departure time choice models in urban transportation systems based on mean field games. *Transport. Sci.*, *56*(6), 1483–1504.

Cai, Z., Mo, D., Geng, M., Tang, W., & Chen, X. M. (2023). Integrating ride-sourcing with electric vehicle charging under mixed fleets and differentiated services. *Transport. Res. E-Log.*, *169*, 102965.

Gao, Y., Liu, W., Zhang, Z., & Du, B. (2025). Managing parking and charging services for electric autonomous vehicles considering vehicle-to-grid facilities. *Transport. Res. A-Pol.*, *201*, 104680.

Guille, C., & Gross, G. (2009). A conceptual framework for the vehicle-to-grid (v2g) implementation. *Energ. Policy*, *37*(11), 4379–4390.

Ke, J., Cen, X., Yang, H., Chen, X., & Ye, J. (2019). Modelling drivers' working and recharging schedules in a ride-sourcing market with electric vehicles and gasoline vehicles. *Transport. Res. E-Log.*, *125*, 160–180.

Lai, Z., Shang, Y., Dong, T., & Li, S. (2025). Integrating valet charging with mobile discharging services: A strengthened integer l-shaped method. *Transport. Sci.*.

Wei, B., Du, B., Saberi, M., Waller, S. T., & Liu, W. (2023). On ride-sourcing services of electric vehicles considering cruising for charging and parking. *Transport. Res. D-Tr. E.*, *118*, 103716.

Wu, M., Yin, Y., & Lynch, J. P. (2025). Multiday user equilibrium with strategic commuters. *Transport. Sci.*, *59*(2), 413–432.

Yan, Y., Lin, X., He, F., & Wang, D. Z. (2024). Dynamic operations of a mobile charging crowdsourcing platform. *Transport. Sci.*, *58*(5), 995–1015.

Yang, X., Abdin, A., & Puchinger, J. (2024). Optimal management of coupled shared autonomous electric vehicles and power grids: Potential of renewable energy integration. *Transport. Res. C-Emer.*, *165*, 104726.

Zhang, Z., Zhang, F., & Liu, W. (2025). Economic analysis of parking, vehicle charging and vehicle-to-grid services in the era of electric vehicles. *Transport. Res. B-Meth.*, *191*, 103133.

Zhu, Z., Ke, J., & Wang, H. (2021). A mean-field markov decision process model for spatial-temporal subsidies in ride-sourcing markets. *Transport. Res. B-Meth.*, *150*, 540–565.