

CONSENT: A Negotiation Framework for Leveraging User Flexibility in Vehicle-to-Building Charging under Uncertainty

Anonymous Author(s)

ABSTRACT

The growth of Electric Vehicles (EVs) creates a conflict in vehicle-to-building (V2B) settings between building operators, who face high energy costs from uncoordinated charging, and drivers, who prioritize convenience and a full charge. To resolve this, we propose a negotiation-based framework that, by design, guarantees voluntary participation, strategy-proofness, and budget feasibility. It transforms EV charging into a strategic resource by offering drivers a range of incentive-backed options for bounded flexibility in their departure time or requested state of charge (SoC). Our framework is calibrated using user survey data and validated with real operational data from a commercial building at the Nissan Advanced Technology Center – Silicon Valley. Simulations show that our negotiation protocol creates a mutually beneficial outcome: lowering the building operator’s costs by over 3.5% compared to an optimized, non-negotiating smart charging policy, while simultaneously reducing user charging expenses by 22% below the utility’s retail energy rate. By aligning operator and EV user objectives, our framework provides a strategic bridge between energy and mobility systems, transforming EV charging from a source of operational friction into a platform for collaboration and shared savings.

KEYWORDS

User Negotiation; Mechanism Design; Sequential Decision-Making; Model Predictive Control (MPC); EV Charging Optimization; Vehicle-to-Building (V2B); Smart Grid

ACM Reference Format:

Anonymous Author(s). 2026. CONSENT: A Negotiation Framework for Leveraging User Flexibility in Vehicle-to-Building Charging under Uncertainty. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 25 pages. <https://doi.org/10.65109/>

1 INTRODUCTION

The rapid growth of electric vehicles (EVs), fueled by sustainability goals and advances in battery technology [2, 28], presents a significant challenge for energy management. When EV charging is uncoordinated, it can trigger sharp demand spikes, inflate energy costs, and threaten grid stability [14]. This has spurred research into coordinated charging, particularly in the vehicle-to-building (V2B) context, where operators can leverage parked EVs for peak shaving [27]. However, creating an effective V2B system requires navigating a set of deeply interconnected challenges that span both technical control and human behavior.

At the core of the control problem lies the interplay between uncertainty and long-term rewards [13]. A charging controller must

operate with incomplete knowledge of future building loads and EV arrivals, yet its decisions are judged against long-term objectives like minimizing the monthly peak demand charge, a sparse and delayed reward. This makes myopic, reactive control insufficient. This technical challenge is further complicated by the need for incentive alignment [5]. The operator’s goal of deferring charging to minimize electricity costs, driven by both energy prices and monthly peak-power demand charges, directly conflicts with the user’s desire for convenience and a full battery [10]. A negotiation must therefore quantify the value of a user’s flexibility; however, this value is itself a function of an uncertain future. Finally, any negotiation is vulnerable to strategic misreporting, where users may exaggerate their needs to extract larger incentives, undermining the system’s integrity. Addressing these issues requires more than just efficient scheduling; it demands a strategic bridge between the building’s stationary energy system and the dynamic mobility needs of its users. While prior work addresses these challenges in isolation, a key challenge remains: the co-optimization of the real-time charging control policy with a strategic, user-facing negotiation policy. This requires a framework that both makes optimal control decisions and generates fair, personalized incentives. While methods like Reinforcement Learning (RL) have shown promise for handling uncertainty, the co-optimization task particularly benefits from approaches that can explicitly handle hard constraints and adapt to new scenarios without extensive re-training. For this reason, we contend that the explicit constraint handling and computational efficiency of Model Predictive Control (MPC) provide a more direct and scalable solution [15, 18, 20]. An MPC-based approach offers robust uncertainty handling through sampling, which we show is comparable to an RL baseline, establishing MPC as the ideal engine for our framework.

Leveraging this MPC foundation, we introduce our negotiation framework, CONSENT, which directly addresses the shortcomings of prior work [5, 9, 13]. To achieve incentive alignment, users are presented with tailored incentives for voluntary adjustments to their requested state-of-charge (SoC) or departure time, grounded in empirical survey data. To prevent strategic misreporting, our mechanism includes formal guarantees of strategy-proofness, budget feasibility, and voluntary participation. By integrating this mechanism and validating with real-world data, we demonstrate that operator savings can be aligned with high user participation and satisfaction. Our contributions include:

- A strategy-proof, budget-feasible negotiation framework for V2B charging that co-optimizes building costs and user incentives via a novel method for generating and evaluating personalized options under uncertainty and encourages voluntary participation (Sections 3 and 4).
- An empirical user study and behavioral model quantifying the trade-offs between charging convenience and monetary incentives in realistic scenarios (Section 5).

- A comprehensive evaluation using real-world operational data from a commercial building at the Nissan Advanced Technology Center – Silicon Valley, with sensitivity analyses demonstrating robustness and effectiveness (Section 6).

2 RELATED RESEARCH

Research in EV charging systems has largely advanced along two parallel streams. The first focuses on the charging control problem, where the goal is to optimize charging schedules to minimize operator costs like peak demand [7, 10]. A wide array of methods has been developed, including rule-based heuristics and mathematical programming [17, 23, 25]. More recent advances use Model Predictive Control (MPC) for real-time optimization, effectively adapting to dynamic loads, renewables, and prices [8, 29]. However, this entire stream of research, including stochastic MPC [20], typically solves the problem from the operator’s perspective by assuming full system control, thereby overlooking the critical challenge of securing voluntary user participation [19, 26].

Recognizing that user compliance cannot be guaranteed, the second stream of research focuses on the socio-technical challenge of user negotiation and cooperation. This work has produced incentive-based frameworks like menu-based pricing [9] and iterative negotiation [12, 27] to provide users with choices. These are often informed by behavioral models that predict user responses to inconvenience and rewards [21], with a key finding being the importance of user heterogeneity [4, 31]. Despite being user-centric, these frameworks are often designed for large-scale aggregators and are rarely co-designed with the real-time operational constraints of a single building.

The limitations of these separate streams highlight a significant gap at their intersection: the co-optimization problem, which requires strategically coupling real-time control with empirically-grounded, user-centric negotiation. Tackling this frontier requires advanced methods capable of handling uncertainty and complex decisions. Reinforcement Learning (RL) has emerged as a powerful tool for learning control policies in such stochastic environments [13, 19]. However, for the specific co-optimization challenge, MPC presents a more suitable approach for this specific challenge. While RL excels in exploration, the co-optimization problem benefits from MPC’s ability to explicitly handle hard constraints (e.g., user choices, charger limits) and efficiently evaluate hypothetical negotiation scenarios. By using Monte-Carlo sampling, MPC can robustly manage uncertainty, establishing it as a powerful engine for a framework that must bridge control and negotiation. Accordingly, our work addresses this co-optimization gap. We develop an MPC-based negotiation framework that operates directly on personalized user requests and is calibrated with empirical survey data, enabling voluntary and incentive-compatible cooperation at the building-user interface.

3 PROBLEM FORMULATION

We consider a V2B scenario where parked EVs can charge and/or discharge. The EVs follow an unknown, but regular and typically predictable pattern of arrivals, departures, and energy usage profiles. The system performs two tasks: (i) offers users cost-saving alternatives by suggesting bounded adjustments to their requested

SoC and departure time; (ii) minimizes total costs by optimally scheduling EV charging/discharging under time-varying energy tariffs (\$/kWh), demand tariffs (\$/kW), charger constraints, and user-requested SoC and departure times. All EVs depart by the end of each day. A notation table and additional details are provided in the supplementary material.

For each user, the process can be divided into three parts: **(i) Charger Allocation:** when an EV arrives, it is assigned a charger if one is available. Otherwise, the user cannot plug in and participate. **(ii) Charging Session Negotiation:** after plugging in, each user submits a charging session request with departure time and SoC. Based on the user’s request, the system provides multiple options for target SoC and departure time, each with an associated charging cost. The user selects one of these options to begin the charging session. **(iii) Charging Actions:** During each session, the system sets charger power at every time step to minimize total cost while meeting all negotiated requirements.

3.1 Specifications

Time Steps. The optimization horizon is discretized into a finite set of time steps \mathcal{T} , each τ hours long. Charging decisions are taken at the start of each time step and held constant for the duration.

Users and their EVs. Let \mathcal{V} be the set of EV users, each associated with a single electric vehicle (EV), and let v stand for both the user and their EV. Each user v has an arrival time t_{arr}^v and a requested departure time t_{req}^v . The full set of time steps when v is connected is $\mathcal{T}^v \subseteq \mathcal{T}$. The EV has a battery capacity bounded by $[e_{\text{min}}^v, e_{\text{max}}^v]$, arrives with initial state of charge e_{arr}^v , and requests a target SoC at departure e_{req}^v , and e_t provide the SoC at time t . The departure SoC may be revised if the negotiation leads to a different agreement.

User Choices. For each user $v \in \mathcal{V}$, their choice $\theta_v \in \Theta$ records the arrival time t_{arr}^v , arrival SoC e_{arr}^v , negotiated departure time $\theta_v^{t_{\text{dep}}}$ and negotiated target SoC $\theta_v^{e_{\text{dep}}}$, which may be adjusted during the negotiation process, based on their requested departure time t_{dep}^v and e_{req}^v respectively. The set of all user choices within the optimization horizon is represented using θ .

Negotiation Options & Flexibility Limits. User flexibility is parameterized by allowable deviations in target SoC, Δe_l^{max} , and departure time, $\Delta t_{\text{dep},l}^{\text{max}}$, where $l \in \{0, 1, \dots, L\}$ indexes discrete flexibility levels. Each flexibility level represents a distinct negotiation option offered to users, with level $l = 0$ corresponding to no deviation from the user’s originally requested departure SoC and time. The flexibility limits are determined from survey responses.

Building’s Power Draw and Energy Use. Let $p_t^b \in \mathcal{P}^b$ denote the building’s instantaneous power draw (kW) at time t . The total power draw $p_t \in \mathcal{P}$ includes both the building load and EV chargers, and the demand cost is computed from p_t . The building’s energy cost at time t is calculated as $p_t^b \cdot \tau$.

Chargers. We consider heterogeneous charger types, differing in control (controlled or uncontrolled), directionality (uni- or bidirectional), and charging rate (Level 1 or 2). Controlled chargers can be switched on/off; uncontrolled chargers draw power whenever connected to an EV. Unidirectional chargers only charge, while bidirectional chargers can also discharge. Let C denote all charger types with c_k units of type k . Charging rates range within $[c_{\text{min}}, c_{\text{max}}]$, where $c_{\text{min}} = 0$ for unidirectional and $c_{\text{min}} < 0$ for bidirectional

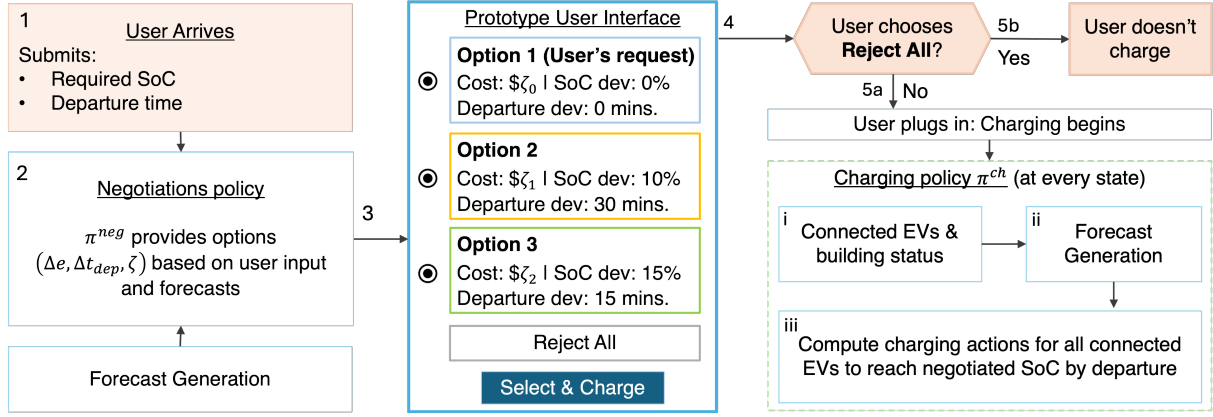


Figure 1: Negotiation workflow: Forecast-based offers on arrival; accepted offers optimally scheduled for negotiated SoC.

chargers. Charging/discharging efficiency η models electrical losses as the ratio of output power to input power.

Charging Rates. The charging rate r_t^v (in kWh) for EV v at time t is a decision variable.

Battery Degradation. The penalty term K_{batt} is applied to any discharging ($r_t^v < 0$) to model battery degradation costs.

Missing SoC Cost. For each EV, departing below the negotiated SoC incurs a penalty of K_{SOC} (\$/kWh).

Electricity Cost. Electricity cost comprises two components: time-varying energy charges $w_t \in \mathcal{W}$ (in \$/kWh) and demand charges K_{DC} (in \$/kW). At each time step t , total energy consumption includes the building load and EV charging. The energy cost at time t is $g_t = w_t (p_t^b \tau + \sum_{v \in \mathcal{V}} r_t^v)$, where p_t^b (kW) is the building load, τ (hr) is the time step duration, and r_t^v (kWh) is the energy delivered to EV v . The energy cost attributed to EV v sums over its connection time as $g_t^v = \sum_{t \in \mathcal{T}^v} w_t r_t^v$.

The building's demand cost reflects the peak power usage (kW) over a billing cycle and can contribute significantly to the total electricity costs, and are typically calculated over certain peak time periods $t \in \mathcal{T}^{\text{peak}} \subseteq \mathcal{T}$. Peak times also incur higher energy charges, incentivizing the shift of energy consumption to off-peak times. Demand cost is computed as $K_{\text{DC}} \cdot p^{\text{max}}$, where $p^{\text{max}} = \max_{t \in \mathcal{T}^{\text{peak}}} (b_t^p + \frac{1}{\tau} \sum_{v \in \mathcal{V}} r_t^v)$ is the maximum total power drawn at any time step in the peak periods. Demand charges are non-additive and require non-myopic optimization.

Building's Policies. Let π^{neg} and π^{ch} denote the negotiation and charging policies, respectively, and define the combined policy as $\pi = (\pi^{\text{neg}}, \pi^{\text{ch}})$. The building's optimal policy $\pi^* = (\pi^{\text{neg},*}, \pi^{\text{ch},*})$ jointly determines negotiation offers and EV charging schedules to minimize total costs, including operational expenses and user incentives, at each decision point.

Total Cost. Our objective is to minimize the building's net expenditure over the billing period under negotiation and charging policies π , which are controlled by the realized user choices θ . For convenience, we omit explicit notation of this dependence, writing

the cost function as

$$J^\pi(\theta | \mathcal{W}) = \sum_{t \in \mathcal{T}} g_t + K_{\text{DC}} \cdot p^{\text{max}} + K_{\text{SOC}} \sum_{v \in \mathcal{V}} (e_{\text{req}}^v - e_{\text{dep}}^v) + K_{\text{batt}} \sum_{v \in \mathcal{V}} \sum_{t \in \mathcal{T}^v} q_t^v \quad (1)$$

where all control parameters and actions (e.g., $g_t, p^{\text{max}}, q_t^v, e_{\text{dep}}^v$) are implicitly determined by the charging policy π , itself governed by the negotiated user choices θ . We drop these indices for clarity and readability.

User's Marginal Utility. θ_v denotes user v 's arrival and departure choices. Let $\theta_{\setminus v}$ be the realized choices of all users who have arrived prior to v . For users arriving after v , let $\hat{\theta}_{\setminus v}$ be random variables with joint distribution ψ , representing all possible future arrivals along with their associated choices and flexibilities. Instead of point forecasts, the controller samples from ψ to estimate expected system costs under uncertainty, as ψ encompasses the stochastic processes of EV arrivals and user decisions (e.g., arrival/departure times and flexibility preferences).

Given the optimal policy π^* (including negotiation and charging decisions), user v 's marginal utility is defined as the expectation under ψ of the cost reduction enabled by the user's participation:

$$U^{\theta_{\setminus v}, \pi^*}(\theta_v) = \mathbb{E}_{\hat{\theta}_{\setminus v} \sim \psi} \left[J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v} | \mathcal{W}') - J^{\pi^*}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v} | \mathcal{W}') \right] \quad (2)$$

where the expectation integrates over all future scenarios varying in arrival numbers, user types, and their flexibility parameters. We assume the quantity \mathcal{W}' represents building loads, energy charges, and demand charge parameters. If $U(\theta_v) > 0$, then serving user v is expected to reduce net building costs (e.g., via peak shaving or load shifting); otherwise, if $U(\theta_v) < 0$, the user's participation is expected to increase costs.

User Cost. The charging cost for user v , incorporating incentives, is denoted by ζ^v . It depends on the user's marginal utility; the following sections outline its computation.

Solution Space. A feasible solution to the V2B charging and negotiation problem consists of:

- **Assignment:** Each arriving EV $v \in \mathcal{V}$ is assigned to a single charger $c \in \mathcal{C}$ for its entire stay.

- **Negotiated User Requirements:** For each user, the system may propose adjustments to their requested SoC and departure time in exchange for an incentive, providing multiple choices of $(e_{\text{req}}^v, t_{\text{dep}}^v, \zeta^v)$ to choose from.
- **Charging/Discharging Schedule:** For each user’s EV v and each time step $t \in \mathcal{T}^v$, the charging policy π^{ch} provides the charging rate r_t^v from the feasible set determined by the assigned charger’s capabilities and efficiency. At every time step, the aggregate building and charging load must not violate power, charger, or infrastructure constraints.

The solution space includes all charger assignments, charging and discharging schedules, negotiated user requirements, user utilities that jointly satisfy system constraints, user participation guarantees, and cost minimization objectives.

4 APPROACH

Due to the stochastic nature of EV arrivals and the complexity of the V2B charging optimization problem, we decompose the task into two subproblems: (i) EV-to-charger assignment for newly arriving EVs and (ii) formulating the user negotiation options.

4.1 EV-Charger Assignment & Charging Policy

Whenever an EV arrives, the system assigns it to an available charger using a first-come, first-served (FCFS) policy. This policy prioritizes higher-power bidirectional chargers and breaks ties by earlier departure times, ensuring fairness among users. As shown in the EV to charger assignment table in the supplementary material, FCFS offers a simple yet effective heuristic compared to alternative strategies. After the assignment, the system formulates user negotiation options.

The negotiation module presents multiple charging options to each user, each with distinct flexibility levels and corresponding user costs to encourage adaptive behavior (detailed later). Based on the user’s selected required SoC and departure time, the system applies an EV charging policy π^{ch} to optimize charging rates for long-term building cost minimization.

The charging policy π^{ch} employs a Monte Carlo Model Predictive Control (MC-MPC) approach that solves a mixed-integer linear program (MILP) over a short-horizon window (e.g., one day). To handle uncertainty, the MILP optimization incorporates multiple Monte Carlo-sampled scenarios of future EV arrivals and building loads. It captures key system dynamics, including state-of-charge (SoC) evolution, while minimizing a total building cost function comprising energy costs, long-term demand charges, and penalties for unmet charging requirements. This rolling-horizon formulation enables robust, near-optimal decisions that dynamically adapt to real-world stochastic conditions. The detailed formulation of the charging policy is provided in the supplementary material.

4.2 Formulating the User Negotiation Options

We model the EV user negotiation process as a sequential decision-making framework triggered by each new user arrival, as illustrated in Fig. 1. The objective is to design a negotiation policy that interacts with users based on their requested charging requirements and departure times, offering alternative options with associated costs to encourage more flexible charging behaviors. This process operates

alongside an underlying charging policy, π^{ch} , which determines each charger’s rate at fixed intervals (e.g., every 15 minutes). The negotiation policy is computed through the following three steps. **Step 1: Option Generation.** Upon each user arrival, the system generates L charging options, each defined by distinct flexibility bounds and corresponding costs.

Step 2: Value Estimation. For each option, a forward simulation estimates the expected user cost, considering building load forecasts, previously arrived EVs, and predicted future arrivals. This evaluation accounts for all possible deviations within the option’s specified limits over the decision horizon.

Step 3: Option Provision. The system then presents the user with the option that maximizes the expected reward within the allowed deviation range, balancing building cost savings, incentive payments, and user inconvenience.

We formulate user negotiation as a Semi-Markov Decision Process (SMDP): $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r)$, where decision epochs occur at the arrival of each new EV. Let t_k denote the arrival time of the k -th EV, and $\Delta t_k = t_{k+1} - t_k$ the random wait until the next EV arrives.

State. At the arrival of EV v during the k -th decision epoch, the system state s_k is defined as

$$s_k = (t_k, p_{t_k}^b, \{e_{t_k}^v \mid v \in \mathcal{V}_k\}, (e_{\text{arr}}^v, e_{\text{req}}^v, t_{\text{req}}^v), p_{\text{max}}^{\text{past}})$$

where t_k is the current time step, $p_{t_k}^b$ is the building’s load at time t_k , $\{e_{t_k}^v \mid v \in \mathcal{V}_k\}$ represents the states of charge of all EVs currently connected, and $(e_{\text{arr}}^v, e_{\text{req}}^v, t_{\text{req}}^v)$ are the newly arrived EV’s initial SoC, requested SoC, and requested departure time, respectively. The highest power draw till epoch k is recorded in $p_{\text{max}}^{\text{past}}$.

Action. At the k -th decision epoch, when a new user v arrives in state s_k , the overall action set $\mathcal{A}(s_k)$ comprises multiple candidate charging options organized by flexibility levels $l \in \{1, \dots, L\}$. Each flexibility level defines a range of allowable deviations in target SoC and departure time. For each flexibility level l , there exists a subset of actions $\mathcal{A}_l = \{a_i = (\Delta e_i, \Delta t_{\text{dep},i}, \zeta_i)\} \subseteq \mathcal{A}(s_k)$, where the specific deviations $\Delta e_i \leq \Delta e_l^{\text{max}}$ and $\Delta t_{\text{dep},i} \leq \Delta t_{\text{dep},l}^{\text{max}}$ vary within the l -th deviation limits, and the associated user cost ζ_i depends on the current state s_k and electricity pricing. The complete action space is the union of all such subsets, $\mathcal{A}(s_k) = \bigcup_{l=1}^L \mathcal{A}_l$. Candidate actions are evaluated via the reward function to select the one that best balances user deviations and charging costs per deviation limit, yielding L negotiation options.

Within the L options, the action set always includes a “No deviation” option ($l = 0$) which maintains the user’s exact SoC and departure time requests, and costs ζ_0 . In addition to the L options, a “Reject all” choice lets the user opt out of facility charging and instead charge externally (e.g., at home or public stations). It allows flexibility to reject offers with unfavorable costs or deviations.

State Transition. The state transition function $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ models system evolution as users respond to charging offers and charging unfolds. User requests update upon choice (e.g., modified SoC or departure), and connected EVs are charged per the policy π^{ch} , which adapts rates to changing conditions. Details of the charging policy π^{ch} are in the supplementary material.

As EVs depart on schedule, chargers free up, and the system state updates to reflect the connected EVs and their SoCs, charger availability, building load, and any new arrivals, enabling accurate

tracking of charging and user decisions. The SMDP state transition kernel $\mathcal{P}(s_{k+1}, \Delta t_k \mid s_k, a_k; \pi^{ch})$ combines: (i) the stochastic EV inter-arrival distribution; and (ii) the deterministic evolution under charging policy π^{ch} during Δt_k . We do not explicitly define the state-action transition function, as our simulator leverages data-driven generative models trained on actual operational records [24].

Reward Function. At each decision epoch k , the reward for action $a_i \in \mathcal{A}_l \subseteq \mathcal{A}(s_k)$ is the marginal utility U associated with a_i in the negotiation option $l \in L$. This utility quantifies the impact of the chosen action on the total system cost, conditioned on the fixed choices of other connected EVs $\theta_{\setminus v}$ and the predicted behavior of future users $\hat{\theta}_{\setminus v}$. Formally, $r(s_k, a_i) = U^{\theta_{\setminus v}, \pi^*}(a_i)$, where π^* denotes the joint optimal negotiation and charging policy. The reward represents the utility of user v , and through this user-centric metric, it facilitates principled optimization of the negotiation policy in the sequential decision-making process.

4.3 Negotiation Policy

We operationalize the negotiation problem by solving the negotiation SMDP using an online, finite-horizon optimization at each decision epoch k (i.e., the arrival of a new EV).

User Cost: For each user v , the cost ζ^v is the energy cost (based on time-of-use price w_t and consumption) minus the utility gained from selecting the negotiation option.

$$\zeta^v = \sum_{t \in \mathcal{T}^v} g_t^v w_t - \alpha \cdot U(\theta_v) \quad (3)$$

where g_t^v is the energy charged to the EV v at time t (charging is positive energy use, and discharging is negative energy use), w_t is the electricity price at time t , and $\sum_{t \in \mathcal{T}^v} g_t^v w_t$ is the user's energy cost to be paid to the power utility. The user's utility $U(\theta_v)$ can be scaled by α to set the share of utility passed from the building operator to the user.

Our negotiation engine selects incentive-backed charging offers for each arriving by integrating a Monte Carlo sampling based model predictive control (MC-MPC). Our approach belongs to the family of stochastic Model Predictive Control (SMPC) methods [16]. Algorithm 1 presents the user negotiation process.

MC-MPC Controller with Deviation Limits. The MC-MPC controller implements the negotiation policy $\pi^{neg} : (s_k, \mathcal{F}) \mapsto \mathcal{A}(s_k)$, where s_k is the current state and \mathcal{F} denotes a set of Monte Carlo-sampled future scenarios. Each scenario simulates anticipated EV arrivals, departures, and building loads over the prediction horizon. To account for user flexibility, the controller systematically evaluates each flexibility level $l \in \{0, 1, \dots, L\}$, where each l specifies allowable bounds for SoC and departure time deviations.

For each flexibility level l , the controller explores the corresponding continuous action subset \mathcal{A}_l (from the SMDP) and solves a Mixed Integer Linear Program (MILP) implementing a Sample Average Approximation [20]. The MILP minimizes the average total cost over the sampled scenarios \mathcal{F} , yielding an optimal candidate

action a_l^* for each l :

$$a_l^* = \min_{\mathcal{A}_l} \frac{1}{|\mathcal{F}|} \sum_{f \in \mathcal{F}} \left[\sum_{i=0}^{H(f)} g_{i,f} + K_{DC} \cdot p_f^{\max} + K_{SOC} \sum_{v \in \mathcal{V}^f} z_f^v + K_{batt} \sum_{v \in \mathcal{V}} \sum_{t \in \mathcal{T}^v} q_{t,f}^v \right] \quad (4)$$

where for each sample f , $g_{i,f}$ is the energy cost, $p_{\max,f}$ is the peak power draw, z_f^v is any unmet SoC by departure for EV v , and $q_{t,f}^v$ is the sum of energy discharged.

The solution is subject to user deviation constraints corresponding to level l , operational constraints, and realized or fixed past user choices. The policy also offers the additional *Reject all* choice, which uses a fixed market-rate price of the nearest available charger. This approach yields robust options that balance user flexibility and cost minimization under uncertainty about future system dynamics.

4.4 Theoretical Guarantees

THEOREM 1 (STRATEGY-PROOFNESS FOR DEPARTURE TIME AND REQUESTED SoC). *The mechanism is strategy-proof with respect to departure time and requested SoC: a user cannot increase their utility by misreporting an earlier departure time or a higher required SoC $\hat{\theta}_v$ than their true preferences θ_v . Formally, truthful reporting always maximizes user utility, i.e., $U(\theta_v) \geq U(\hat{\theta}_v)$.*

THEOREM 2 (BUDGET FEASIBILITY). *Any payment ζ^v to user $v \in \mathcal{V}$ must satisfy $\zeta^v \geq -U(\theta_v)$ to ensure that the building's total cost after incentives does not exceed the baseline.*

THEOREM 3 (VOLUNTARY PARTICIPATION). *For all rational users v , the user choice mechanism guarantees $Y_1^v \geq 0$. Thus, no user is worse off by participating, and some benefit strictly if $U(\theta_v) > 0$.*

Proofs. Complete proofs of Theorems 1–3 are provided in Appendix D in the supplementary material.

5 SURVEY AND USER MODELING

To quantify user flexibility, we conducted an anonymized stated-preference survey with 28 university participants. Respondents specified both their maximum tolerable deviations in departure time and SoC, as well as the minimum compensation required (their perceived inconvenience cost) to accept specific deviation scenarios in departure (Δt_{dep}), and requested SoC (Δe).

(i) *Deviation limits for negotiation options:* Survey responses (see supplementary material) are used to group users into L negotiation options, indexed by l , each with calibrated upper bounds on SoC and departure time deviations, Δe_l^{\max} and Δt_l^{\max} . These bounds establish the menu of negotiation options that the system offers.

(ii) *User choice and inconvenience cost model:* Each survey respondent v provides data across j scenarios, each characterized by deviations Δe_j^v and $\Delta t_{\text{dep},j}^v$, along with a self-reported required discount, which we are terming as ‘‘inconvenience’’ I_j^v . These observations are used to estimate inconvenience sensitivities per respondent through the linear model, $I_j^v = w_1^v \Delta e_j^v + w_2^v \Delta t_{\text{dep},j}^v$. Thus, for no deviation, $I_j^v = 0$. To generalize across the population and reduce dimensionality, respondents v are clustered into N user types, indexed by n , each characterized by average inconvenience weights (w_1^n, w_2^n). For each

Algorithm 1: Online V2B Negotiation Policy

Input: Current state S_k ; number of future samples N ; number of flexibility levels L

Output: Sequence of negotiated choices $\Theta = \{\theta_1, \theta_2, \dots\}$

```
1 for each new EV arrival  $k = 1, 2, \dots$  do
2   Observe the current system state  $S_k$  at arrival time  $t_k$ ;
3   Future trajectory set  $\mathcal{F} \leftarrow \emptyset$ ;
4   for  $i = 1$  to  $N$  do
5     Sample a future trajectory  $f$  from the generative model
6     given  $S_k$ ;
7     Add  $f$  to  $\mathcal{F}$ ;
8   Option set  $\mathcal{L}_k \leftarrow \emptyset$ ;
9   for each flexibility level  $l = 0$  to  $L$  do
10    Solve the MILP (Eq. (4)) over scenarios  $\mathcal{F}$  to find the
11    optimal option  $a_l^*$  for level  $l$ ;
12    Add  $a_l^*$  to the option set  $\mathcal{L}_k$ ;
13 Present the option set  $\mathcal{L}_k$  (plus the 'Reject all' choice) to the
14 arriving user  $v_k$ ;
15 Receive the user's choice  $\theta_k \in \mathcal{L}_k \cup \{\text{reject}\}$ ;
16 Update  $v_k$ 's requirements based on the negotiated choice  $\theta_k$ ;
17 Apply the online charging policy (Algorithm in Appendix) for
18 all time steps between  $t_k$  and the next arrival  $t_{k+1}$  to evolve
19 the system to state  $S_{k+1}$ ;
```

user type n , the estimated inconvenience for negotiation option l is computed as $I_l^n = w_1^n \Delta e_l + w_2^n \Delta t_{\text{dep},l}$ where Δe_l and $\Delta t_{\text{dep},l}$ are the SoC and departure-time deviations for option l .

User acceptance during negotiation is modeled via a logit-based choice function [5]. The utility of acceptance for option l by user type n for charging cost ζ_l and given an external charging cost \bar{E} , is

$$Y_l^n = \left[\bar{E} - (\zeta_l + I_l^n) \right] \quad (5)$$

The probability that type n user accepts option l is $P_l^n = \frac{\exp(Y_l^n) + \epsilon}{\sum_{l=0}^L \exp(Y_l^n)}$, with a small noise term, ϵ , for numerical stability.

This framework enables the system to estimate acceptance likelihoods personalized by inferred user types and their characteristic inconvenience costs, allowing design and optimization of incentive-aware negotiation menus aligned with user preferences. Summary statistics of the survey and clustering results are in the supplementary material.

6 EXPERIMENTS

We evaluate our framework using a modular simulator [24], calibrated with operational data from a Nissan Advanced Technology Center – Silicon Valley building. It models multiple chargers, dynamic tariffs, stochastic EV arrivals and preferences, load forecasts, and charger constraints for online and ahead-of-time operation.

Hardware. Experiments ran on a 32-core, 6.2 GHz, 128 GB RAM machine using IBM CPLEX [11] for MILP.

Runtime and Complexity. Runtime scales as $O(L \cdot \mathcal{F} \cdot H \cdot \mathcal{V}^p)$, where L is the number of negotiation options, \mathcal{F} the number of future samples, H the prediction horizon, \mathcal{V} the number of EVs, and p the solver's empirical polynomial factor.

Data Collection. We use data from a commercial site in Santa Clara, California (May 2023 – November 2024), including building power usage, charger activity, and vehicle telemetry. All EVs used support bidirectional charging, and charging is modeled with a three-segment piecewise linear SoC curve. We sampled 100 such episodes, evenly split into 50 for future sampling and 50 for testing. Only weekdays are analyzed, as employee presence and demand charges are negligible on weekends. Figures 8 and 9 in the supplementary material show EV arrival/departure and SoC distributions. The site has controllable, 10 unidirectional (0–20 kW) and 5 bidirectional (± 20 kW) chargers. Silicon Valley Power's time-of-use and demand charge rates [22] define \mathcal{W} where peak hours (6 AM to 10 PM) are \$0.178/kWh and off-peak hours are \$0.137/kWh. Demand charge K_{DC} is \$11.67/kW. The penalty for missing SoC is set at $K_{\text{SOC}} = \$0.5/\text{kWh}$, and battery degradation penalty is set to $K_{\text{batt}} = \$0.05/\text{kWh}$ for discharging. The external charging cost is set to $\bar{E} = \$0.30/\text{kWh}$, reflecting the average residential EV charging rate in the area, including amortized charger installation costs. More details are provided in the supplementary material.

Evaluation Metrics. Each scenario is evaluated using the following metrics. **Building cost** measures the total monthly cost incurred by the operator. **User cost** is the average cost per kilowatt-hour delivered to users, including incentives. **Monthly user cost** denotes cumulative user expenditure over the month. **Reject (%)** indicates the fraction of users unable or unwilling to charge at the building. **Forecasting.** EV arrivals, departures, and SoC requests are forecasted with generative models trained on nearly two years of telemetry. The building operator confirmed EV schedules are independent of building load and past charging, allowing decoupled pre-sampling. Detailed EV behavior and building load models, with performance metrics, are reported in the supplementary material. For experimenting, we sample $\mathcal{F} = 10$ future samples for π^{neg} and $\mathcal{F} = 30$ for π^{ch} .

Negotiation Options. Based on the user survey, we define the negotiation options used in the experiment. Each user is presented with $L = 3$ alternatives in addition to their exact request ($l = 0$). Each option is represented as a tuple $(\Delta e_l^{\text{max}}, \Delta t_{\text{dep},l}^{\text{max}})$, indicating the maximum deviations in SoC and departure time. Specifically, option $l = 0$ is (0, 0), option $l = 1$ is (6.25%, 30 min), option $l = 2$ is (10%, 15 min), and option $l = 3$ is (20%, 105 min). An additional "Reject all" option allows users to opt out of charging at the building. **User Types.** Survey analysis (see supplementary material) identifies $N = 4$ user types, defined by time and SoC sensitivity. Without ground-truth labels, users are assigned types based on survey proportions.

Type	SoC Sensitivity (w_1)	Time Sensitivity (w_2)	Share
i	0.0489 ± 0.02	0.1250 ± 0.01	10.7%
ii	0.0133 ± 0.01	0.0346 ± 0.02	53.6%
iii	0.0362 ± 0.01	0.0673 ± 0.01	32.1%
iv	0.0000 ± 0.00	0.1083 ± 0.00	3.6%

Baseline Policies. We evaluate our framework against a series of progressively sophisticated baselines detailed in Table 1. The analysis begins with the **Building-Only** reference, which establishes the ideal operational cost without any EV load. We then model the common real-world scenario of **Uncoordinated Charging** [3], where EVs charge for free upon arrival. Building upon this, the

Policy Description	Building Cost (\$/month)	User Cost (\$/kWh)	Monthly User Cost (\$/month)	Reject (%)
<i>Reference Baseline</i>				
1 Building-Only Reference	8484 ± 2185	—	—	—
<i>Non-Negotiation Baselines</i>				
2 Uncoordinated Charging [3] + Free Cost	9005 ± 2184	0.00	—	—
3 Smart Charging + Free Cost [13]	8785 ± 2144	0.00	—	—
4 Smart charging + Fixed User Cost [22]	8457 ± 2144	0.178	328 ± 61	—
5 Smart Charging + Utility-based Cost	8413 ± 2183	0.183	372 ± 64	—
<i>Negotiation-Based Policies</i>				
6 Menu-based Negotiation [5, 9]	8479 ± 2153	0.157	309 ± 63	34.89
7 CONSENT (Our Approach)	8480 ± 2171	0.138	254 ± 55	21.99
<i>Ablation Study</i>				
8 CONSENT (without Reject Option)	8461 ± 2155	0.165	299 ± 60	—

Table 1: Comparative Evaluation of V2B Policies. The table progresses from a building-only reference to increasingly sophisticated baselines, culminating in our proposed negotiation framework (CONSENT). Lower values are better for all the metrics.

Smart Charging + Free Cost policy [13] isolates the benefit of centralized coordination by introducing an intelligent scheduling algorithm (π^{ch}) while keeping charging free for users. To introduce a simple economic model, we evaluate the **Smart Charging + Fixed Cost** policy, which sets the user price based on the local utility’s peak energy rate [22]. Advancing this, the **Smart Charging + Utility-based Cost** policy represents a purely operator-centric model that pairs the scheduling algorithm with our utility based pricing function (Eq. (3)) to maximize building savings. Finally, to benchmark against state-of-the-art interactive systems, we include the **Menu-based Negotiation** policy [5, 9], which introduces user choice and the ability to reject offers, serving as a direct precursor to our proposed framework.

7 RESULTS

We evaluate our framework, CONSENT, which integrates smart charging, utility-based cost generation tailored to user flexibility, and the option for users to reject offered choices. Its performance is compared against a series of progressively complex baseline policies, with results summarized in Table 1. For reference, the building’s operational cost without any EV charging load (Policy 1) is nearly \$8,484 per month, representing an ideal peak-shaving benchmark. **Baseline Policy Performance.** The real-world baseline of uncoordinated free charging (Policy 2) [3] raises the building’s monthly cost by over 6%. A smart scheduling algorithm (Policy 3) [13, 19] mitigates this increase but still costs more than the no EV reference. **Smart Charging + Fixed Cost** policy [22] further reduces the building’s cost, bringing it nearly to the ideal no EV baseline. Finally, using our utility-based, demand-aware user pricing (Policy 5) enables the building to achieve its lowest operational cost among all policies. However, this purely operator-centric optimization leads to the highest per-unit user cost, leading to negotiation-based approaches. **Negotiation Framework Performance.** Building on these baselines, our first negotiation-based comparison is Policy 5, which implements a state-of-the-art, menu-based approach consistent with prior work [5, 9]. While this method reduces user costs, rejection rate is high, at nearly 35%, indicating misaligned incentives.

Our proposed framework, CONSENT (Policy 6), strikes the best overall balance between operator and user objectives, and offers a **12% lower user cost** at \$0.138/kWh and substantially reduces

the rejection rate from 34.89% to **21.99%**, compared to Menu-based Negotiation. This user cost is over 61% lower than average public charging rates in the United States [1]. Simultaneously, the framework reduces building operator costs by 5.83% compared to uncoordinated charging and by nearly 3.5% compared to an optimized smart-charging policy with free user access (Policy 3).

To isolate the importance of user choice, we performed an ablation study by removing the reject option (Policy 7). The results show that while forcing participation slightly reduces the building’s cost to \$8461/month, it does so at the expense of a significantly higher user cost (\$0.165/kWh). This trade-off highlights the critical role of providing user autonomy; the reject option is essential for achieving the most favorable outcome for users. Generating all options for a user takes ~ 0.8 s on average, with each charging option computation requiring ~ 0.99 s. Our method is scalable over larger problem instances, as shown in Figure 2.




Figure 2: Comparison of computational runtime per decision across varying EV arrival intensities.

Sensitivity to User Flexibility Profiles. To evaluate the robustness of our framework, we conduct a sensitivity analysis, with results shown in Table 2. We test three policies against five simulated user populations, each representing a different flexibility profile: a baseline derived from our survey, populations that are generally more or less flexible, and populations specifically sensitive to delays in departure time or reductions in SoC. For our approach, CONSENT, a clear trade-off emerges. When interacting with a *More Flexible* user population (lower sensitivity), the building operator can achieve greater savings, reducing costs from \$8480 to \$8311 per month. This increased system flexibility also leads to more agreeable offers, lowering the rejection rate from 21.99% to 19.33%. However, the average user cost per kWh increases, as the

User Profile Scenario	Bldg. Cost (\$/mo)	User Cost (\$/kWh)	Monthly User Cost (\$/mo)	User Reject (%)
<i>CONSENT (with Reject option)</i>				
Baseline (Survey)	8480 ± 2171	0.138	254 ± 55	21.99
More Flexible (↓25% w_1, w_2)	8311 ± 2098	0.159	293 ± 63	19.33
Range-Sens. (↑50% w_1)	8311 ± 2120	0.156	287 ± 62	23.09
Time-Sens. (↑50% w_2)	8395 ± 2151	0.155	285 ± 61	22.87
Less Flexible (↑25% w_1, w_2)	8492 ± 2171	0.166	306 ± 66	22.65
<i>CONSENT without Reject option</i>				
Baseline (Survey)	8461 ± 2155	0.165	304 ± 50	—
More Flexible (↓25% w_1, w_2)	8349 ± 2100	0.195	359 ± 59	—
Range-Sens. (↑50% w_1)	8485 ± 2110	0.188	346 ± 57	—
Time-Sens. (↑50% w_2)	8412 ± 2111	0.190	350 ± 57	—
Less Flexible (↑25% w_1, w_2)	8503 ± 2180	0.201	370 ± 61	—
<i>Menu-based Negotiation [5, 9]</i>				
Baseline (Survey)	8479 ± 2153	0.157	289 ± 52	34.89
More Flexible (↓25% w_1, w_2)	8352 ± 2121	0.176	324 ± 58	33.50
Range-Sens. (↑50% w_1)	8501 ± 2119	0.181	333 ± 60	36.63
Time-Sens. (↑50% w_2)	8497 ± 2101	0.182	335 ± 60	36.81
Less Flexible (↑25% w_1, w_2)	8521 ± 2181	0.193	355 ± 64	36.46

Table 2: Sensitivity of negotiation outcomes to different user flexibility profiles where we compare our framework, CONSENT, against a menu-based and an ablated baseline.

system optimizes for its own objectives. Conversely, a *Less Flexible* population limits optimization opportunities, resulting in higher user costs and a slightly higher rejection rate. Notably, making users more sensitive to range (*Range-Sensitive*) provides more cost-saving potential for the building than making them sensitive to time (*Time-Sensitive*).

This analysis highlights the value of our framework. The ablation study, *Our Approach without Reject option*, shows that while the building can achieve even lower costs by forcing participation, it does so at the expense of consistently higher user costs across all profiles. Furthermore, compared to the *Menu-based Negotiation* baseline, our method consistently achieves a much lower rejection rate across all scenarios (e.g., 19.33% vs. 33.50% for more flexible users), demonstrating a superior alignment of incentives. These findings confirm that our approach is robust and effectively balances operator and user objectives across diverse user populations.

Utility Sharing	Building's Cost (\$/mo)	User's Cost (\$/kWh)	Reject(%)
100% sharing	8480.11 ± 2171.28	0.138 ± 0.03	21.99
90% sharing	8469.79 ± 2173.71	0.148 ± 0.02	22.14
75% sharing	8451.28 ± 2171.93	0.158 ± 0.02	23.78
50% sharing	8427.71 ± 2163.53	0.170 ± 0.01	24.06
25% sharing	8411.49 ± 2155.76	0.176 ± 0.01	25.33
10% sharing	8397.71 ± 2147.44	0.178 ± 0.00	26.88

Table 3: Performance across utility-sharing ratios α (100%–10%) from (3): lower sharing reduces building cost but increases user costs and simulated rejections, highlighting trade-offs in operational savings and user acceptance.

Sensitivity to Cost-Sharing Levels. Table 3 summarizes the negotiation framework’s performance across cost-sharing ratios from 100% to 10%, by tuning α in Eq. (3). As α decreases, building costs

fall, reaching peak savings at 10% sharing, while user costs, satisfaction, and participation also decline. Notably, even with 90% utility sharing, building costs drop below the building’s baseline load without EVs (Policy 1 in Table 1), while user costs remain lower than all other policies. This shows how strategic cost allocation drives both operational savings and user affordability, underscoring the mechanism’s robustness and adaptability.

8 CONCLUSION

We introduce CONSENT, a novel negotiation framework built on a stochastic Model Predictive Control (SMPC) foundation that acts as a strategic bridge between building energy systems and EV mobility needs. Empirically validated with real-world data, it co-optimizes operator and user objectives via a novel method to generate personalized options. This results in a 22% user cost reduction and a significantly lower rejection rate than negotiation baselines, all while achieving higher building savings than free to the user, purely operator-centric policies, successfully transforming V2B charging into a platform for collaboration.

Our strategy-proof guarantee ensures fairness in any single negotiation by making misreporting unprofitable. Future work could extend this concept from single-shot interactions to a dynamic, repeated-game setting. While our MPC-based framework adapts to new statistical patterns, a game-theoretic approach would model users as rational agents who strategically adapt their behavior over time. Developing mechanisms to anticipate these shifts and find a stable, efficient equilibrium is a key direction for future research.

REFERENCES

- [1] AAA. 10-09-2025. AAA Fuel Prices. <https://gasprices.aaa.com/ev-charging-prices/>
- [2] Ali Saadon Al-Ogaili, Tengku Juhana Tengku Hashim, Nur Azzammudin Rahmat, Agileswari K Ramasamy, Marayati Binti Marsadek, Mohammad Faisal, and Mahammad A Hannan. 2019. Review on scheduling, clustering, and forecasting strategies for controlling electric vehicle charging: Challenges and recommendations. *Ieee Access* 7 (2019), 128353–128371.
- [3] Mohammed Al-Saadi, Bartosz Patkowski, Maciej Zaremba, Agnieszka Karwat, Mateusz Pol, Lukasz Chelchowski, Joeri Van Mierlo, and Maitane Berecibar. 2021. Slow and fast charging solutions for li-ion batteries of electric heavy-duty vehicles with fleet management strategies. *Sustainability* 13, 19 (2021), 10639.
- [4] Natascia Andrenacci and Maria Pia Valentini. 2023. A literature review on the charging behaviour of private electric vehicles. *Applied Sciences* 13, 23 (2023).
- [5] Sangjae Bae, Teng Zeng, Bertrand Travacca, and Scott Moura. 2021. Inducing Human Behavior to Maximize Operation Performance at PEV Charging Station. *IEEE Transactions on Smart Grid* 12, 4 (2021), 3353–3363.
- [6] Mengmeng Cai, Manisa Pipattanasomporn, and Saifur Rahman. 2019. Day-ahead building-level load forecasts using deep learning vs. traditional time-series techniques. *Applied energy* 236 (2019), 1078–1088.
- [7] Kristoffer Christensen, Zheng Grace Ma, and Bo Nørregaard Jørgensen. 2025. A scoping review on electric vehicle charging strategies with a technical, social, and regulatory feasibility evaluation. *Renewable and Sustainable Energy Reviews* 211 (2025), 115300.
- [8] Alessandro Di Giorgio, Francesco Liberati, and Silvia Canale. 2014. Electric vehicles charging control in a smart grid: A model predictive control approach. *Control Engineering Practice* 22 (2014), 147–162.
- [9] Arnob Ghosh and Vaneet Aggarwal. 2017. Control of charging of electric vehicles through menu-based pricing. *IEEE Transactions on Smart Grid* 9, 6 (2017).
- [10] Zhanwei He, Javad Khazaei, and James D Freihaut. 2022. Optimal integration of Vehicle to Building (V2B) and Building to Vehicle (B2V) technologies for commercial buildings. *Sustainable Energy, Grids and Networks* 32 (2022), 100921.
- [11] IBM. [n.d.]. IBM CPLEX Python API Documentation.
- [12] Komal Khan, Islam El-Sayed, and Pablo Arboleya. 2022. Multi-issue negotiation evs charging mechanism in highly congested distribution networks. *IEEE Transactions on Vehicular Technology* 71, 6 (2022), 5743–5754.
- [13] Fangqi Liu, Rishav Sen, Jose Paolo Talusan, Ava Pettet, Aaron Kandel, Yoshinori Suzue, Ayan Mukhopadhyay, and Abhishek Dubey. 2025. Reinforcement Learning-based Approach for Vehicle-to-Building Charging with Heterogeneous

- Agents and Long Term Rewards. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems* (Detroit, MI, USA) (AAMAS '25). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1345–1353.
- [14] Weicheng Liu, Xujiang Shi, Jianfeng Zhao, Xiao-Ping Zhang, and Ying Xue. 2021. Electric Vehicle Charging Simulation Framework Considering Traffic, User, and Power Grid. *Journal of Modern Power Systems and Clean Energy* 9, 3 (2021).
- [15] Graham McClone, Avik Ghosh, Adil Khurram, Byron Washom, and Jan Kleissl. 2024. Hybrid Machine Learning Forecasting for Online MPC of Work Place Electric Vehicle Charging. *IEEE Transactions on Smart Grid* 15, 2 (2024).
- [16] Ali Mesbah. 2016. Stochastic Model Predictive Control: An Overview and Perspectives for Future Research. *IEEE Control Systems Magazine* 36, 6 (2016), 30–44. <https://doi.org/10.1109/MCS.2016.2602087>
- [17] Yorie Nakahira, Niangjun Chen, Lijun Chen, and Steven H Low. 2017. Smoothed least-laxity-first algorithm for EV charging. In *Proceedings of the Eighth International Conference on Future Energy Systems*. 242–251.
- [18] MI Saleem, S Saha, U Izhar, and L Ang. 2024. A stochastic MPC-based energy management system for integrating solar PV, battery storage, and EV charging in residential complexes. *Energy and Buildings* 325 (2024), 114993.
- [19] Rishav Sen, Yunuo Zhang, Fangqi Liu, Jose Paolo Talusan, Ava Pettet, Yoshinori Suzue, Ayan Mukhopadhyay, and Abhishek Dubey. 2025. Online Decision-Making Under Uncertainty for Vehicle-to-Building Systems. In *Proceedings of the ACM/IEEE 16th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2025)* (Irvine, CA, USA) (ICCPSS '25). Association for Computing Machinery, New York, NY, USA, Article 20, 12 pages. <https://doi.org/10.1145/3716550.3722024>
- [20] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. 2021. *Lectures on stochastic programming: modeling and theory*. SIAM.
- [21] Mahla Shariatzadeh, Marta AR Lopes, and Carlos Henggeler Antunes. 2025. Electric vehicle users' charging behavior: A review of influential factors, methods and modeling approaches. *Applied Energy* 396 (2025), 126167.
- [22] Silicon Valley Power. 21-07-2025. Commercial Rates and Fees. <https://www.siliconvalleypower.com/businesses/rates-and-fees>
- [23] John A. Stankovic, Krithi Ramamritham, and Marco Spuri. 1998. *Deadline Scheduling for Real-Time Systems: Edf and Related Algorithms*. Kluwer Academic Publishers, USA.
- [24] Jose Paolo Talusan, Rishav Sen, Ava Pettet, Aaron Kandel, Yoshinori Suzue, Liam Pedersen, Ayan Mukhopadhyay, and Abhishek Dubey. 2024. OPTIMUS: Discrete Event Simulator for Vehicle-to-Building Charging Optimization. In *2024 IEEE International Conference on Smart Computing (SMARTCOMP)*. <https://doi.org/10.1109/SMARTCOMP61445.2024.00050>
- [25] Kevin Tanguy, Maxime R Dubois, Karol Lina Lopez, and Christian Gagné. 2016. Optimization model and economic assessment of collaborative charging using Vehicle-to-Building. *Sustainable Cities and Society* 26 (2016), 496–506.
- [26] Jacob Thran, Jakub Marecek, Robert N Shorten, and Timothy C Green. 2025. Reserve Provision from Electric Vehicles: Aggregate Boundaries and Stochastic Model Predictive Control. *IEEE Transactions on Power Systems* (2025).
- [27] B. Wang et al. 2018. An EV Charging Scheduling Mechanism Based on Price Negotiation. *Future Internet* 10, 5 (2018), 40.
- [28] Ji Wu, Hao Su, Jinhao Meng, and Mingqiang Lin. 2023. Electric vehicle charging scheduling considering infrastructure constraints. *Energy* 278 (2023), 127806.
- [29] Lei Yang, Xinbo Geng, Xiaohong Guan, and Lang Tong. 2023. EV charging scheduling under demand charge: A block model predictive control approach. *IEEE Transactions on Automation Science and Engineering* 21, 2 (2023).
- [30] Baran Yildiz, Jose I Bilbao, and Alistair B Sproul. 2017. A review and analysis of regression and machine learning models on commercial building electricity load forecasting. *Renewable and Sustainable Energy Reviews* 73 (2017), 1104–1122.
- [31] Maximilian Zähringer, Teresa Junior, and Lennart Adenaw. 2024. Watt matters most—Survey data results of private passenger vehicle owners and commercial vehicle drivers. *Data in Brief* 52 (2024), 109942.

A APPENDIX

Assignment	Tie Breaker	Building Cost with Smart Charging (\$)
Bidirectional	Departure	8479.80±1154.23
Bidirectional	Capacity	8479.80±1155.43
Bidirectional	Random	8480.36±1155.18
Random	Random	8481.72±1170.43
Random	Departure	8482.87±1167.17
Random	Capacity	8483.46±1164.88
Unidirectional	Random	8583.07±1182.91
Unidirectional	Departure	8583.29±1182.85
Unidirectional	Capacity	8583.30±1182.84

Table 4: Charger Assignment and tiebreaker comparisons with an oracle, optimal policy (MILP). Assigning to Bidirectional chargers first and then breaking ties by assigning them to the EV that departs later results in the lowest bill. Lower is better.

B CHARGING POLICY

We formulate the V2B charging policy π^{ch} using Model Predictive Control (MPC), which utilizes a Monte-Carlo sampling and a Mixed-Integer Linear Program (MILP). The Monte-Carlo sampling provides a range of possible future trajectories in terms of EV arrivals and user requests, while the MILP aims to minimize the total cost while satisfying user charging requests and operational constraints. The model captures discrete charger assignments, models battery SoC dynamics using piecewise linear segments, and incorporates realistic electricity tariffs that include both time-of-use energy prices and demand charges. This formulation ensures that charging schedules respect user-specified arrival times, departure times, and requested SoC targets, while optimizing energy usage to flatten demand peaks and reduce total cost for the building operator.

The MPC framework enables real-time implementation in a dynamic environment with uncertain future arrivals and building loads. The MPC controller operates in a rolling horizon fashion: at each decision epoch (e.g., every 15 minutes), it solves the MILP using up-to-date forecasts of building load, EV arrivals, and user requests. By repeatedly re-optimizing at each time step as new information becomes available, the MPC policy adapts charging and discharging actions to minimize expected cumulative cost over the remaining horizon.

EV to Charger Assignment function. At each time step t , chargers and EVs are matched one-to-one: each charger serves at most one EV, and each EV connects to at most one charger based on the first-come-first-served (FCFS) assignment policy. The assignment policies are summarized in Table 4. We select the best-performing approach: Bidirectional-first with departure time as the tie-breaker. Formally, this defines a mapping $\phi_t : \mathcal{V}_t \leftrightarrow \mathcal{C}_t$, where \mathcal{V}_t and \mathcal{C}_t are the sets of EVs and chargers active at t . If $\phi_t(v) = c$, then $\phi_t^{-1}(c) = v$, respecting infrastructure constraints.

Charging Rates. In addition to the charging rate limits of the chargers, to capture the typical slowdown in charging as an EV battery fills, we model charging using piecewise-linear segments $\delta \in \Delta$. Each segment is defined by a maximum charging rate $r_{\max}^{v,\delta}$

valid over SoC intervals $[e_{\min}^{v,\delta}, e_{\max}^{v,\delta}]$. During discharging, these piece-wise functions are not considered, and the rate of discharge is considered to be linear at all battery levels.

Decision Variables

- $a_t^{v,c} \in \{0, 1\}$: Binary assignment variable indicating whether electric vehicle v is assigned to charger of type c at time step t .
- $y_t^{v,\delta} \in \{0, 1\}$: Indicator variable equal to 1 if the state of charge (SoC) of EV v lies within segment δ of the piecewise-linear SoC curve at time t .
- $x_t^{v,c,\delta} \in \{0, 1\}$: Binary variable coupling the SoC segment δ and charger assignment $a_t^{v,c}$ to regulate feasible charging rates.
- $r_t^v \in [c_{\min}, c_{\max}]$: Continuous variable indicating charging (or discharging) power for vehicle v at time step t , where bounds depend on charger directionality.
- $e_t^v \in [e_{\min}^v, e_{\max}^v]$: Continuous variable denoting the state-of-charge (SoC) of EV v at the start of time step t .
- $g_t \in [0, \infty)$: Energy cost incurred by building and charging load during time step t , constrained so that net energy recovered from vehicles never exceeds current building energy use.
- $q_t^v \geq 0$: Auxiliary continuous variable penalizing excessive discharging from vehicle v at time t .
- $z_v \geq 0$: Continuous variable tracking the SoC deficit at the time of vehicle v 's departure; used to penalize unmet charging targets and maintain user satisfaction.
- $p^{\max} \geq 0$: Peak building power draw, defined as the maximum instantaneous power consumption across all time steps $\max_t(p_t)$.

Constraints The following constraints are modeled the MILP formulation:

EV to Charger Assignment. For the duration of the EV v , $\mathcal{T}^v \subseteq \mathcal{T}$ we already assumed EV to charger assignment is done exogenously, and is captured by $a_{v,c,t}$.

To ensure that an EV remains continuously assigned to the same charger throughout its stay, we impose:

$$\forall v \in \mathcal{V}, \forall c \in \mathcal{C}, \forall t \in \mathcal{T}^v \setminus \{t_{\text{arr}}\} : a_{t-1}^{v,c} - a_t^{v,c} = 0 \quad (6)$$

Modeling Charging Dynamics via Piecewise Linear Curves. We model EV battery dynamics using piecewise-linear approximations of the charging curve, divided into Δ SoC segments. For each EV v , the SoC at any time must lie within exactly one segment which ranges between $[e_{\min}^{v,\delta}, e_{\max}^{v,\delta}]$ for $\delta \in \Delta$ segments, as indicated by $y_t^{v,\delta}$:

$$\forall v \in \mathcal{V}, \forall t \in \mathcal{T}^v, \forall \delta \in \Delta : \begin{cases} e_t^v \geq e_{\min}^{v,\delta} \cdot y_t^{v,\delta} - (1 - y_t^{v,\delta}) \\ e_t^v \leq e_{\max}^{v,\delta} \cdot y_t^{v,\delta} + M_1(1 - y_t^{v,\delta}) \end{cases} \quad (7)$$

where M_1 is a large constant, and we set it to e_{\max}^v to relax the bounds when a segment is inactive. Choosing a smaller value, such as choosing $M_1 = e^{v,\delta_{\max}}$ would cancel the indicator variable $y_t^{v,\delta}$ and invalidate the constraint; we therefore retain $M_1 = e_{\max}^v$.

Each EV must belong to exactly one linear segment at any time:

$$\forall v \in \mathcal{V}, \forall t \in \mathcal{T} : \sum_{\delta \in \Delta} y_t^{v,\delta} = 1 \quad (8)$$

We define auxiliary variables $x_t^{v,c,\delta}$ to link charger assignment $a_t^{v,c}$ and SoC linear segment $y_t^{v,\delta}$. These variables are used to limit

Variable	Description
$t \in \mathcal{T}, \mathcal{T}^{\text{peak}} \subseteq \mathcal{T}$	Set of discretized time steps of uniform length τ (e.g., 0.25 hr), Set of peak periods (for demand cost)
$v \in \mathcal{V}$	Set of EV users in the system
$[t_{arr}^v, t_{req}^v] \in \mathcal{T}^v \subseteq \mathcal{T}$	Arrival and required departure time of EV user v
$[e_{min}^v, e_{max}^v]$	Battery capacity constraints for EV v SoC
$e_t^v, e_{arr}^v, e_{req}^v$	State of charge (SoC) of user v 's EV at time step t , at arrival, and the required SoC at departure.
$\theta_v = \{t_{arr}^v, e_{arr}^v, \theta_{t_{dep}}^v, \theta_{e_{dep}}^v\}$	User parameters: arrival time, arrival SoC, departure time, departure SoC
$c \in C, c_k$	Set of charger types, and their counts
c_{min}, c_{max}, η	Minimum and maximum charging rates supported by the charger, and its efficiency (charging/discharging)
$\phi_t : \mathcal{V}_t \leftrightarrow C_t$	One-to-one mapping between user's EVs and chargers at time t
p_t^b	Building's power load (kW) (without chargers) at time step t
w_t	Energy price (\$/kWh) at time t
K_{DC}	Demand charge rate (\$/kW of peak demand)
K_{SOC}	Penalty coefficient for unfulfilled SoC in objective
K_{batt}	Penalty coefficient for battery degradation due to discharging

Table 5: Key Input Variables

Charging decision variables:	
Variable	Description
$a_t^{v,c} \in \{0, 1\}$	Binary variable; 1 if EV v is assigned to charger c at time t , otherwise 0
r_t^v	Charging/discharging rate (kW) for EV v at time step t (decision variable)
g_t^v, g_t	Energy draw by EV of user v and the total (building + EV charging) at time t
$y_t^{v,\delta}$	Indicates the current SoC segment in piece-wise linear SoC curve for EV v
$x_t^{v,c,\delta}$	Links maximum charging rate to SOC segment
p_t	Total power draw (building + EV charging) at time t
p^{max}	Peak power over peak periods (for demand charges)
q_t^v	Auxiliary variable tracking discharge to limit battery wear
z^v	SoC shortfall (absolute difference between required and achieved SoC at departure) for EV v
Additional Negotiation decision variables:	
$\lambda_t^v \in \{0, 1\}$	Indicates departure time step of EV v
$\sigma_{soc}^v, \sigma_{dep}^v$	Deviation from requested SoC and departure time for EV v
$U(\theta_v)$	Marginal utility of user v defined as cost reduction due to v 's participation
ζ	User charging cost computed as energy cost minus shared utility
L	Set of generated options for negotiations

Table 6: Key Decision Variables

the charging energy r_t^v in each segment of the SoC curve:

$$\forall v \in \mathcal{V}, \forall c \in C, \forall t \in \mathcal{T}, \forall \delta \in \Delta :$$

$$x_t^{v,c,\delta} \leq a_t^{v,c}; \quad x_t^{v,c,\delta} \leq y_t^{v,\delta}; \quad x_t^{v,c,\delta} \geq a_t^{v,c} + y_t^{v,\delta} - 1 \quad (9)$$

$$\forall v \in \mathcal{V}, \forall t \in \mathcal{T}^v : r_t^v \leq \eta \sum_{c \in C} \sum_{\delta \in \Delta} r_{\text{max}}^{v,\delta} \cdot x_t^{v,c,\delta} \quad (10)$$

where η denotes the round-trip efficiency of charging and discharging, and $r_{\text{max}}^{v,\delta}$ represents the maximum charging rate for EV v within SoC segment δ , expressed as a fraction of the charger's capacity c_{max} .

The minimum charging rate is set to be the minimum rate supported by the connected charger:

$$\forall v \in \mathcal{V}, \forall t \in \mathcal{T}^v : r_t^v \geq \eta \sum_{c \in C} c_{\text{min}} \cdot a_{v,c,t} \quad (11)$$

The SoC update over time follows:

$$\forall v \in \mathcal{V}, \forall t \in \mathcal{T}^v : e_{t_n}^v = \begin{cases} e_{t_{arr}}^v + r_{t_n}^v, & \text{if } n = 0 \\ e_{t_{n-1}}^v + r_{t_n}^v, & \text{otherwise} \end{cases} \quad (12)$$

where, e_0^v is the initial charge the EV arrives with.

To discourage excessive discharging, the total discharge for each EV is tracked using:

$$q_t^v \geq -r_t^v \quad (13)$$

Electricity Cost and SoC Completion. The energy cost to the building during each time step combines EV charging loads and the building's base load under the time-of-use price:

$$\forall t \in \mathcal{T} : g_t = \sum_{v \in \mathcal{V}} w_t (c_t^v + p_t^b / \tau) \quad (14)$$

To ensure user satisfaction, each EV must reach its desired SoC before departure, and deviations are penalized:

$$\begin{aligned} \forall v \in \mathcal{V} : z_v &= \left| e_{\text{req}}^v - e_{\text{dep}}^v \right| \\ \implies z_v \geq e_{\text{req}}^v - e_{\text{dep}}^v ; z_v &\geq e_{\text{dep}}^v - e_{\text{req}}^v \end{aligned} \quad (15)$$

Demand Cost Calculation. To compute the demand cost, we track the power use of the building and chargers combined. The demand cost is calculated on the maximum power.

$$\forall t \in \mathcal{T} : p_t = p_t^b + \frac{1}{\tau} \sum_{v \in \mathcal{V}} r_t^v \quad (16)$$

$$\forall t \in \mathcal{T} : p_t^{\max} \geq p_t \quad (17)$$

Objective Function. The overall objective is to minimize the total electricity cost (energy cost + demand cost) and penalties for missed SoC targets:

$$\begin{aligned} \min \sum_{t \in \mathcal{T}} g_t + K_{\text{DC}} \cdot p_t^{\max} + K_{\text{SOC}} \cdot \sum_{v \in \mathcal{V}} z_v \\ + K_{\text{batt}} \sum_{v \in \mathcal{V}} \sum_{t \in \mathcal{T}^v} q_t^v \end{aligned} \quad (18)$$

where K_{DC} is the demand charge rate, K_{SOC} is the compensation for unfulfilled SoC (per kWh) and controls the penalty for unfulfilled SoC targets. K_{batt} is the penalty rate per kWh discharged to limit excessive battery wear.

B.1 Online Control

In real-world deployment, we consider the problem of optimizing EV charging decisions in a dynamic environment, where the static optimization method described earlier becomes inapplicable due to the absence of oracle inputs. To handle this online decision-making process under uncertainty, we propose a Monte Carlo Model Predictive Control (MC-MPC) framework, which belongs to the family of stochastic Model Predictive Control (SMPC) methods [16]. This framework extends our static optimization approach to an online setting by leveraging predictive models to simulate future scenarios.

Online EV Charging Optimization. To support real-time operation, we reformulate the offline MILP problem as a sequential decision-making process, where the system evolves over discrete time steps. At each time step t , the system is characterized by the current state: $S_t = (\{e_t^v, e_{\text{req}}^v, e_{\text{min}}^v, e_{\text{max}}^v, t_{\text{dep}}^v, \phi_t(v)\}_{v \in \mathcal{V}_t}, b_t^p, p_{\text{past}}^{\max})$ where, $\mathcal{V}_t \subseteq \mathcal{V}$ denotes the set of EVs present by time t , and for each EV $v \in \mathcal{V}_t$: e_t^v is the current state-of-charge (SoC), e_{req}^v is the requested SoC at departure, e_{min}^v and e_{max}^v are the SoC bounds, t_{dep}^v is the scheduled departure time, $\phi_t(v)$ indicates the charger assigned to v , and The building state includes the current energy usage b_t^p and the past peak demand p_{past}^{\max} .

Based on the current state, the policy selects an action: $A_t = (r_t^v \mid v \in \mathcal{V}_t)$, where r_t^v is the assigned charging or discharging energy for each EV v present at time t . The system then transitions

to the next state s_{t+1} based on the applied charging actions, EV arrivals and departures, battery dynamics (as defined in Eqs. (6)-(17)) and the updated building load.

MC-MPC Controller. We define the MC-MPC controller as a policy function $\pi_{\text{ch}} : (S_t, \mathcal{F}) \rightarrow A_t$, where S_t is the current system state and \mathcal{F} is a set of Monte Carlo-sampled future trajectories. Given these inputs, the controller selects a feasible action A_t that minimizes the expected cumulative cost over the remainder of the billing period, based on outcomes simulated from \mathcal{F} . Each trajectory $f \in \mathcal{F}$ simulates the system's evolution over a prediction horizon $[t, H(t)]$, where $H(t)$ denotes the final time step of the current decision period (the end of the day for our problem). The simulation includes future EV arrivals, their departure times and SoC target, as well as building load profiles, using the generation models (described in Section 6). Let $\mathcal{V}^f = \mathcal{V}_t \cup \mathcal{V}_{t+1:H(t)}^f$ represent the set of all EVs (existing and future) in trajectory f .

At each time step, we aim to determine the current action A_t by solving a short-term static optimization problem, which, given the current EV status specified in S_t , minimizes the expected cost across all sampled future trajectories \mathcal{F} , coupled at the current state. While each trajectory f may follow its own optimal actions from $t+1$ to $H(t)$, the action at the current time step A_t must remain consistent across all trajectories. This leads to the following optimization formulation:

$$\begin{aligned} \min_{A_t, \{A_{t+1}^f, \dots, A_{H(t)}^f\}_{f \in \mathcal{F}}} \mathbb{E}_{f \in \mathcal{F}} \left[\sum_{k=t}^{H(t)} g_{k,f} + w^d p_f^{\max} \right. \\ \left. + K_{\text{SOC}} \sum_{v \in \mathcal{V}^f} z_v^f + K_{\text{batt}} \sum_{v \in \mathcal{V}} \sum_{t \in \mathcal{T}^v} q_{t,f}^v \right] \end{aligned} \quad (19)$$

subject to the EV-charger assignment feasibility constraints. Here, A_t is the common action applied at time t , while $\{A_{t+1}^f, \dots, A_{H(t)}^f\}$ are trajectory-specific future actions. The term g_k^f denotes the energy cost at time step k in trajectory f , p_f^{\max} is the peak demand over the trajectory (including the historical peak p_{past}^{\max}), and z_v^f represents the SoC shortfall for EV v upon its departure in trajectory f .

THEOREM 4. Temporal decomposition. *Minimizing daily peak demand independently each day leads to a minimized monthly demand cost, without requiring multi-day coupling.*

PROOF. Our MC-MPC approach minimizes the monthly total cost by decomposing the problem into daily control tasks and optimizing each day independently. This is supported by (i) the system assumption that EVs do not persist across days in the building, and (ii) the property of the demand charge structure that the monthly peak can occur on any day, allowing each day's peak to contribute independently to the monthly maximum.

Let \mathcal{D} denote the set of days in the monthly billing cycle, and for each $d \in \mathcal{D}$, let \mathcal{T}_d denote the time steps belonging to day d . Define the full monthly time horizon as $\mathcal{T}^{\text{month}} = \bigcup_{d \in \mathcal{D}} \mathcal{T}^d$. The monthly peak demand is given by: $P_{\text{max}}^{\text{month}} = \max_{t \in \mathcal{T}^{\text{month}}} (\pi_t)$, where π_t stores the instantaneous total power use at time t . By

properties of the maximum operator over disjoint sets, we have:

$$\max_{t \in \mathcal{T}^{\text{month}}} \pi_t = \max_{t \in \bigcup_{d \in \mathcal{D}} \mathcal{T}^{(d)}} \pi_t = \max_{d \in \mathcal{D}} (\max_{t \in \mathcal{T}^{(d)}} \pi_t)$$

We can define the daily peak load as $P_{\max}^d = \max_{t \in \mathcal{T}^{(d)}} (\pi_t)$ and it follows:

$$P_{\max}^{\text{month}} = \max_{d \in \mathcal{D}} P_{\max}^d.$$

We formally show that minimizing daily peak loads independently P_{\max}^d minimizes the overall monthly peak P_{\max}^{month} . \square

Action Refinement. To further reduce the monthly demand cost while meeting EV SoC requirements, we introduce an action refinement step defined as a function $f_{\text{refine}} : (A_t, \hat{P}_{\max}) \rightarrow \tilde{A}_t$, which adjusts the action A_t based on the estimated peak power \hat{P}_{\max} . If the current aggregate power (building load plus charging) is below \hat{P}_{\max} , charging rates can be safely increased without incurring additional demand costs. This proactive adjustment helps fulfill SoC targets, discharge at opportune times and reduce the likelihood of higher peaks later in the billing period as EVs approach departure. Specifically, we estimate the gap between the predicted peak power and the current total charging rate: $P_{\text{diff}} = \hat{P}_{\max} - (\sum_{v \in \mathcal{V}_t} r_t^v / \tau + b_t^p)$. If $P_{\text{diff}} > 0$, we distribute this remaining capacity equally across all connected EVs while respecting their charging constraints and SoC limits. The final charging rate \tilde{A}_t for each EV $v \in \mathcal{V}_t$ is adjusted by:

$$\tilde{r}_t^v = \begin{cases} r_t^v + \min\left(\frac{P_{\text{diff}}}{|\mathcal{V}_t|} \tau, e_{\text{req}}^v, \phi(v)_{\text{max}}\right), & \text{if } \phi(v)_{\text{min}} = 0 \\ r_t^v + \min\left(\frac{P_{\text{diff}}}{|\mathcal{V}_t|} \tau, e_{\text{max}}^v, \phi(v)_{\text{max}}\right), & \text{if } \phi(v)_{\text{min}} < 0 \end{cases} \quad (20)$$

for $v \in \mathcal{V}_t$ where $\phi(v) = 0$ denotes a uni-directional charger, and $\phi(v) < 0$ denotes a bi-directional charger. Algorithm 2 provides an overview of our V2B charging optimization framework.

Algorithm 2: Online V2B Charging Optimization

Input: Initial state S_0 ; number of samples N ; billing period

T ; estimated monthly peak power \hat{P}_{\max}

Output: Charging decisions $\{A_0, A_1, \dots, A_T\}$

```

1 for  $t = 0$  to  $T$  do
2   Check EV departures and free assigned chargers;
3   Assign chargers to new arrivals using FIFO; prioritize
   bidirectional chargers;
4   Observe current state  $S_t$ ;
5   Future trajectory set  $\mathcal{F} \leftarrow \emptyset$ ;
6   for  $i = 1$  to  $N$  do
7     Sample trajectory  $f$  from the generative model
   given  $S_t$ ;
8     Simulate from  $t+1$  to  $H(t)$  and add  $f$  to  $\mathcal{F}$ ;
9   Get  $A_t$  by solving the optimization in Eq. (19);
10  Refine action to  $\tilde{A}_t = f_{\text{refine}}(A_t, \hat{P}_{\max})$  by Eq. (20);
11  Apply  $\tilde{A}_t$  and update system state to  $S_{t+1}$ ;

```

C OPTIMIZATION FRAMEWORK FOR NEGOTIATION

In this section, we present the optimization framework that underpins our negotiation-based Model Predictive Control (MPC)

approach for managing and providing options in the V2B electric vehicle charging environment. The framework formulates the optimization problem as a Mixed-Integer Linear Program (MILP) that jointly considers user charging preferences (deviation limits), building operational constraints, and dynamic electricity pricing and demand charges.

The following subsections detail the MILP formulation, scenario sampling procedures, and the solution methodology that iteratively adapts charging and incentive offers at each decision epoch.

Decision Variables

We define the following decision variables, distinguishing between those reused from the charging policy and those introduced for the negotiation policy. We reuse some variables from the charging policy, and define more new variable specific to the negotiations case. Variables inherited from the charging policy are:

- $a_t^{v,c} \in \{0, 1\}$: Indicates whether electric vehicle (EV) v is assigned to charger c at time step t .
- $r_t^v \geq 0$: Charging power assigned to EV v at time step t .
- $e_t^v \geq 0$: State of charge (SoC) of EV v at time step t .
- t_{dep}^v : Optimized departure time for EV v .
- $e_{\text{dep}}^v \geq 0$: SoC of EV v at its (possibly negotiated) departure time.

In the negotiation problem, we require additional variables since the departure time and requested state of charge (SoC) are not fixed. Instead, for each option $l \in L$, we know only the allowable deviation limits. These limits are used to identify the optimal combination of user utility, departure time, and SoC.

- $\lambda_t^v \in \{0, 1\}$: Indicates whether EV v departs at time step t .
- $\sigma_{\text{soc}}^v \in [0, e_l^{\text{max}}]$: Deviation of SoC at departure from the required SoC for EV v , bounded by limit l .
- $\sigma_{\text{dep}}^v \in [0, t_{\text{dep},l}^{\text{max}}]$: Deviation of actual departure time from the required departure time for EV v , bounded by limit l .

Constraints

- (1) As discussed earlier, at each time step t , chargers and EVs are matched one-to-one: each charger serves at most one EV, and each EV connects to at most one charger based on the first-come-first-served (FCFS) assignment policy. Formally, this defines the mapping $\phi_t : \mathcal{V}_t \leftrightarrow \mathcal{C}_t$,
- (2) Similar to equation (6), the charger assigned to EV v must remain the same throughout its duration of stay. However, since the exact departure time is unknown, we introduce λ to determine the optimal departure time:

$$\forall v \in \mathcal{V}, \forall c \in \mathcal{C}, \forall t \in \mathcal{T}^v, t > t_{\text{arr}}^v : \quad (21)$$

$$a_{t-1}^{v,cp} \leq a_t^{v,cp} + \lambda_t^v \quad (22)$$

- (3) To prevent assignments after departure:

$$\forall v \in \mathcal{V}, \forall c \in \mathcal{C}, \forall t \in \mathcal{T}^v : a_{v,c,t} \leq 1 - \sum_{t=t_{\text{arr}}^v}^t \lambda_t^v \quad (23)$$

- (4) Charging is done only when assigned, and follows the exact constraints from the charging policy. Thus, we apply equations (7) - (12) to maintain the charging and discharging rates in each SoC segment of the piece-wise linear SoC curve.

- (5) Any discharge accounts towards battery wear and is controlled similar to the charging policy, using equation (13).
(6) Since the problem decides the optimized departure time step within the duration of stay, it is tracked using:

$$\forall v \in \mathcal{V} : t_{\text{dep}}^v = \sum_{t \in \mathcal{T}^v} t \cdot \lambda_{v,t} \quad (24)$$

- (7) We also need to track and make sure that each vehicle departs exactly once, within the maximum limit of its duration \mathcal{T}^v :

$$\forall v \in \mathcal{V} : \sum_{s \in \mathcal{T}^v} \lambda_t^v = 1 \quad (25)$$

- (8) Within the allowed deviation limits, we need to find the optimal deviations in SoC and departure time. The following equations keep track of these deviations:

$$\forall v \in \mathcal{V} : \sigma_{\text{soc}}^v = e_{\text{req}}^v - e_{\text{dep}}^v \quad (26)$$

$$\forall v \in \mathcal{V} : \sigma_{\text{dep}}^v = t_{\text{dep}}^v - \bar{t}_{\text{req}}^v \quad (27)$$

- (9) Energy cost calculation for the building:
We follow the same constraints in equations (14) - (15) to calculate the energy cost.
(10) Demand cost calculation for the building:
Equations (16) - (17) are used to calculate the demand cost.

Objective Function. The overall objective is again the same as in the charging optimization, i.e., to minimize the total cost cost in eq. (1), to minimize energy, demand cost, penalties for missed SoC targets, and excessive discharge:

$$\min \sum_{t \in \mathcal{T}} g_t + K_{\text{DC}} \cdot p_{\text{max}} + K_{\text{SOC}} \cdot \sum_{v \in \mathcal{V}} z_v + K_{\text{batt}} \sum_{v \in \mathcal{V}} \sum_{t \in \mathcal{T}^v} q_t^v \quad (28)$$

where K_{DC} is the demand charge rate, K_{SOC} is the compensation for unfulfilled SoC (per kWh) and controls the penalty for unfulfilled SoC requirements. K_{batt} provides the battery degradation penalty for discharging.

Practical Considerations: This approach does not require perfect foresight; it only relies on the latest forecasts, updating its policy as new data arrive.

At every step, MPC with Monte Carlo sampling ensures that the negotiation process is both flexible and tractable. The system computes user-specific incentives and charging schedules that align with both real-time operational requirements and user behavior, all while maintaining scalability to large fleets and adaptability to new data or planning advancements.

D PROOF OF THEOREMS

STRATEGY-PROOFNESS FOR DEPARTURE TIME AND SoC

The mechanism is strategy-proof with respect to departure time and requested SoC: a user cannot increase their utility by misreporting an earlier departure time or a higher required SoC $\bar{\theta}_v$ than their true preferences θ_v . Formally, truthful reporting always maximizes user utility, i.e., $U(\theta_v) \geq U(\bar{\theta}_v)$.

PROOF. Consider a user v whose utility is defined as the marginal reduction in system cost due to their participation:

$$U(\theta_v) = J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v}) - J^{\pi^*}(\theta_v, \theta_v, \hat{\theta}_{\setminus v}).$$

Here $J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v})$ denotes the expected cost of the system without user v (but including the other users), and $J^{\pi^*}(\theta_v, \theta_v, \hat{\theta}_{\setminus v})$ the cost with user v participating truthfully.

Departure Time. Suppose user v reports a departure time $\bar{t}_{\text{dep}} \in \bar{\theta}_v$ that is earlier than their true departure time $t_{\text{dep}} \in \theta_v$ ($\bar{t}_{\text{dep}} < t_{\text{dep}}$). This act constrains the set of possible charging actions for the system because the charging schedule for v must now conclude by \bar{t}_{dep} rather than t_{dep} . Formally, the set of time steps available for charging is reduced from the interval $[t_{\text{arr}}, t_{\text{dep}}]$ to $[t_{\text{arr}}, \bar{t}_{\text{dep}}]$.

Accordingly, the optimization is conducted over a reduced set of time steps for user v , thereby limiting the system's temporal flexibility to leverage lower-priced periods or to achieve better peak load flattening. Any load shifting or peak smoothing achievable within the shortened horizon $[t_{\text{arr}}, \bar{t}_{\text{dep}}]$ is also feasible over the extended horizon $[t_{\text{arr}}, t_{\text{dep}}]$. Thus, allowing the user to remain until t_{dep} (with $t_{\text{dep}} > \bar{t}_{\text{dep}}$) provides additional opportunities for load shifting and peak flattening, potentially reducing overall costs. Since the system always minimizes the total cost within the available set of charging actions, restricting the window cannot lower this minimum cost. Thus,

$$J^{\pi}(\theta_{\setminus v}, \bar{\theta}_{\setminus v}, \hat{\theta}_{\setminus v}) \geq J^{\pi}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v}) \text{ and,}$$

$$U(\bar{\theta}_v) = J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v}) - J^{\pi^*}(\theta_v, \bar{\theta}_v, \hat{\theta}_{\setminus v}). \quad (29)$$

Subtracting the utilities,

$$U(\theta_v) - U(\bar{\theta}_v) = J^{\pi}(\theta_{\setminus v}, \bar{\theta}_v, \hat{\theta}_{\setminus v}) - J^{\pi}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v}). \quad (30)$$

Therefore,

$$U(\theta_v) - U(\bar{\theta}_v) \geq 0, \quad (31)$$

which implies

$$U(\theta_v) \geq U(\bar{\theta}_v). \quad (32)$$

Thus, the utility for misreporting is less than or equal to that for reporting truthfully. We prove that reporting the true departure time maximizes user utility.

State of Charge (SoC). Similarly, if user v reports a higher required SoC $\bar{e}_{\text{req}}^v \in \bar{\theta}_v$ at departure than their true requirement $e_{\text{req}}^v \in \theta_v$, i.e., ($\bar{e}_{\text{req}}^v > e_{\text{req}}^v$), the system must deliver a greater amount of energy within the same (or possibly even a reduced) time window. This imposes a stricter set of feasible charging schedules, as the minimum total energy to be supplied increases. The increased constraint limits the system's ability to optimize cost, and so

$$J^{\pi}(\theta_{\setminus v}, \bar{\theta}_{\setminus v}, \hat{\theta}_{\setminus v}) \geq J^{\pi}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v}) \text{ and,}$$

Hence, similar to the previous proof, by following steps (29)-(31), we can prove that the user cannot achieve greater utility by exaggerating their SoC requirement, and,

$$U(\theta_v) \geq U(\bar{\theta}_v).$$

Calculation of Utility Difference. Combining both cases, for any misreported parameter $\bar{\theta}_v$, the utility difference obeys

$$U(\theta_v) - U(\bar{\theta}_v) \geq 0.$$

Thus, no user benefits by tightening their reported constraints; truthful reporting of departure time and SoC is optimal. \square

BOUNDING UTILITY BY TOTAL SAVINGS

Any payment ζ^v to user $v \in \mathcal{V}$ must satisfy $\zeta^v \geq -U(\theta_v)$ to ensure that the building's total cost after incentives does not exceed the baseline.

PROOF. We begin from the budget feasibility constraint, which requires that the building's cost with user v and payment ζ^v does not exceed the cost without v :

$$J^{\pi^*}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v}) - \zeta^v \leq J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v}). \quad (33)$$

Rearranging terms:

$$-\zeta^v \leq J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v}) - J^{\pi^*}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v}) = U(\theta_v). \quad (34)$$

Thus, we obtain the bound:

$$\zeta^v \geq -U(\theta_v). \quad (35)$$

which is valid through Eq. (3).

This ensures that the payment to user v does not exceed the cost savings they bring. In the limiting case (equality), we recover the identity:

$$J^{\pi^*}(\theta_{\setminus v}, \theta_v, \hat{\theta}_{\setminus v}) + U(\theta_v) = J^{\pi^*}(\theta_{\setminus v}, \hat{\theta}_{\setminus v}). \quad (36)$$

\square

VOLUNTARY PARTICIPATION

For all rational users v , the user choice mechanism guarantees $Y_l^v \geq 0$. Thus, no user is worse off by participating, and some benefit strictly if $U(\theta_v) > 0$.

PROOF. If $U(\theta_v) \geq 0$, then from Eq. 3, the user cost for option l is less than the charging cost, $\zeta_l^v \leq \sum_{t \in \mathcal{T}^v} g_t^n w_t$. This ensures the user has a charging price lower than the time-of-use energy cost.

Also, recall the user satisfaction metric $Y_l^n = \bar{E} - (\zeta_l + I_l^n)$ from eq. 5. Whenever

$$\zeta_l + I_l^n > \bar{E} \iff Y_l^n < 0,$$

the charging options at the building are strictly more expensive than the external alternative. A rational user therefore rejects the contract and charges externally at cost \bar{E} . Consequently, the maximum amount a user can ever pay at the building is \bar{E} , and participation in the building's negotiation mechanism can never leave them worse off. \square

E USER SEGMENTATION FROM SURVEY

E.1 Flexibility limits for L Negotiation Options

To identify distinct user groups based on their flexibility preferences, we conducted a survey asking respondents about two key behavioral dimensions: their maximum acceptable departure delay and their maximum acceptable reduction in state-of-charge (SoC), or range flexibility, when offered incentives. These questions were designed to capture users' willingness to adjust their charging behavior under different conditions. All the survey questions are shared below.

Clustering Methodology

After collecting the survey responses, we preprocessed the data for clustering analysis. Responses for maximum departure delay, initially provided as ordinal categories from 1 to 9, were converted to corresponding time values in minutes (0, 15, 30, 45, 60, 75, 90, 105, 120 minutes). Similarly, range flexibility responses, originally one-hot encoded, were transformed into percentage values (0%, 5%, 10%, 20%) representing the maximum acceptable SoC reduction.

Data preprocessing also involved removing incomplete responses, such as those where users selected "I need more info" for range flexibility or provided no response. These could not be meaningfully incorporated into the distance-based clustering algorithm. This filtering resulted in a final dataset of 27 valid responses across both dimensions.

Given the different scales of the two variables (minutes versus percentages), we applied standardization using z-score normalization to ensure equal weighting in the clustering algorithm. This step is critical for K-means clustering, as the algorithm is sensitive to variable scales and may otherwise be dominated by the variable with larger numerical range.

Finally, we applied K-means clustering to the standardized data to segment users into homogeneous groups. These clusters provide valuable insights for developing targeted charging strategies tailored to the flexibility profiles of different user segments.

Optimal Cluster Determination

To identify the optimal number of clusters, we applied two widely-used validation techniques: the *silhouette method* and the *elbow method*. Both approaches help assess cluster quality by analyzing solutions over a range of cluster counts ($k = 2$ to $k = 7$).

The *silhouette score* quantitatively measures how well each data point fits within its assigned cluster. For a data point i , its silhouette value $s(i)$ is given by:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

where $a(i)$ is the average distance from i to all other points in its own cluster, and $b(i)$ is the lowest average distance from i to points in any other cluster (i.e., the nearest neighboring cluster). The silhouette score ranges from -1 to 1 , where a higher value indicates better-defined and more cohesive clusters.

The *elbow method* involves plotting the total within-cluster sum of squared distances (WCSS) for different values of k . As k increases, WCSS decreases because clusters contain fewer points and become more compact. However, after a certain point, the rate of decrease sharply slows, forming an "elbow" which suggests that additional clusters beyond this point yield limited improvement and may not be meaningful.

Number of Clusters (k)	Silhouette Score
2	0.498
3	0.588
4	0.547
5	0.563
6	0.589
7	0.596

In our analysis (Figure 3), the highest silhouette score was observed at $k = 7$, suggesting that this value offers the most well-defined partitioning. However, $k = 3$ yields a comparable silhouette score and provides greater applicability to our problem. The elbow

method also supports $k = 3$ as a suitable choice (see Figure 4). Therefore, we considered both mathematical criteria and domain interpretability when selecting the optimal number of clusters.

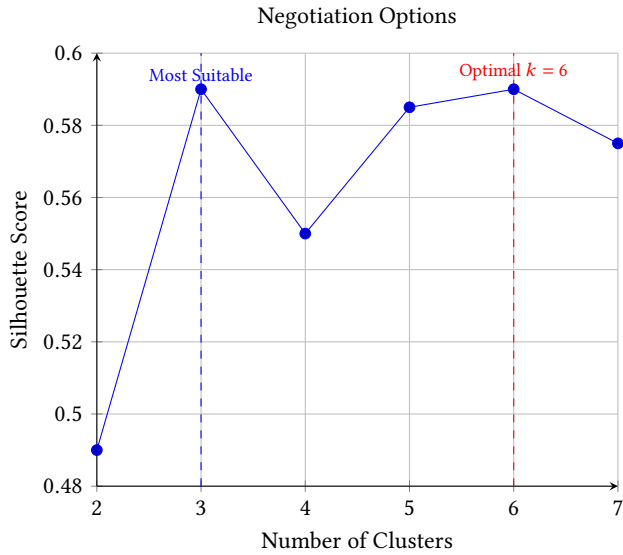


Figure 3: Silhouette method shows $k = 6$ as the optimal choice (highest score). However, $k = 3$ is numerically very close, and is most suitable as users might get overwhelmed with more options.

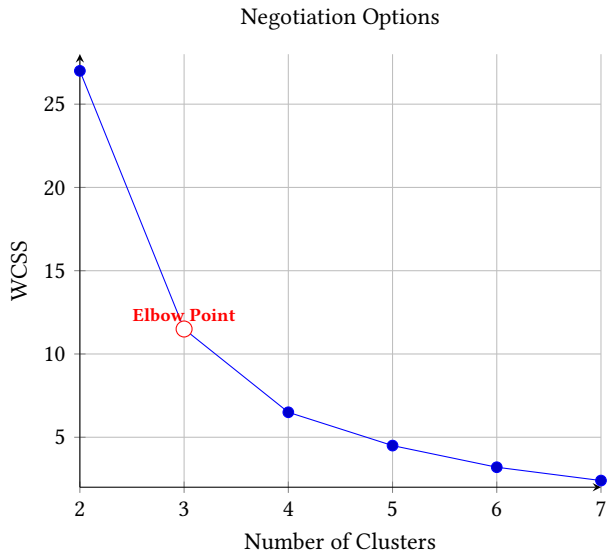


Figure 4: Elbow method shows a sharp decrease until $k = 3$, indicating the elbow point.

From a practical perspective, three clusters align well with common behavior groups in transportation research, representing low,




Figure 5: Clustered user deviation limits data provides Negotiation Options and flexibility limits

medium, and high flexibility users. This approach offers useful insights for managing charging infrastructure without adding the complexity that more clusters would bring.

E.2 Cluster Characterization and Representative Points

The final three-cluster solution effectively segmented users into distinct behavioral profiles, as summarized in Table 7, and is visualized in Figure 6. Each cluster was characterized by its centroid, the representative point that minimizes the within-cluster variance and serves as the archetypal user for that segment.

Table 7: Negotiation Options (Flexibility Analysis) Summary

Cluster	Count	Dep. Flexibility (min)		SoC Flexibility (%)	
		Mean	Std	Mean	Std
1	7	27.86	18.22	20.00	0.00
2	16	25.31	20.29	6.25	3.87
3	4	112.50	15.00	20.00	0.00

The cluster analysis revealed three distinct user segments:

- Cluster 1 - Low Flexibility Profile:** This segment ($n=16$, 59.3% of sample) constitutes the largest user group, characterized by similar moderate departure delay tolerance (mean = 25.31 minutes) but substantially lower range flexibility (mean = 6.25%). The higher standard deviation in departure delay (20.29 minutes) indicates greater heterogeneity within this cluster regarding time preferences, while the low standard deviation in range flexibility (3.87%) suggests consistent preference for minimal SoC reduction.
- Cluster 2 - Moderate Delay, High Range Flexibility:** This segment ($n=7$, 25.9% of sample) represents users with moderate departure delay tolerance (mean = 27.86 minutes) but uniformly high range flexibility (20%). These users demonstrate willingness to accept significant SoC reduction while maintaining relatively conservative time constraints.
- Cluster 3 - High Delay, High Range Flexibility:** This smallest segment ($n=4$, 14.8% of sample) exhibits the highest




Figure 6: Cluster selection and segmentation of user types. Left: The Elbow method plot shows inertia (WCSS) against the number of clusters, with an identified elbow at $k = 4$. Center: Silhouette scores for different cluster counts, indicating the best silhouette value at $k = 4$. Right: Final clustering of survey respondents into four user types based on SoC sensitivity (w_1) and time sensitivity (w_2), with cluster centroids marked by black crosses. Each color represents a distinct user type.

Scenario Type	Deviation Values	Example Base Cost (\$)
Departure delay (Δt_{dep}) (minutes)	0, 15, 30, 60	10
SoC reduction (Δe) (kWh)	0, 5, 10, 20	6 – 10

Table 8: User Type Characteristics: Survey input points

departure delay tolerance (mean = 112.50 minutes) combined with maximum range flexibility (20%). The low standard deviation in departure delay (15.00 minutes) indicates homogeneous preferences for extended charging times within this group.

The cluster centroids capture typical user behaviors in each group. These centroids represent the “typical” user within each segment and can help guide charging infrastructure design and operation by enabling tailored services that match the needs of different user types.

E.3 User Choice Modeling

Building the Inconvenience Cost Model

We use an anonymized stated-preference survey with 28 workplace EV users to quantify flexibility preferences and the monetary incentives required to accept various charging deviations. Each respondent indicated both how much they would need to save to tolerate additional departure delay and to accept reduced state-of-charge (SoC) at departure. The scenarios presented included:

- **Departure delay** (Δt_{dep}): Options for 15, 30, and 60 minutes, each with a base charging cost $\zeta = \$10$.
- **SoC reduction** (Δe): Reductions from an initial target of 80% SoC to 75% (5% reduction), 70% (10%), and 60% (20%), with scenario-specific base costs: $\zeta_{70\%} = \$8$, $\zeta_{75\%} = \$9$, $\zeta_{60\%} = \$6$.

Respondents mapped their categorical willingness to accept each deviation into corresponding discount percentages (ranging from 0% for “no incentive needed” to 40% for “extremely unlikely”), then to

dollar amounts atop the relevant base cost. This processing yielded 168 valid scenario-user-incentive tuples ($\Delta e^{v,j}$, $\Delta t_{\text{dep}}^{v,j}$, $I^{v,j}$). All the survey questions are shared below.

Step 1: Individual Inconvenience Cost Modeling

For each participant v , we modeled their required incentive as a linear function of SoC reduction and time delay:

$$I^{v,j} = w_1^v \Delta e^{v,j} + w_2^v \Delta t_{\text{dep}}^{v,j} + \epsilon^{v,j}$$

where w_1^v (\$/% SoC) and w_2^v (\$/min) represent sensitivities to SoC and time deviations, and $\epsilon^{v,j}$ is a residual. We fit (w_1^v, w_2^v) for all 28 respondents with at least two usable scenarios using ordinary least squares, constraining the intercept to zero. This linear model explained a moderate share of variation in user responses (mean $R^2 = 0.57$; median 0.81; quartiles: 0.36/0.81/0.90), but also revealed that user heterogeneity is substantial.

Step 2: User Type Clustering and Summary

To capture the diversity of user preferences at the population level, we standardized the fitted weights and applied K-means clustering to the survey respondents. After evaluating cluster solutions ranging from $N = 2$ to $N = 8$, we found that $N = 4$ clusters provided the best representation of distinct user types. The optimal number of clusters was selected based on both the Elbow and Silhouette methods, with a WCSS of 15.25 and a silhouette score of 0.424. Both metrics indicate a fair level of cluster separation, suggesting reasonably distinct but not perfectly separated clusters. Figure 6 displays the WCSS and silhouette scores, and the resulting clusters.

Type	SoC Sensitivity (w_1)	Time Sensitivity (w_2)	Share
i	0.0489 ± 0.02	0.1250 ± 0.01	10.7%
ii	0.0133 ± 0.01	0.0346 ± 0.02	53.6%
iii	0.0362 ± 0.01	0.0673 ± 0.01	32.1%
iv	0.0000 ± 0.00	0.1083 ± 0.00	3.6%

These clusters reveal distinct user flexibility profiles:

- Type i users are strongly averse to deviations in both departure time and SoC, representing highly inflexible behavior.
- Type ii users are the most flexible, showing the least sensitivity to changes in either parameter.
- Type iii users are moderately sensitive, willing to tolerate some deviation in both time and SoC, but not as much as Type II.
- Type iv users are a small minority who uniquely accept SoC deviations but remain inflexible regarding departure times.

This segmentation highlights clear differences in how users value their charging preferences and supports the development of more tailored management strategies.

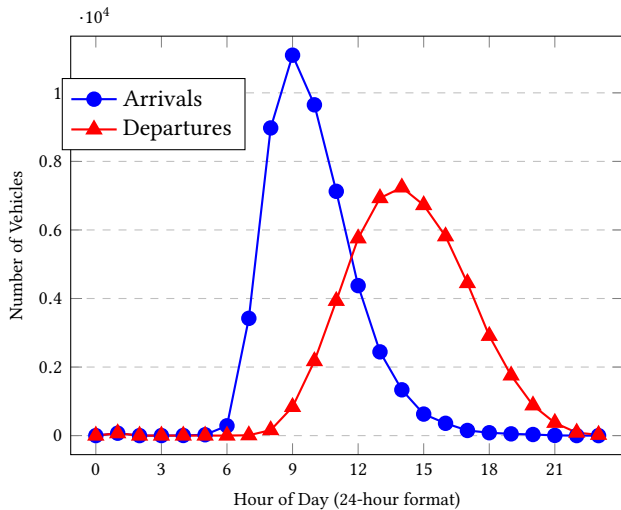


Figure 7: Hourly distribution of car arrivals and departures over all samples.

F EXPERIMENTS: MORE DETAILS

To model the piece-wise linear SoC charging curves, we break the charging into 3 segments. The charging power is kept constant at maximum (20 kW) for $\text{SoC} \leq 83\%$, then linearly decreasing as $-\frac{4}{3} \cdot \text{SoC} + 130$ for $83-90\%$, and $-\text{SoC} + 100$ for $90-100\%$. Discharging assumes constant power with linear SoC decline. Electricity prices follow Silicon Valley Power rates: \$0.178/kWh peak (6 AM–10 PM), \$0.137/kWh off-peak, and an \$11.67/kWh demand charge [22].

F.1 Forecast Generation.

For each control step, the MC-MPC framework requires forecasts of building load and EV behavior (arrivals, departures, SoC requirements). Building load forecasts are generated using standard autoregressive models fitted to the past month of building-meter

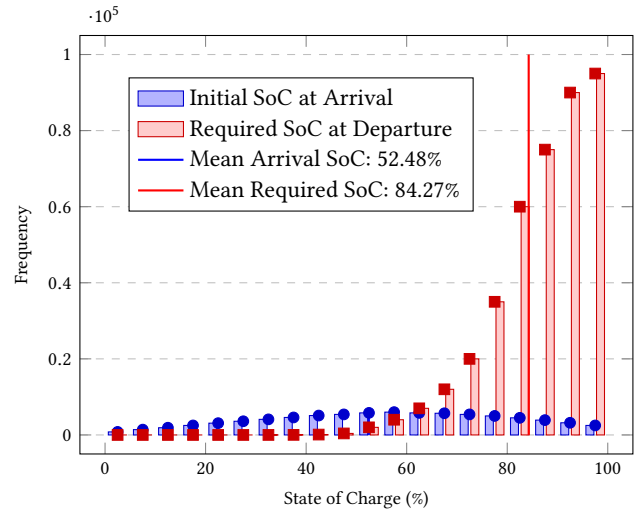


Figure 8: Overlay of Initial SoC at Arrival and Required SoC at Departure (bins shifted for clarity).

data; electricity prices follow pre-published power utility rates and are assumed to be known a day ahead.

EV Behavior Forecast. Figure 7 presents the ground truth data for electric vehicle (EV) arrivals and departures, while Figure 8 illustrates the distributions of State of Charge (SoC) at both arrival and departure. EV behavior is modeled using generative probabilistic models trained on twenty-two months of real-world telemetry data. EV behavior is modeled with generative probabilistic models trained on twenty-two months of telemetry data. Arrival counts per 15-minute interval are predicted using a zero-inflated negative binomial (ZINB) model, with input features including date, day of week, hour, and cyclical encodings of time-of-day (sine and cosine of hour). Arrival SoC is modeled via a beta regression using day of week and cyclical time features. Departure SoC is predicted using a second beta regression conditioned on the arrival SoC, day of week and time features. Stay durations are drawn from a Weibull survival model with event rates dependent on day of week and arrival time. We confirmed with the building operator that EV arrival and departure times depend solely on user schedules, not on building load or prior charging, so we presample arrivals, initial SoC, and dwell durations. This decoupling means charging decisions will not affect future availability, letting the controller focus exclusively on meeting SoC targets at minimum cost.

Model diagnostics. Table 9 lists concise goodness-of-fit statistics for each generative component that feeds the Monte-Carlo SMPC.

- **Arrival counts.** The zero-inflated negative-binomial model explains roughly one-third of the deviance (pseudo- $R^2 = 0.34$), confirming that calendar and diurnal features capture meaningful variation in arrivals.
- **Arrival & departure SoC.** Both beta regressions achieve strongly negative AIC values, indicating parsimonious yet well-fitting density estimates for battery state distributions.

Table 9: Diagnostic metrics for generative EV behavior models

Model	LL	AIC	Note
Arrival count (ZINB)	-2526	5057	Pseudo- $R^2 = 0.34$
Arrival SoC (Beta)	+128	-241	Matches histogram well
Departure SoC (Beta)	+909	-1801	Cond. on arrival improves fit
Stay duration (Weibull)	-104	-201	C-index = 0.71

- **Stay duration.** The Weibull survival model attains an AIC near -200 and a concordance index around 0.71 , reflecting useful predictive skill for dwell times.

These diagnostics support the credibility of the sampled future scenarios used by the MC-MPC controller.

Building Load Forecast. Building demand can usually be predicted to within $\approx 2-10\%$ mean absolute percentage error (MAPE) over day-ahead horizons [6, 30]. We therefore create a “forecast” by adding zero-mean Gaussian noise, tuned to that error band, to the realized load. This keeps EV arrivals and departures as the dominant source of uncertainty and avoids entangling them with exogenous load-forecast errors. A fully coupled forecasting-control stack is left for future work.

G RESULTS: MORE DISCUSSION

We begin the discussion by presenting a comprehensive evaluation of our proposed MC-MPC charging framework. This includes detailed performance metrics across key dimensions such as building cost, missing SoC, and peak shaving. By highlighting both aggregate trends and scenario-specific behaviors, we aim to demonstrate the robustness and adaptability of our approach under varying operational conditions.

G.1 MC-MPC as Charging Policy

Baseline Approaches. We evaluate the performance of our online approach by comparing it against various methods including real-world charging procedures, several smart heuristic approaches, and a reinforcement learning-based policy. We provide a brief description of the baselines here.

- **MaxCharge:** This approach simulates current real-world charging, where all connected EVs are charged at the fastest possible rate.
- **ReqCharge:** Similar to *MaxCharge*, however, this policy only charges all connected EVs as quickly as possible to their requested SoC, e_{req}^v .
- **Least Laxity First (LLF):** A heuristic that prioritizes charging EVs with the least amount of idle time before their departure [17].
- **Early Deadline First (EDF):** A heuristic that prioritizes charging EVs with the earliest departure time [23].
- **Reinforcement Learning (RL):** We use a Deep Deterministic Policy Gradient (DDPG) algorithm to manage charging actions, as detailed in prior work [13]. A separate model is trained for each month.

We evaluate all approaches on 50 test episodes per month on four key metrics: *Total Cost*, *Missing SoC*, the count of *Cars Under Required SoC*, and *Peak Shaving* relative to the MaxCharge policy.

Comparative Performance Analysis. Table 10 reports the scaled monthly electricity costs for each policy. The offline Oracle (MILP), while infeasible in practice, provides a theoretical lower bound on performance. Our MC-MPC policy performs exceptionally well, with an 8-month average cost of \$8785, which is within 1.6% of the oracle’s benchmark. This demonstrates a near-optimal balance of foresight and real-time adaptability.

The RL policy achieves a competitive average cost of \$8808. While it secures lower costs in some summer months by implicitly learning temporal demand patterns, our MC-MPC proves more stable and cost-effective on average. Both intelligent policies significantly outperform the heuristics; MC-MPC is approximately 1.3% more efficient than LLF [17] and EDF [23]. The gap is even wider when compared to rigid, real-world policies, with our approach outperforming ReqCharge and MaxCharge by 2.2% and 5.4%, respectively.

Beyond monthly costs, reliability and grid impact are critical. Through Table 11, it is observed that MC-MPC and RL consistently meet all EV energy demands, resulting in zero missing SoC across all test episodes. In contrast, heuristic approaches like LLF and EDF fail to fully charge several vehicles each month. Furthermore, MC-MPC achieves the highest peak shaving among all practical methods, outperforming RL slightly and surpassing the heuristics by a wide margin. These metrics confirm that MC-MPC offers the most robust and reliable solution, delivering near-optimal cost savings without compromising user satisfaction or grid stability.

Experiments Under Uncertainty. Uncertainty is a fundamental aspect of real-world EV charging operations. Vehicle arrivals and departures are often unpredictable, user preferences can vary dynamically, and building loads fluctuate throughout the day. To evaluate the robustness of different control strategies under these practical conditions, we assess performance across five representative uncertainty scenarios. These scenarios encompass both aleatoric uncertainty, arising from inherent randomness, and epistemic uncertainty, stemming from model misalignment or deployment shifts. Tables 12 and 13 summarize each strategy’s ability to handle uncertainty, measured in terms of total building cost and the incidence of missing state-of-charge (SoC) values.

Aleatoric uncertainty:

- **Departure-time uncertainty:** EV departure deviation modeled as a Gaussian perturbation with mean as the initial departure time, and 2-hour standard deviation applied to departure times, given they have enough time to charge to meet the required SoC.
- **Out-of-distribution (OOD) SoC requests:** Departure SoC targets are perturbed by adding a random offset, uniformly sampled within $\pm 15\%$ of the originally generated SoC value, simulating deviations from the assumed generative model.
- **Building load noise:** The building’s baseline demand is perturbed by sampling zero-mean noise from a uniform distribution with support $\pm 15\%$ of the average daily load. This models day-level fluctuations in building load without adding a systematic bias.

Table 10: Comparative analysis of different real-time charging policies, with costs scaled for consistency with Table 1. The results highlight the superiority of the MC-MPC strategy, which achieves the lowest average electricity cost among all practical methods, performing within a small margin of the theoretical Oracle (MILP) benchmark. Crucially, MC-MPC accomplishes this economic efficiency without compromising service quality, delivering perfect charging reliability (zero unfulfilled SoC). It also provides the most effective peak shaving, demonstrating a robust balance of cost savings, reliability, and grid stability.

Policy	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	8 months
<i>Oracle (MILP)</i>	7700 ± 118	8293 ± 101	8300 ± 114	12015 ± 87	8716 ± 68	9278 ± 92	8072 ± 111	6639 ± 110	8645 ± 1367
MC-MPC	7776 ± 100	8516 ± 148	8622 ± 122	12121 ± 114	8828 ± 76	9384 ± 96	8169 ± 119	6862 ± 136	8785 ± 1429
RL	7817 ± 188	8446 ± 108	8512 ± 154	12052 ± 105	8896 ± 101	9398 ± 94	8221 ± 123	6899 ± 134	8808 ± 1441
LLF	7928 ± 146	8613 ± 135	8684 ± 208	12294 ± 92	8896 ± 118	9452 ± 114	8320 ± 134	7013 ± 142	8899 ± 1457
EDF	7933 ± 159	8609 ± 136	8687 ± 206	12294 ± 94	8898 ± 120	9453 ± 116	8321 ± 135	7014 ± 138	8900 ± 1457
ReqCharge	7945 ± 130	8667 ± 138	8731 ± 138	12428 ± 155	9009 ± 108	9548 ± 131	8385 ± 104	7085 ± 166	8975 ± 1482
MaxCharge	8345 ± 240	9002 ± 261	9112 ± 279	12496 ± 199	9313 ± 242	9615 ± 156	8670 ± 271	7506 ± 265	9257 ± 1389

Table 11: Base-case performance of Real-Time Policies for all months

*Higher peak shaving is better. Note that *Oracle (MILP)* is only provided as a reference; it is an omniscient control strategy and is not practically feasible.

Scenario	<i>Oracle (MILP)</i>	MC-MPC	RL	LLF	EDF	ReqCharge	MaxCharge
Missing SoC/month(in kWh)	0.0 ± 0.0	0.00 ± 0.00	0.00 ± 0.00	3.66 ± 6.04	3.64 ± 5.81	0.00 ± 0.00	0.00 ± 0.00
Cars with unfulfilled SoC	0.0 ± 0.0	0.00 ± 0.00	0.00 ± 0.00	0.70 ± 0.97	0.69 ± 0.93	0.00 ± 0.00	0.00 ± 0.00
Peak Shaving* (kW)	28.37 ± 17.11	27.91 ± 16.82	27.42 ± 16.42	15.04 ± 16.15	14.95 ± 16.14	9.25 ± 17.02	N/A

Table 12: Electricity Cost (USD) of all Policies under Various Uncertainty Scenarios.

Scenario	<i>Oracle (MILP)</i>	MC-MPC	RL	LLF	EDF	ReqCharge	MaxCharge
Departure-time uncertainty	7086 ± 1411	7334 ± 1480	7357 ± 1445	7428 ± 1423	7433 ± 1420	7534 ± 1477	8059 ± 1409
OOD departure SoC	6882 ± 1401	7073 ± 1503	7101 ± 1453	7203 ± 1464	7203 ± 1463	7246 ± 1514	7414 ± 1467
Building load noise	8281 ± 1489	8534 ± 1627	8519 ± 1529	8605 ± 1584	8606 ± 1586	8699 ± 1639	8964 ± 1585
Mid-day price spike	7040 ± 1378	7238 ± 1521	7248 ± 1499	7348 ± 1501	7349 ± 1500	7395 ± 1560	7580 ± 1506
Fleet scale-up	7089 ± 1487	7288 ± 1560	N/A	7348 ± 1501	7349 ± 1500	7395 ± 1560	7580 ± 1506

Table 13: Cumulative Missing SoC per month (kWh) over all Policies under Various Uncertainty Scenarios

Scenario	<i>Oracle (MILP)</i>	MPC	RL	LLF	EDF	ReqCharge	MaxCharge
Departure-time uncertainty	9.31 ± 23.09	9.31 ± 23.09	10.29 ± 23.09	15.12 ± 23.55	15.43 ± 23.55	9.31 ± 23.09	9.31 ± 23.09
OOD departure SoC	0.49 ± 1.07	0.49 ± 1.07	9.45 ± 5.44	3.95 ± 6.08	3.85 ± 5.91	0.49 ± 1.07	0.49 ± 1.07
Building load noise	0.00 ± 0.00	26.50 ± 18.45	26.50 ± 18.45	2.94 ± 5.09	2.96 ± 5.13	0.00 ± 0.00	0.00 ± 0.00
Mid-day price spike	7040 ± 1378	7238 ± 1521	7248 ± 1499	7348 ± 1501	7349 ± 1500	7395 ± 1560	7580 ± 1506
Fleet scale-up	0.00 ± 0.00	0.00 ± 0.00	N/A	2.94 ± 5.13	2.96 ± 5.16	0.00 ± 0.00	0.00 ± 0.00

Note that *Oracle (MILP)* is only provided as a reference; it is an omniscient control strategy that is not practically feasible.

Across all scenarios in Table 12, our proposed MC-MPC controller exhibits consistent and stable performance, incurring only 2–3% additional electricity cost compared to the offline MILP oracle. This reflects its ability to plan over sampled futures and adjust to emerging disturbances, without requiring retraining or explicit reconfiguration. While the reinforcement learning (RL) approach

demonstrates strong performance under some scenarios, particularly under building load noise, they show higher variance and less reliable SoC fulfillment in others, such as out-of-distribution SoC requests. This brittleness under distribution shift is expected: RL policies are typically trained on a fixed data distribution and can

overfit to nominal dynamics, making them less robust in deployment. Heuristic methods (EDF, LLF) and fixed rules (ReqCharge, MaxCharge) suffer in both cost and reliability. They lack foresight and adaptability, resulting in higher cumulative shortfalls in energy delivery and frequent under-charging of vehicles, especially when operating outside their nominal range.

Epistemic and structural uncertainty:

- **Mid-day price spike:** Electricity tariffs are artificially increased by 20% during the mid-day peak period (12 PM to 4 PM), simulating a sudden price surge.
- **Fleet scale-up:** The number of chargers and simultaneously connected EVs is increased by a factor of 2, stressing the system’s scalability and responsiveness.

Table 13 further illustrates MC-MPC’s advantage: even under temporal or preference variability, it meets SoC targets without compromise. This is achieved by jointly optimizing shared first-stage decisions across multiple plausible futures, rather than reacting myopically or committing prematurely. Such lookahead, combined with real-time feedback and constraint-awareness, enables MC-MPC to generalize across operational scenarios.

Overall, these results highlight the strength of MC-MPC as a general-purpose, real-time decision-making framework: it scales to large fleets, adapts to non-stationarity, and remains resilient under uncertainty, offering a strong alternative to brittle learning methods and inflexible heuristics in V2B coordination.

Table 14: MC-MPC Hyperparameter Search Results

Param	Description	Min	Max	Step	Best
\mathcal{F}	Sample size	5	50	5	30
P_{DC}	Demand penalty	0	50	5	10
P_{SOC}	SoC penalty	0	20	1	1
q	Max rate change	0	10	2.5	5

G.2 User Choices and Negotiation Behavior

We aim to understand how simulated users interact with the menu of charging options in our negotiation framework, through Figures 9, 10, and 11.

Figure 9 indicates that most users gravitate toward the *exact* fulfillment option (Option 0), suggesting a strong preference for maintaining their original charging request without deviation. Nonetheless, while the majority remain inflexible, a substantial subset is willing to accept one of the flexible options (Choices 1, 2, or 3) or even to reject all offers. This highlights the dual reality faced by system designers: exact fulfillment cannot be relied upon exclusively for system flexibility, yet there remains meaningful engagement with negotiated alternatives, providing valuable operational leeway for building operators.

Operational Deviations: Departure and SoC

Figure 10 shows the distribution of departure time deviations across all users. The sharp spike at zero deviation confirms that most users are not shifted from their requested departure, with only small fractions accepting nontrivial delays. Notably, the presence of a secondary group experiencing larger delays underscores that operational flexibility is selectively leveraged: the system offers minimal

inconvenience to most, while extracting valuable participation from a subset of more flexible users. The wide tail further emphasizes the heterogeneity in user flexibility and tolerated inconvenience.




Figure 9: Histogram of choice distribution for all users




Figure 10: Histogram of departure deviation for all users




Figure 11: Histogram of SoC deviation for all users

Similarly, Figure 11 shows the distribution of SoC deviations at departure. Note that these deviations are represented as negative values relative to the user’s requested SoC, and thus the data appear primarily in the second quadrant of the plot. Nearly all users experience negligible SoC shortfalls, shown by the pronounced peak near zero deviation. However, isolated secondary peaks at moderate and larger SoC deficits reveal that some negotiated solutions do require a minority of users to accept tangible compromises. The majority’s satisfaction is preserved, while building-level optimization gains




Figure 12: User cost and distribution upon changing the user sensitivity weights using CONSENT

cost-effective flexibility by dispersing minor inconveniences among the most adaptable users.

Analysis of Users Choosing to “Reject” Figure 13 presents the distribution of user cost and user satisfaction Y (as defined in eq. (5)) for each negotiation option, including those users who choose to reject all offers. The left figure shows that user costs remain relatively consistent across options $l = 0$ through $l = 3$, with only moderate variation. Crucially, even for users selecting the “Reject All” option, the average simulated user cost is still notably lower than the external charging benchmark cost.

This outcome is statistically robust due to the large number of Monte Carlo samples used in the simulation, ensuring convergence of empirical averages to expected values as per the law of large numbers. Consequently, the distribution of costs observed accurately reflects the underlying probabilistic model of user behavior and system dynamics, validating that participation in the negotiation process offers economic benefits on average, even when users ultimately reject all negotiated offers.

In the right figure, user satisfaction reveals a more nuanced pattern. Users who choose either the exact solution ($l = 0$) or the “Reject All” option tend to report the lowest satisfaction scores, reflecting limited perceived value or compensation for accommodating flexibility. However, in several cases, higher flexibility options (such as $l = 3$) could have delivered slightly higher user satisfaction, suggesting that some users may not be fully accounting for the additional utility available through more flexible solutions.

This result underscores two important points for negotiation-based charging strategies. First, user cost savings alone are not sufficient to guarantee satisfaction and participation, especially for those with limited willingness to defer or compromise. Second, improved communication or interface design could help users recognize and capitalize on more rewarding (but flexible) options, ultimately improving their experience and system participation rates.

Ultimately, these findings highlight the importance of aligning incentive structures and user communication, ensuring that users

not only realize economic savings, but also perceive adequate compensation for their flexibility, especially among more discriminating or conservative user segments.

Complementing these metrics, Figure 12 illustrates the distribution of user choices across negotiation options under the three scenarios. Key observations include:

- **Exact Choice (Choice 0):** The share of users selecting their exact requested profile decreases notably when sensitivity increases (from 4.1% to 1.8%), while it slightly increases under reduced sensitivity (to 4.7%). This indicates that higher sensitivity discourages strict adherence to original preferences.
- **Flexible Deviations (Choices 1–3):** The uptake of flexible negotiated options fluctuates moderately, with the most flexible option (Choice 3) gaining share as sensitivity increases (from 59.1% to 63.2%) and losing some share under reduced sensitivity (to 59.1%), suggesting that even sensitive users may accept greater deviations in specific cases.
- **Rejection (Choice 5):** Rejection rates rise from 20.5% at baseline to 21.0% with increased sensitivity and drop to 19.9% with reduced sensitivity, reinforcing that users who perceive greater inconvenience are more likely to reject offers.

In summary, these findings demonstrate that user education or interventions that effectively reduce perceived sensitivity to flexibility can lower system costs and increase participation, while increases in sensitivity generally have the opposite effects. The observed shifts in choice distribution emphasize the behavioral mechanisms influencing negotiation outcomes. These trade-offs should be carefully considered when designing demand-side management programs to optimize both system efficiency and user satisfaction.

More Details on Sensitivity to Cost-Sharing Levels. Table 3 evaluates our negotiation framework under varying proportions of cost-sharing between the building and users, from full (100%) sharing down to minimal (10%) sharing of savings. As the building’s share decreases, we observe a steady decrease in building




Figure 13: User cost and utility distribution by negotiation option (for users who chose to reject)

costs, reaching a minimum reduction of approximately 6.5% at 10% sharing, representing higher direct building savings.

Conversely, user costs increase with lower sharing percentages, indicating a shift of cost burden to users. User satisfaction and system utilization also decline gradually as sharing decreases, reflecting reduced incentives and acceptance under higher individual costs.

This sensitivity analysis highlights the trade-offs inherent in cost allocation within the negotiation design: higher building savings come at the expense of user costs and engagement. The selection of an appropriate sharing level thus balances operational savings with user participation and satisfaction, providing flexibility to adapt to deployment priorities.

Overall, these results complement our core findings by demonstrating the negotiation mechanism’s robustness and tunability across diverse cost-sharing preferences.

Interpretation and System-Level Tradeoffs

Collectively, these empirical distributions underscore several important findings:

- **User Heterogeneity:** Most users are relatively inflexible, but a significant flexible minority provides enough operational slack for meaningful negotiation outcomes.
- **Limited Impact on Most Users:** The negotiation framework maintains negligible inconvenience for the majority, ensuring high overall acceptance and satisfaction.
- **Minority Flexibility Drives Savings:** Substantial cost savings and peak-shaving benefits are made possible by the minority willing to accept time or SoC deviations. The system strategically allocates inconvenience with minimal overall dissatisfaction.
- **Participation–Cost Tradeoff:** As further quantified in sensitivity tables and rejection analysis, user participation and satisfaction are maximized when incentives are substantial and deviation boundaries are carefully set; overly aggressive cost-shifting suppresses engagement.

In summary, the distributions make plain that our negotiation-based framework achieves a robust balance: it delivers operational savings comparable to more coercive or top-down methods, but

maintains high user acceptance and satisfaction by leveraging heterogeneous flexibility and tailoring minor concessions primarily to those most willing or able to absorb them. These results complement the core findings on cost-sharing and demonstrate the framework’s adaptability for real-world deployment.

H STATED-PREFERENCE SURVEY

H.1 Section A

Thank you for choosing to take this survey. Your perspective is valuable - even if you don’t currently drive an EV. To help, we’ll provide a set of simple assumptions for you to base your responses on. Please answer using your best understanding and estimates. We’re developing an EV charging system that lets you earn discounts by being a little flexible, either by departing a bit later or accepting a slightly lower battery level. Although you can charge at home or at public stations, this survey focuses on how workplace and school charging can reward that flexibility. Your feedback will help us create a more sustainable, grid-friendly charging experience for everyone. To give you an idea of how this system works, let’s walk through an example. Imagine arriving at work/school each morning, parking in your usual spot, and plugging in your EV. On a EV companion app, you enter the driving range (or battery charge level) you need when you leave and your planned departure time. Right before charging begins, the app presents three tailored offers on your phone. A few possible choices:

- Leave on time with a full charge - regular price, say \$10.
- Leave a bit later, or take slightly less range - think of it like filling up a bit less fuel on a gas car in exchange for a moderate discount, say \$8 (The discount is greater than the amount of charge/fuel you give up).
- Leave up to an hour later, or accept a more noticeable reduction in range - like topping off your gas tank with less fuel but paying much less, say \$6 (The discount is much greater than the amount of charge/fuel you give up).
- If you think the price is too high, you can reject and not choose to charge at school/work. The nice thing is, you always have the option to come back next time and try a different choice.

When you pick one of the offers, you're making a single-day "deal" for that charging session.

1. How do you typically commute to your workplace or school?

- I drive (my own car or a shared one)
- Public transportation (bus, train, tram, etc.)
- Bicycle/Bike
- Walk
- Motorcycle/Scooter
- Rideshare (e.g., Uber, Lyft)
- I primarily work/study remotely and do not have a regular commute
- Other

2. What is your experience with electric vehicles (EVs)?

- I currently own, regularly drive and charge an EV.
- I have driven and charged an EV one or more times.
- I have an understanding of how EV charging works and have been a passenger in one.
- I have not had any personal experience driving or charging an EV.
- I am not sure.

3. What time do you typically arrive at your work or school on a regular day?

(Time hh:mm input)

4. And, how many hours would you typically spend at work or school on such a day?

(Numerical input)

5. How many miles do you usually travel to and from work or school?

(Numerical input)

H.2 Section B

For the following questions, you can base your answers on these Scenario Assumptions:

You own an EV, and a full battery gives about 300 miles of range (like a full gas tank). Work/school provides EV charging stations for employee/student use. You have a charger installed at home (used for cost comparisons). Charging prices get lower the more flexibility (departure time and EV driving range) you accept. Your anonymous answers in this survey tell us how much flexibility and compensation feel fair, so we can provide the user with fair compensation for their flexibility and design a user-friendly system. Please remember that your honest feedback in answering all the questions is valuable, whether positive or negative; we're not testing you, and we welcome all insights to help us improve.

6. On most days, what's the minimum EV range (in miles) you need before your next charging session at work/school?

- Less than 20 miles (very short direct commute home)
- 20 to 39 miles
- 40 to 59 miles
- 60 to 79 miles
- 80 to 99 miles

- 100 to 119 miles
- More than 120 miles
- I find it difficult to estimate this.

7. On a typical day, what is the maximum delay you would comfortably accept to reduce your charging costs?

- 0 minutes, leave exactly on time
- 15 minutes
- 30 minutes
- 45 minutes
- 60 minutes
- 75 minutes
- 90 minutes

8. Driving a 300-mile range EV, imagine you arrive with 30% charge and want your car charged to 80% by departure time. Which discount options would you consider fair if you left with less charge?

- I insist on my ideal 80% (240 miles range) charge.
- Accept leaving with ~75% charge (approx. 225 miles range, so 15 miles less) for a small benefit.
- Accept leaving with ~70% charge (approx. 210 miles range, so 30 miles less) for a fair benefit.
- Accept leaving with ~60% charge (approx. 180 miles range) for a significant benefit.
- I need more information to answer.

9. Please specify the ideal cost reduction (in dollars) you would expect for the benefit level you chose in the previous question, Q8 (small, fair, or significant benefit):

(If you are answering using a mobile device, tap on the slider square to activate it, and then move it)

10. As a quick check to confirm you're following along, please select option B below.

- A
- B
- C
- D

H.3 Section C

In this section, we ask you to imagine owning an EV with a 300-mile battery. You arrive with a 30% charge (about 90 miles) and aim to reach 80% (240 miles) by the time you leave. A full charge in this scenario normally costs \$10.

You will now be asked:

How much compensation would you consider fair if the system asked you to either delay your departure time or accept a slightly lower charge? This is part of a smart negotiation system: The higher the compensation you request, the less likely it is that the system accepts your offer. In some cases, no compensation may be offered if the system determines that your requested compensation is too high.

Note: The \$10 cost already covers battery degradation and maintenance. Any discount you select will be applied in addition to this amount. This representative \$10 cost is typically lower than public charging rates and may also be more cost-effective than home

charging in certain markets. Your answers below help us understand what people consider fair, enabling us to design systems that work better for both users and the grid.

11. 30 minutes delay: If you were asked to delay your departure from work/school by 30 minutes and a full charging session normally costs \$10.00, which of these discounts would feel fair?

- No monetary incentive needed - I would do it to support sustainability efforts.
- 5% off - Almost always accepted
- 10% off - Very likely to be accepted
- 20% off - Likely to be accepted
- 25% off - May be accepted in some cases
- 30% off - Unlikely to be accepted
- 35% off - Very unlikely to be accepted
- 40% off or more - Extremely unlikely to be accepted
- I would not delay my departure for 30 minutes for an incentive of this scale.
- Not applicable / Cannot estimate.

12. 15 minutes delay: If you were asked to delay your departure from work/school by 15 minutes and a full charging session normally costs \$10.00, which of these discounts would feel fair?

- No monetary incentive needed - I would do it to support sustainability efforts.
- 5% off - Almost always accepted
- 10% off - Very likely to be accepted
- 20% off - Likely to be accepted
- 25% off - May be accepted in some cases
- 30% off - Unlikely to be accepted
- 35% off - Very unlikely to be accepted
- 40% off or more - Extremely unlikely to be accepted
- I would not delay my departure for 15 minutes for an incentive of this scale.
- Not applicable / Cannot estimate.

13. 60 minutes delay: If you were asked to delay your departure from work/school by 60 minutes, and a full charging session normally costs \$10.00, which of these discounts would feel fair?

- No monetary incentive needed - I would do it to support sustainability efforts.
- 5% off - Almost always accepted
- 10% off - Very likely to be accepted
- 20% off - Likely to be accepted
- 25% off - May be accepted in some cases
- 30% off - Unlikely to be accepted
- 35% off - Very unlikely to be accepted
- 40% off or more - Extremely unlikely to be accepted
- I would not delay my departure for 60 minutes for an incentive of this scale.
- Not applicable / Cannot estimate.

14. 70% instead of 80%: Imagine your EV has a 300-mile battery. You plug in with 30% charge (90 miles) and want to charge to 80% (240 miles) for \$10. Today, the system offers to charge you instead to 70% (210 miles) for \$8. You would still have enough range to get home and run errands. What is the smallest extra discount off this \$8 charge that would feel fair for the reduced range?

- No monetary incentive needed - if the reduced range is sufficient.
- 5% off - Very likely to be accepted
- 10% off - Likely to be accepted
- 15% off - May be accepted in some cases
- 20% off - Unlikely to be accepted
- 25% off - Very unlikely to be accepted
- 30% off or more - Extremely unlikely to be accepted
- I would not accept less range for an incentive of this scale.
- Not applicable / Cannot estimate.

15. 75% instead of 80%: Similar to Q14, you plug in with 30% charge and want 80% charge and the charging cost is shown as \$10. Today, the system provides you an alternative and offers to charge you instead to 75% (225 miles) for \$9. You would still have enough range to get home and run errands. What is the smallest extra discount off this \$9 charge that would feel fair for the reduced range?

- No monetary incentive needed - if the reduced range is sufficient.
- 5% off - Very likely to be accepted
- 10% off - Likely to be accepted
- 15% off - May be accepted in some cases
- 20% off - Unlikely to be accepted
- 25% off - Very unlikely to be accepted
- 30% off or more - Extremely unlikely to be accepted
- I would not accept less range for an incentive of this scale.
- Not applicable / Cannot estimate.

16. 60% instead of 80%: Similar to Q14, you plug in with a 30% charge and want an 80% charge, and the charging cost is shown as \$10. Today, the system provides you another alternative to 60% (180 miles) for \$6. You would still have enough range to get home and run errands. What is the smallest extra discount off this \$6 charge that would feel fair for the reduced range?

- No monetary incentive needed - if the reduced range is sufficient.
- 5% off - Very likely to be accepted
- 10% off - Likely to be accepted
- 15% off - May be accepted in some cases
- 20% off - Unlikely to be accepted
- 25% off - Very unlikely to be accepted
- 30% off or more - Extremely unlikely to be accepted
- I would not accept less range for an incentive of this scale.
- Not applicable / Cannot estimate.