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

Rishav Sen
Vanderbilt University
Nashville, USA
rishav.sen@vanderbilt.edu

Fangqi Liu
Vanderbilt University
Nashville, USA
fangqi.liu@vanderbilt.edu

Jose Paolo Talusan
Vanderbilt University
Nashville, USA
jose.paolo.talusan@vanderbilt.edu

Ava Pettet
Nissan Advanced Technology Center - Silicon Valley
Santa Clara, USA
ava.pettet@nissan-usa.com

Yoshinori Suzue
Nissan Advanced Technology Center - Silicon Valley
Santa Clara, USA
yoshinori.suzue@nissan-usa.com

Mark Bailey
Nissan Advanced Technology Center - Silicon Valley
Santa Clara, USA
Mark.Bailey@nissan-usa.com

Ayan Mukhopadhyay
William and Mary
Williamsburg, USA
amukhopadhyay@wm.edu

Abhishek Dubey
Vanderbilt University
Nashville, USA
abhishek.dubey@vanderbilt.edu

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:
Rishav Sen, Fangqi Liu, Jose Paolo Talusan, Ava Pettet, Yoshinori Suzue, Mark Bailey, Ayan Mukhopadhyay, and Abhishek Dubey. 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, 9 pages. https://doi.org/10.65109/ODFH4798

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/ODFH4798

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 [13]. 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 [12]. 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 [9]. 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 [14, 17, 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, 8, 12]. 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 [6, 9]. A wide array of methods has been developed, including rule-based heuristics and mathematical programming [16, 23, 25]. More recent advances use Model Predictive Control (MPC) for real-time optimization, effectively adapting to dynamic loads, renewables, and prices [7, 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 [8] and iterative negotiation [11, 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, 30]. 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 [12, 19]. However, for the specific co-optimization challenge, MPC presents a more suitable approach. 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. For simplicity, we assume all EVs depart by the end of each day. A notation table and additional details are provided in the supplementary material [18].

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. Else, 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), where v denotes both the user

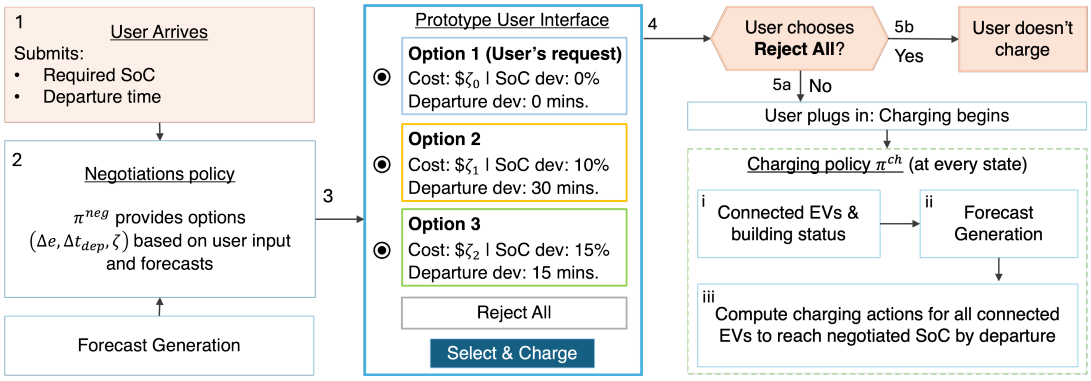

Figure 1: Negotiation workflow: Forecast-based offers on arrival; accepted offers optimally scheduled for negotiated SoC.

and their associated EV. Each user v has an arrival time t^v_{arr} and a requested departure time t^v_{req}. 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 $[e^v_{\min}, e^v_{\max}]$, arrives with initial state of charge e^v_{arr}, and requests a target SoC at departure e^v_{req}, 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 $\theta_v \in \boldsymbol{\theta}$ records the arrival time t^v_{arr}, arrival SoC e^v_{arr}, 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^v_{dep} and e^v_{req} respectively. The set of all user choices within the optimization horizon is represented using $\boldsymbol{\theta}$.

Negotiation Options & Flexibility Limits. User flexibility is parameterized by allowable deviations in target SoC, $\Delta e^{\max}_l$, and departure time, $\Delta t^{\max}_{\text{dep},l}$, where $l \in \{0, 1, \ldots, L\}$ 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 $p^b_t \in \mathcal{P}^b$ 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 $p^b_t \cdot \tau$.

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 $[c_{\min}, c_{\max}]$, where $c_{\min} = 0$ for unidirectional and $c_{\min} < 0$ for bidirectional chargers. Charging/discharging efficiency η models electrical losses as the ratio of output power to input power.

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

Battery Degradation. The penalty term K_{batt} is applied to any discharging ($r^v_t < 0$) 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 $g_t = w_t\left(p^b_t \tau + \sum_{v\in\mathcal{V}} r^v_t\right)$, where p^b_t (kW) is the building load, τ (hr) is the time step duration, and r^v_t (kWh) is the energy delivered to EV v. The energy cost attributed to EV v sums over its connection time as $g^v_t = \sum_{t\in\mathcal{T}^v} w_t\, r^v_t$.

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 $t \in \mathcal{T}^{peak} \subseteq \mathcal{T}$. Peak times also incur higher energy charges, incentivizing the shift of energy consumption to off-peak times. Demand cost is computed as $K_{\text{DC}} \cdot p^{\max}$, where $p^{\max} = \max_{t\in\mathcal{T}^{peak}}\left(b^p_t + \frac{1}{\tau}\sum_{v\in\mathcal{V}} r^v_t\right)$ 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 $\pi = (\pi^{neg}, \pi^{ch})$. The building's optimal policy $\pi^* = (\pi^{neg,*}, \pi^{ch,*})$ 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 $\boldsymbol{\theta}$. For convenience, we omit explicit notation of this dependence, writing the cost function as

$$J^\pi(\boldsymbol{\theta} \mid \mathcal{W}) = \sum_{t\in\mathcal{T}} g_t + K_{\text{DC}} \cdot p^{\max} + K_{\text{SOC}} \sum_{v\in\mathcal{V}} \left(e^v_{\text{req}} - e^v_{\text{dep}}\right) + K_{\text{batt}} \sum_{v\in\mathcal{V}} \sum_{t\in\mathcal{T}^v} q^v_t \tag{1}$$

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

User's Marginal Utility. $\boldsymbol{\theta}_v$ denotes user v's arrival and departure choices. Let $\boldsymbol{\theta}_{\setminus v}$ be the realized choices of all users who have arrived prior to v. For users arriving after v, let $\hat{\boldsymbol{\theta}}_{\setminus v}$ 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:

$$U^{\boldsymbol{\theta}_{\setminus v},\pi^*}(\theta_v) = \mathbb{E}_{\hat{\boldsymbol{\theta}}_{\setminus v}\sim\psi}\Big[J^{\pi^*}(\boldsymbol{\theta}_{\setminus v}, \hat{\boldsymbol{\theta}}_{\setminus v} \mid \mathcal{W}') - J^{\pi^*}(\boldsymbol{\theta}_{\setminus v}, \theta_v, \hat{\boldsymbol{\theta}}_{\setminus v} \mid \mathcal{W}')\Big] \tag{2}$$

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 $c \in C$ 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 $(e^v_{\text{req}}, t^v_{\text{dep}}, \zeta^v)$ 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^v_t 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 [18], 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 in Section 4.2). 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 adapt to real-world stochastic conditions as new information arrives. The charging policy is detailed in the supplementary material [18].

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

$$s_k = \big(t_k,\ p^b_{t_k},\ \{e^v_{t_k} \mid v \in \mathcal{V}_k\},\ (e^v_{\text{arr}}, e^v_{\text{req}}, t^v_{\text{req}}), p^{\text{past}}_{\max}\big)$$

where t_k is the current time step, $p^b_{t_k}$ is the building's load at time t_k, $\{e^v_{t_k} \mid v \in \mathcal{V}_k\}$ represents the states of charge of all EVs currently connected, and $(e^v_{\text{arr}}, e^v_{\text{req}}, t^v_{\text{req}})$ 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 $p^{\text{past}}_{\max}$.

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 $l \in \{1, \ldots, L\}$. 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 $\Delta e_i \leq \Delta e_l^{\max}$ and $\Delta t_{\text{dep},i} \leq \Delta t_{\text{dep},l}^{\max}$ 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 $\mathcal{P} : \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ 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 [18].

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 $\boldsymbol{\theta}_{\setminus v}$ and the predicted behavior of future users $\hat{\boldsymbol{\theta}}_{\setminus v}$. Formally, $r(s_k, a_i) = U^{\boldsymbol{\theta}_{\setminus v}, \pi^*}(a_i)$, 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.

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

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 $U(\boldsymbol{\theta}_v)$ 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 user 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 [15]. 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 $l \in \{0, 1, \ldots, L\}$, 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:

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

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 $\tilde{\theta}_v$ than their true preferences θ_v. Formally, truthful reporting always maximizes user utility, i.e., $U(\theta_v) \geq U(\tilde{\theta}_v)$.*

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 $Y_l^v \geq 0$. Thus, no user is worse off by participating; users who prefer external charging may rationally reject all offers. Some benefit strictly if $U(\theta_v) > 0$.*

Proofs. Complete proofs of Theorems 1–3 are provided in Appendix D in the supplementary material [18].

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 $\Theta = \{\theta_1, \theta_2, \dots\}$

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

5 SURVEY AND USER MODELING

To quantify user flexibility, we conducted an anonymized stated-preference survey with 28 university participants. Respondents specified their maximum tolerable deviations in departure time and SoC. They also provided 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 [18]) are used to group users into L negotiation options, indexed by l, each with calibrated upper bounds on SoC and departure time deviations, $\Delta e_l^{\max}$ and $\Delta t_l^{\max}$. 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 per-respondent inconvenience sensitivities via 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 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

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

The probability that type n user accepts option l is $P_l^n = \frac{\exp(Y_l^n)+\varepsilon}{\sum_{l=0}^{L}\exp(Y_l^n)}$, 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 [18].

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 [10] for MILP.

Runtime and Complexity. The per-decision runtime scales as $O(L \cdot |\mathcal{F}| \cdot T(n_c, n_b))$, where $n_c = O(H \cdot |\mathcal{V}|)$ continuous and $n_b = O(S \cdot H \cdot |\mathcal{V}|)$ binary variables encode charging rates, SoC trajectories, and the S-segment piecewise linear SoC approximation. With fixed S=3, IBM CPLEX solves each instance efficiently; empirical runtimes are shown in Figure 2.

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 (for MC-MPC) and 50 for testing. Only weekdays are analyzed, as employee presence, EV usage, and demand charges are negligible on weekends. Figures 8 and 9 in supplementary material [18] 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_{SOC} = \$0.5/kWh, and battery degradation penalty is set to K_{batt} = \$0.05/kWh for discharging. The external charging cost is set to $\bar{E}$ = \$0.30/kWh, reflecting the average residential EV charging rate in the area, including amortized charger installation costs. More details are provided in the supplementary material [18].

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 [18].

	Policy Description	Building Cost ($/month)	User Cost ($/kWh)	Monthly User Cost ($/month)	Reject (%)
Reference Baseline					
1	Building-Only Reference	8484 ± 2185	—	—	—
Non-Negotiation Baselines					
2	Uncoordinated Charging [3] + Free Cost	9005 ± 2184	0.00	—	—
3	Smart Charging + Free Cost [12]	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	—
Negotiation-Based Policies					
6	Menu-based Negotiation [5, 8]	8479 ± 2153	0.157	309 ± 63	34.89
7	**CONSENT (Our Approach)**	8480 ± 2171	**0.138**	**254 ± 55**	**21.99**
Ablation Study					
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.

For experimenting, we sample $\mathcal{F}^{\text{neg}} = 10$ future samples for π^{neg} and $\mathcal{F}^{\text{ch}} = 30$ 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 $(\Delta e_l^{\max}, \Delta t_{\text{dep},l}^{\max})$, indicating the maximum deviations in SoC and departure time. Specifically, option $l = 0$ is $(0, 0)$, option $l = 1$ is $(6.25\%, 30\text{ min})$, option $l = 2$ is $(10\%, 15\text{ min})$, and option $l = 3$ is $(20\%, 105\text{ min})$. An additional "Reject all" option allows users to opt out of charging at the building.

User Types. Survey analysis (see supplementary material [18]) 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 ***Smart Charging + Free Cost*** policy [12] 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, 8], 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) [12, 19] mitigates this increase but still costs more than the no EV reference. Introducing a simple, fixed price for users (Policy 4) based on the utility's energy prices [22] further reduces the building's cost, bringing it nearly to the ideal no EV baseline. Finally, using the utility-based (energy and demand charges) 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, because a portion of the increased demand charge is passed to the user.

Negotiation Framework Performance. Building on the previous baselines, the Policy 6 introduces negotiation options to the user, which implements a state-of-the-art, menu-based approach consistent with prior work [5, 8]. While this method reduces user costs, the rejection rate is high, at nearly 35%, indicating misaligned incentives.

Our proposed framework, CONSENT (Policy 7), 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 all" option (Policy 8). 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

User Profile Scenario	Bldg. Cost ($/mo)	User Cost ($/kWh)	Monthly User Cost ($/mo)	Reject (%)
CONSENT (with Reject option)				
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
CONSENT without Reject option				
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	—
Menu-based Negotiation [5, 8]				
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.

for achieving the most favorable outcome for users. Generating all options for a user takes ~0.8s on average, with the charging rate calculation at each time step requiring ~0.99s. Our method is scalable over larger problem instances, as shown in Figure 2.

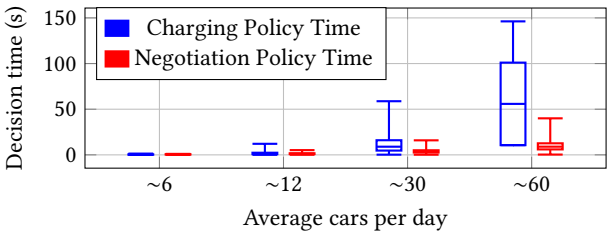

Figure 2: Comparison of computational runtime per decision across varying EV arrival intensities.

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 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 Shared (α)	Bldg. Cost($/mo)	User Cost($/kWh)	Reject(%)
100%	8480.11 ± 2171.28	**0.138 ± 0.03**	**21.99**
90%	8469.79 ± 2173.71	0.148 ± 0.02	22.14
75%	8451.28 ± 2171.93	0.158 ± 0.02	23.78
50%	8427.71 ± 2163.53	0.170 ± 0.01	24.06
25%	8411.49 ± 2155.76	0.176 ± 0.01	25.33
10%	**8397.71 ± 2147.44**	0.178 ± 0.00	26.88

Table 3: Performance across utility-sharing ratios α (100%–10%) from Eq. (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 forecast-driven method to generate personalized incentive 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 the strategy-proofness guarantee 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.

9 ACKNOWLEDGEMENTS

This material is based upon work sponsored by Nissan Advanced Technology Center-Silicon Valley, National Science Foundation under grant 2238815 and Vanderbilt Center for Sustainability, Energy and Climate. Results presented in this paper were obtained using the Chameleon Testbed supported by the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Nissan or National Science Foundation or Vanderbilt Center for Sustainability, Energy and Climate.

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, Łukasz Chełchowski, 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] 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.

[7] 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.

[8] Arnob Ghosh and Vaneet Aggarwal. 2017. Control of charging of electric vehicles through menu-based pricing. *IEEE Transactions on Smart Grid* 9, 6 (2017).

[9] 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.

[10] IBM. [n.d.]. IBM CPLEX Python API Documentation.

[11] 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.

[12] 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.

[13] 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).

[14] 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).

[15] 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

[16] 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.

[17] 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.

[18] Rishav Sen, Fangqi Liu, Jose Paolo Talusan, Ava Pettet, Yoshinori Suzue, Mark Bailey, Ayan Mukhopadhyay, and Abhishek Dubey. 2026. CONSENT: A Negotiation Framework for Leveraging User Flexibility in Vehicle-to-Building Charging under Uncertainty. *arXiv preprint arXiv:2601.01581* (2026).

[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) *(ICCPS '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 Ruszczynski. 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] 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.