IMPERIAL

LAIDLAW SCHOLARS

Network-Aware Dynamic Pricing for EV Charging on Low-Voltage Networks

Anay Gupta[1],
Dr Zohaib Akhtar[1]

[1] Electrical and Electronic Engineering,
Imperial College London

1. Background and motivation

Electricity reaches a London street through a distribution substation transformer and a set of low-voltage feeders, sized decades ago on the assumption that no household would ever draw more than a few kilowatts at once. A representative feeder carries about 75 homes within a firm limit near 101 kW.

A home EV charger draws 7.4 kW, comparable to the rest of the house combined. When enough cars charge at once, the transformer and its feeders exceed their rating; sustained overloading eventually trips the circuit, causing local outages. Reinforcing the network is slow and expensive; instead, charging demand can be controlled by pricing it by time of use.

Britain's dynamic tariffs are indexed to the wholesale market and carry no information about the local network. Every price-responsive car receives the same signal at the same moment, so the tariff itself synchronises demand: a price meant to lower bills can rebuild the peak it was designed to remove.

Unlike the tariffs sold today, a charging tariff must be cheap for the household and safe for the transformer. This study builds such a tariff.

2. Aims and objectives

This study builds a tariff scheme that considers network safety, and tests it against the tariffs sold today.

1. Model, from the bottom up, when and how much London's private electric cars charge.
2. Design a day-ahead tariff on 30-minute bands whose price is set by the capacity left on the feeder.
3. Test whether it is Pareto-efficient against flat pricing, Octopus Go and Octopus Agile, on network risk and household cost together.

3. Modelling charging demand

Demand is built session by session, in four steps. Monte Carlo simulations of 1,000 days give the distribution of load.

Step 1: Number of charging sessions

$$n \sim \text{Poisson}(Nk)$$

N is London's private electric-car fleet and k the share plugged in on a given day: most cars charge every few days, not nightly.

Step 2: Session attributes

$$s \sim \text{LN}(\mu_D, \sigma_D), \qquad \text{SOC} \sim \text{LN}(\mu_S, \sigma_S), \qquad t_{\text{start}} \sim \sum_{j=1}^{3} w_j \mathcal{N}(\mu_j, \sigma_j^2)$$

Three quantities are drawn independently per session and used to model the load curve: the distance driven s, the charge remaining on arrival SOC, and the moment the car is plugged in t_{start}.

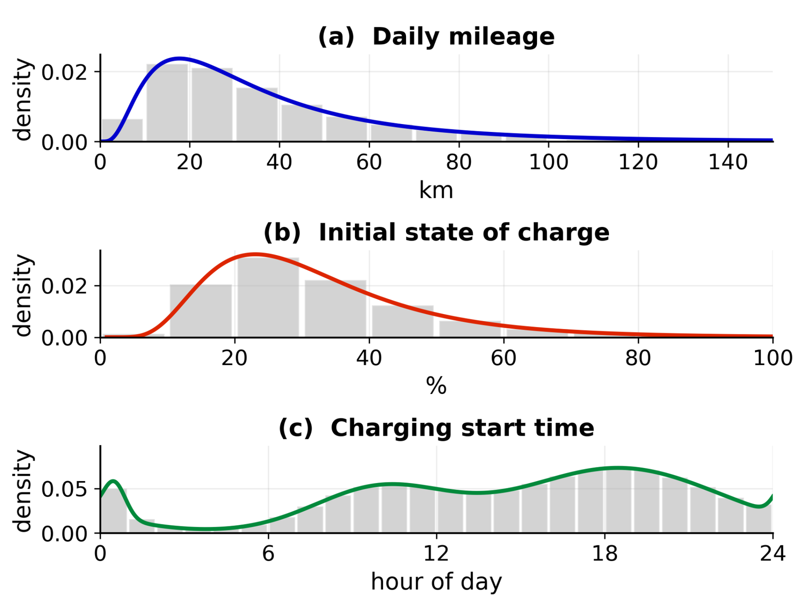

Figure 1: Model inputs. Grey bars are the observed data; coloured curves are the fitted distributions from which each session is drawn. Panels (a) and (b) are log-normal, with medians of 30 km and 33%; panel (c) is a three-component Gaussian mixture with modes at 00:30, 10:00 and 18:30.

Step 3: Energy requirement

$$E_d = \tfrac{s}{100} W_{100}, \qquad E_c = \tfrac{E_d}{k}, \qquad E_n = (1 - \text{SOC})\, E_c, \qquad t_c = \tfrac{E_n}{P\eta}$$

Because only a fraction k of the fleet charges on a given day, each charging session replaces $1/k$ days of driving, lengthening the average session from 44 minutes to 3.5 hours. The daily mileage determines energy expended daily E_d and per session E_c, thus reflecting in the charging load E_n.

Step 4: Aggregation

$$L(t) = \sum_j P_j \mathbf{1}[t \in [t_{\text{start},j},\, t_{\text{start},j} + t_{c,j})]$$

Each session is a rectangular pulse of power P_j running from plug-in until the battery is full. Summing them on a one-minute grid gives the aggregate charging curve of Figure 2.

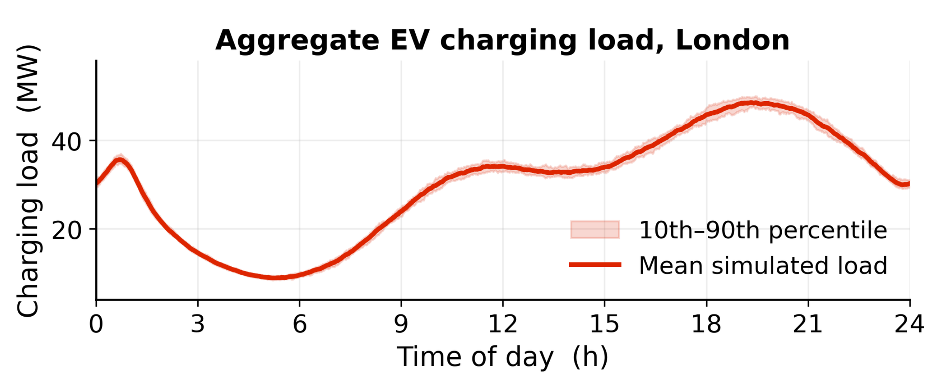

Figure 2: Output of the demand model and input to all subsequent analysis: 29,583 sessions and 720 MWh on an average day, within 0.7% of a closed-form energy balance.

Table 1: Principal model parameters. Sources: DfT and GLA; SP Energy Networks ESDD-02-012; DESNZ, 2024.

Quantity	Value
Private electric cars, London	138,240
Share charging each day, k	0.214
Consumption, W_{100}	18 kWh/100 km
Charging efficiency, η	0.92
Feeder firm limit	118.8 / 100.9 kW
Household background demand	3,240 kWh/yr

4. Design of the tariff

Go and Agile are exogenous: their prices are set outside the distribution network and convey no information about local conditions. The tariff developed here sets price from the feeder's remaining capacity.

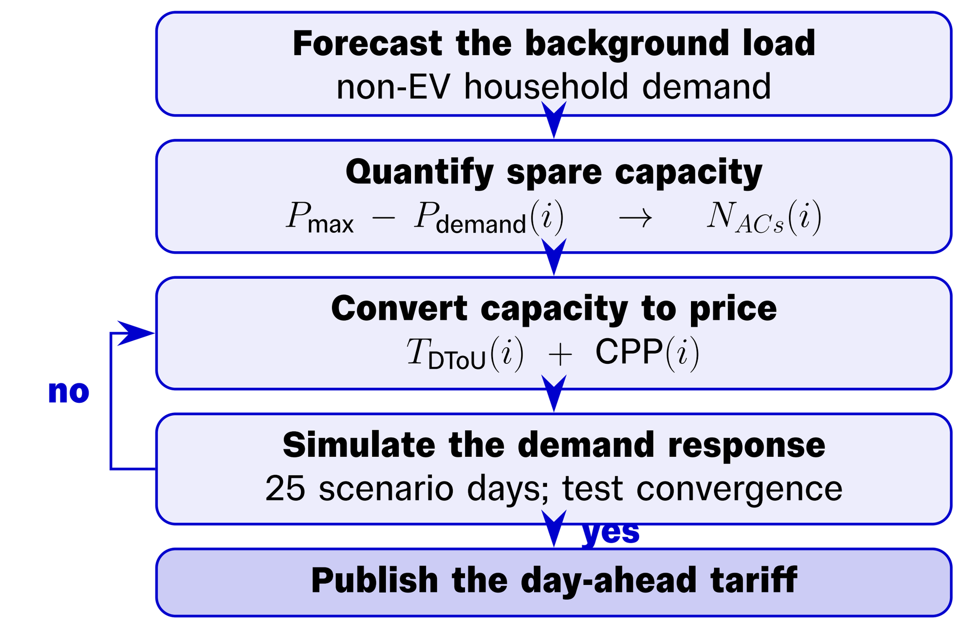

Figure 3: The tariff is not computed once. It is recomputed until consistent with the demand it induces.

Step 1: Quantifying spare capacity

$$P_{\text{surplus}}(i) = P_{\max} - P_{\text{demand}}(i), \qquad N_{ACs}(i) = \left\lfloor \frac{P_{\text{surplus}}(i)}{P_{\text{charger}}} \right\rfloor$$

P_{surplus} is the headroom under the feeder limit in half-hour i; dividing by one charger gives N_{ACs}, the households that could charge at once without breaching it — 3 at the evening peak, 11 overnight.

Step 2: Converting capacity to price

$$T_{\text{DToU}}(i) = C_1 \left| P_{\text{demand}}(i) \right|, \qquad C_1 = \frac{1.1\,\text{Bill}_{\text{static}}}{\sum_i \left| P_{\text{demand}}(i) \right| P_{\text{ch}} N_{ACs}(i)}$$

Price follows the load already on the wires, so charging is cheapest when the feeder is emptiest. Coefficient C_1 scales the curve, set so the tariff recovers 10% more profit.

Step 3: Critical-peak surcharge

$$u(i) = \frac{\hat{D}_{\text{EV}}(i)}{P_{\text{surplus}}(i)}, \qquad \text{CPP}(i) = \begin{cases} 0 & u < 0.8 \\ 25 \text{ p/kWh} & 0.8 \le u < 1.0 \\ 200 \text{ p/kWh} & u \ge 1.0 \end{cases}$$

u is the fraction of headroom forecast charging would consume. A penalty is levied for when overloading is risked by breaching the headroom. As the surcharge guards a rare event, u is assessed at the 90th percentile of the forecast.

Step 4: Solving for a consistent price

Consistency between price and demand

A published price changes the demand it was computed from: a cheap 02:00 draws every responsive car to 02:00. The tariff must satisfy a fixed-point condition,

$$D^{*} = \mathcal{S}(\mathcal{T}(D^{*})), \qquad D^{(m+1)} = (1-\alpha) D^{(m)} + \alpha\, \overline{\mathcal{S}(\mathcal{T}(D^{(m)}))},$$

reached by damped iteration. The tariff is then self-limiting: filling a cheap interval raises its own price and pushes the next car elsewhere.

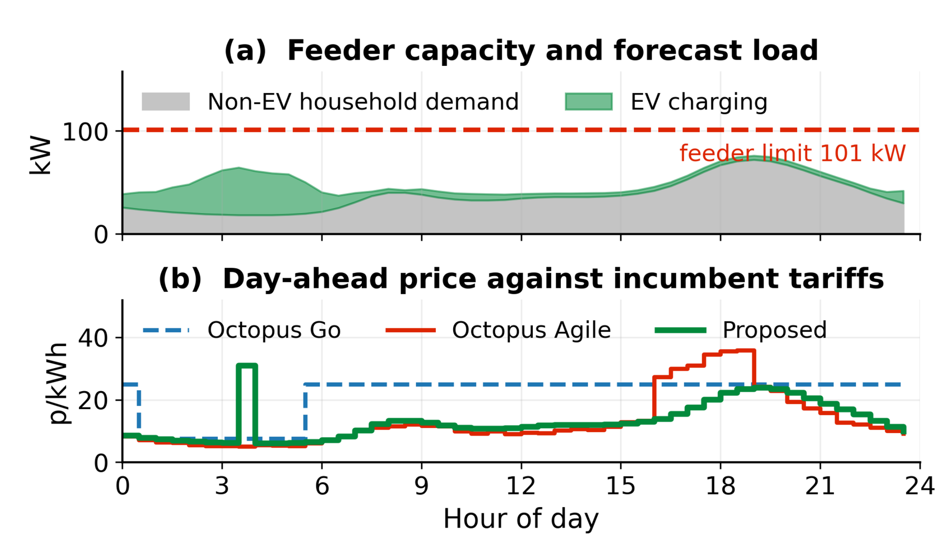

Figure 4: Spare capacity determines the price. Unlike Go and Agile, the proposed tariff is cheapest when the feeder is least loaded.

5. Comparative evaluation

Four schemes were compared over 2,000 simulated days, at six levels of household EV penetration and on two feeder cases. Eighty per cent of drivers respond to price, and 75% of those act as fixed timers where prices tie.

Principal result

Across the four schemes the average feeder peak varies by 12%, whereas the probability of exceeding the rating varies by a factor of four. Pricing acts on the tail, not the centre, of the demand distribution.

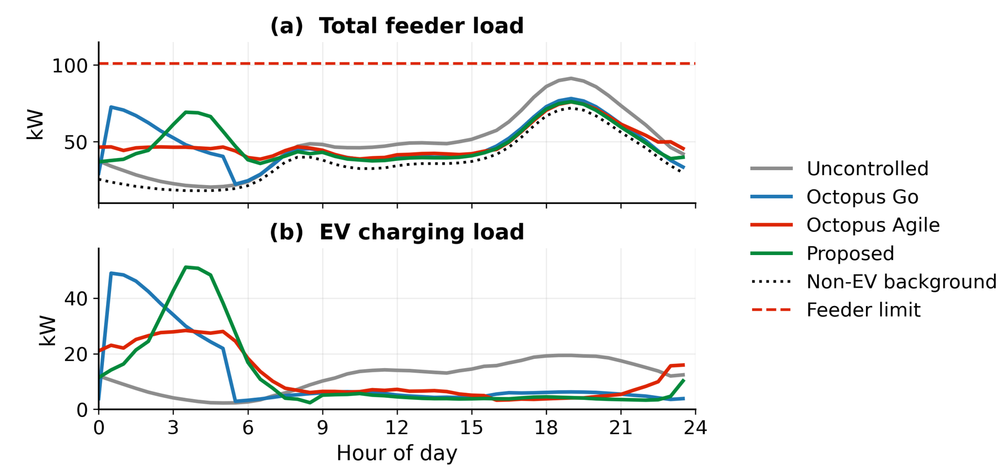

Figure 5: Mean load profiles at 60% penetration. Octopus Go concentrates the fleet into a single spike at 00:30, its entire cheap window being tied in price. The proposed tariff offers no such tie, and the same energy is distributed across the night.

Table 2: Constrained feeder, 60% of households owning an electric car, 2,000 days, firm limit 100.9 kW.

Scheme	Mean peak kW	CVaR_{95} kW	Days over limit	Cost p/kWh
Uncontrolled	94.7	125.6	30.2%	26.11
Octopus Go	85.2	118.0	9.9%	13.58
Octopus Agile	86.3	123.9	13.8%	**7.61**
Proposed	**82.5**	**115.2**	**7.5%**	9.06

Agile is cheapest on average but the most volatile: its day-to-day standard deviation is 6.16 p/kWh against 3.55 here. Its peak price is therefore the higher of the two, 17.4 p/kWh at the 95th percentile against 15.7.

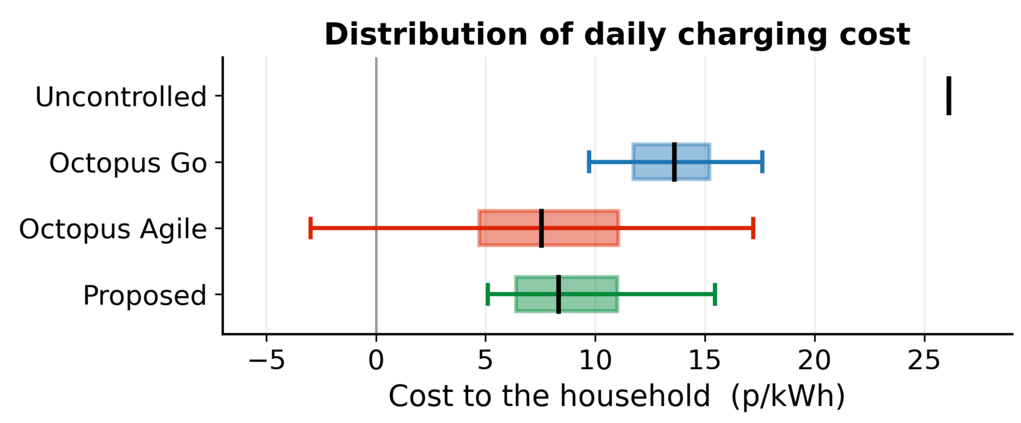

Figure 6: Daily cost to a charging household. Boxes span the quartiles, whiskers the 5th to 95th percentile, the rule the median. Agile's median is the lowest but its spread is by far the widest, and on the dearest days it is the most expensive of the three dynamic schemes. Its cheapest days are negative, because wholesale prices occasionally are.

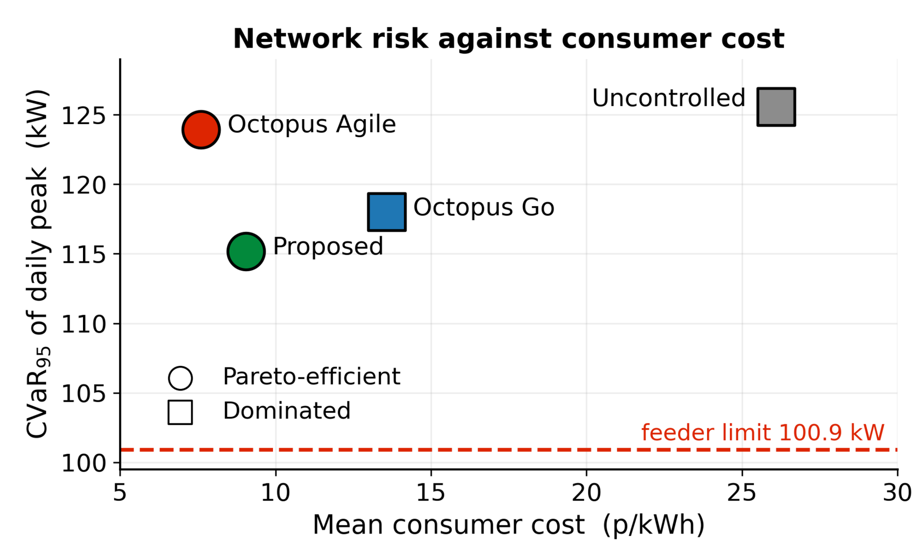

Figure 7: Risk against cost at 60% penetration. Uncontrolled charging and Octopus Go are dominated on both criteria; Agile and the proposed tariff are distinct points on a common frontier.

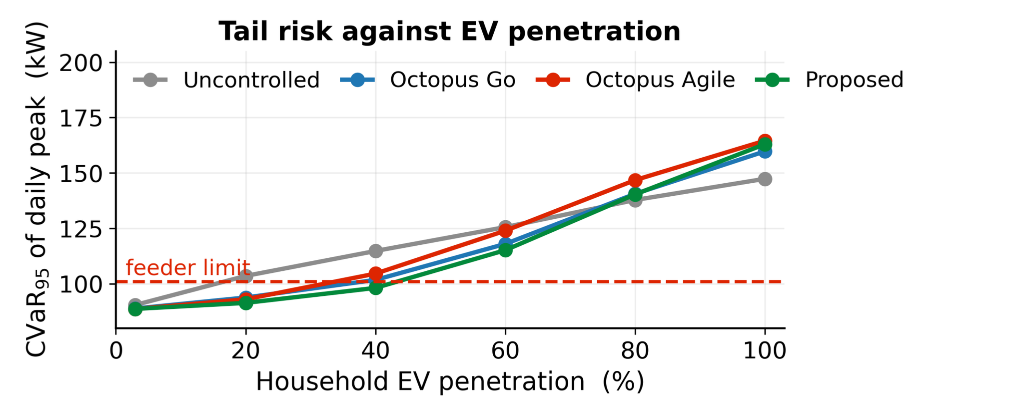

Figure 8: Limit of price-based control. Below 20% penetration the feeder absorbs every scheme; above 80% none holds risk within the rating, and reinforcement or storage is required.

6. Conclusions

The proposed tariff is the day-ahead scheme of Section 4: a price set from feeder headroom, with a critical-peak tier, solved to a fixed point. At 60% car ownership on the constrained feeder, the proposed tariff:

1. **Carries the lowest network risk of any tariff tested.** CVaR_{95} of 115.2 kW, against 118.0 for Go, 123.9 for Agile and 125.6 uncontrolled. Days over the limit fall to 7.5%, a fourfold gain on uncontrolled charging.
2. **Outperforms Octopus Go outright.** Safer on every risk measure and 33% cheaper for households: 9.06 against 13.58 p/kWh. No trade-off to weigh.
3. **Against Octopus Agile, a little average cost buys safety and predictability.** Agile is 16% cheaper on average, but its price varies almost twice as much from day to day, and its peak price is the higher of the two.
4. **Gains from the fixed point, not the peak penalty.** Solving the price against the demand it induces makes a cheap interval raise its own price. Go and Agile lack that feedback, so Agile synchronises despite a low average price.

7. Acknowledgements

I thank the Laidlaw foundation for funding this project, and Yogi Agrawal, Swati Gupta, Vinay Gupta and Tanishka Gupta for their support throughout this project.