

The Rate Wall

Why Some DC Fast Chargers Need Utilization No US Market Has Reached to Break Even, While Others Profit at Average Traffic

James Lightfoot

RightSite, United States

Correspondence: james.lightfoot@rightsite.net

Working paper, July 2026

Abstract

DC fast charging (DCFC) is widely assumed to become profitable more or less automatically as electric-vehicle adoption raises station utilization. This paper tests that assumption against the rate structures operators actually face. Using a library of 241 commercial electricity rate schedules built from utilities' own published tariffs, whose investor-owned utilities alone serve about 50% of United States electricity customers, we model the effective all-in cost of electricity for a 350 kW fast charger as a function of station load factor. At the industry's real-world utilization (16.4% of operating hours occupied by sessions, corresponding to a station load factor near 5%), the median all-in electricity cost is \$0.35/kWh, but the distribution has a heavy right tail: 27% of tariffs exceed the national average retail charging price of \$0.49/kWh, and 14% exceed \$0.60/kWh on electricity alone, before any capital or operating cost. The driver is tariff design, not the underlying price of energy: schedules with conventional demand charges carry a median cost of \$0.40/kWh at a 5% load factor, against \$0.23/kWh for EV-specific schedules that remove or cap the demand charge. A within-schedule test, removing only the demand charge from each conventional rate, attributes a median \$0.30/kWh to the demand charge itself (Section 5.2). Because a demand charge bills peak power regardless of throughput, low utilization is converted directly into high per-kWh cost. Two rate-design remedies, no-demand EV rates and demand-cap riders, flatten cost across utilization. Under a full capital-and-operating pro-forma, an average-traffic station clears its costs in about a third of territories at typical funding and half under strong subsidy, falling to roughly a fifth when capital or financing costs are stressed. Utilization is the master variable, but utility rate design sets how punishing low utilization is.

Highlights

- 241 US commercial tariffs priced for a 350 kW fast charger from published rates.
- At real-world 5% load factor, 27% of tariffs cost more than the retail price.
- Demand charges, not energy price, drive the gap: \$0.40 versus \$0.23/kWh median.
- No-demand EV rates and demand-cap riders flatten cost across utilization.
- At average traffic, a third of territories break even; half with strong subsidy.

Keywords: electric vehicles; DC fast charging; electricity tariffs; demand charges; load factor; rate design; charging infrastructure economics

JEL: Q41; Q48; L94; R48

1. Introduction

Public DC fast charging is the backbone of long-distance and no-home-charging electric-vehicle (EV) travel, and its build-out is a stated goal of federal and state policy, with national needs assessments projecting rapid required growth in fast-charging ports [9]. The dominant narrative for its economics is a utilization story: fast chargers lose money today because they sit idle most of the time, and they will become profitable as EV adoption fills them up. That narrative is half right. Utilization is indeed the single most important driver of unit cost. But it is incomplete in a way that matters for siting, investment, and policy, because it treats the electricity bill as a passive function of how much

energy flows through the site. It is not. The structure of the commercial tariff, and in particular the demand charge, determines how severely a given level of utilization is punished.

A demand charge bills the site for its highest instantaneous power draw in the billing period, in dollars per kilowatt, independent of how many kilowatt-hours the site ultimately sells. A 350 kW charger sets a large peak the first time a single vehicle pulls a fast charge, and it pays for that peak whether it then serves that one car or a thousand. At low utilization the demand charge is spread over very little energy, so the effective cost per kilowatt-hour rises sharply. This is not a market failure or an accident: it is how commercial rates were designed, for a century of loads that were nothing like a fast charger.

This paper quantifies that effect across the United States. We assemble a library of 241 commercial rate schedules directly from utilities' published tariffs, deliberately modeling the specific schedule that a roughly 350 kW, low-load-factor fast charger would actually take, including the EV-specific and demand-relief rates that a growing number of utilities now offer. We then compute the effective all-in cost of electricity as a function of station load factor and ask three questions. First, at the utilization the industry actually sees, how expensive is electricity for a fast charger, and how much does that cost vary across the country? Second, how much of that variation is explained by tariff design rather than by regional energy prices? Third, once realistic capital and operating costs and subsidy are added, in what share of territories does an average-traffic station pay for itself?

Prior work established that demand charges raise fast-charging costs and that utilization is decisive [1, 2, 5], and recent studies find US fast-charging profitability to be marginal and business-model dependent at current utilization [10, 11]. The closest prior national pricing, Borlaug et al. (2020) [1], computed a levelized cost of charging across more than four thousand commercial rates in the public Utility Rate Database, but did so before EV-specific commercial rates were widely filed and without isolating them. What has been missing is a current, schedule-level, EV-rate-aware survey, at national scale, of how far apart utility territories sit once the fast charger is placed on the rate it really takes. The answer is that they sit very far apart, and that the spread is a policy variable.

2. Background: the economics of a fast charger's electricity bill

A commercial electricity bill for a fast-charging site has three components that matter here. The first is an energy charge, in dollars per kilowatt-hour, often varying by time of use. The second is a demand charge, in dollars per kilowatt of monthly peak, sometimes subject to a ratchet that carries a summer peak forward for up to a year. The third is a fixed service charge. For most commercial customers the energy charge dominates, because their load is fairly steady and their load factor (average power divided by peak power) is high. A fast charger is the opposite kind of customer: a spiky, low-load-factor load whose peak is set by its nameplate but whose average draw, at today's usage, is a small fraction of that peak.

The consequence is that the demand charge, benign for a factory, is the defining cost for a fast charger. The National Renewable Energy Laboratory has shown that demand charges can dominate fast-charging costs at low utilization [2], and utility-scale analyses of real charging fleets reached the same conclusion earlier [5]. In response, a growing set of utilities have introduced EV-specific commercial schedules that either remove the demand charge for a phase-in period, replace it with a per-kilowatt-hour subscription, or cap billed demand at an implied load factor through a demand-relief rider. These are precisely the rates a fast charger should be modeled on, and a survey that prices fast chargers on default general-service schedules will systematically misstate both the level and the dispersion of cost. Our library is built to take the EV rate wherever one exists.

3. Data

The core dataset is a library of 241 commercial rate schedules, each encoded from the utility's own published tariff sheet rather than from a third-party aggregator, and each chosen to be the schedule a 350 kW fast-charging site would actually take in that territory. By ownership the library comprises 151 investor-owned schedules (99 in vertically integrated territories and 52 in retail-choice territories, where a competitive supply price is added to the utility

delivery charge), 56 municipal, and 34 cooperative. A utility that serves several states, or offers more than one schedule a fast charger could take, contributes a schedule for each, so the schedule count exceeds the number of distinct utilities. Investor-owned utilities serve roughly three-quarters of US electricity customers nationally [7]; the library's 85 distinct investor-owned utilities account for on the order of 80 million retail customers, roughly 50% of the approximately 162 million United States electricity customers reported by the Energy Information Administration for 2023 [8]. This is a deliberately conservative count: it credits each utility once and does not add the full customer base of multi-subsidary holding companies whose subsidiaries share a name, and the modeled municipals and cooperatives add further coverage not counted here. Coverage is broadest among the largest utilities and thins among small rural systems. The library is therefore a purposive sample of the roughly 3,000 US electric utilities, not a probability sample, so the percentages below describe the library rather than estimating a national population parameter; we treat each schedule as one observation in the headline results and report unweighted and customer-weighted robustness checks in Section 5.6. The schedule selected for each territory follows the utility's own applicability rules: among the rates open to a commercial customer at the site's service voltage and demand level, we take the EV-specific or demand-relief schedule where the site is eligible, and otherwise the general-service schedule a 350 kW load falls under.

Each card records the time-of-use energy rates, any demand charge (annual and monthly components), the fixed service charge, and, for deregulated retail-choice territories, the competitive generation or default-supply price that must be added to the utility delivery charge to obtain an all-in price. Where a utility offers an EV-specific or demand-cap schedule, the card carries the parameters of that schedule, including demand-cap divisors for EV riders. Where a rate level could not yet be confirmed against a live tariff, the source library flags it as a best estimate; these estimates are retained in the main results, because the demand-charge structure that drives the finding is confirmed even where a price level is pending, and Section 5.6 reports a robustness check that excludes every flagged schedule. Tariffs are point-in-time as of mid-2026; supply components in deregulated states reset periodically and are a source of drift discussed in Section 7. The full library is proprietary and is not released; the most expensive territories are named individually in Section 5.3 and Table 2, while all other utility-specific results are reported by state and ownership type.

4. Methods

4.1 Effective cost of electricity

For a site with billed peak power P operating at load factor LF over a year of $H = 8,760$ hours, annual throughput is $P \cdot H \cdot LF$ kilowatt-hours. The effective all-in cost of electricity per kilowatt-hour is

$$c(LF) = e_{blend} + g + [D \cdot b(LF)] / (P \cdot H \cdot LF) + S / (P \cdot H \cdot LF) \quad (1)$$

where $e_{blend} = \sum_p f_p r_p$ is the energy charge blended over time-of-use periods p using the fraction f_p of charging energy in each period, g is the competitive generation adder (zero for bundled utilities), D is the demand charge in dollars per kilowatt-year (any annual charge plus twelve times the monthly charge), $b(LF)$ is billed demand in kilowatts, and S is the annual service charge. The period weights f_p are derived from a fixed public-charging session profile (Appendix B) applied to each tariff's own time-of-use windows, so all schedules are compared on the same load shape. The arithmetic follows standard commercial-tariff accounting.

For a conventional schedule, billed demand equals the site peak, $b(LF) = P$, and the demand term collapses to $D/(H \cdot LF)$: it rises steeply as utilization falls. This single term is the rate wall. We take billed demand at nameplate in every month, a conservative choice that also captures the ratcheted schedules in the library (where a demand ratchet carries the peak forward, monthly billed demand stays at or near nameplate regardless); we do not model partial-ratchet relief, so results are an upper bound on the demand burden for non-ratcheted schedules; Section 5.6 reports a sensitivity at 0.8 times nameplate. For a demand-cap rider with divisor N (kilowatt-hours per kilowatt of billed demand), billed demand is

$$b(LF) = \min(P, (P \cdot H \cdot LF / 12) / N) \quad (2)$$

so that below a threshold load factor $LF^* = N/730$ the cap binds and the demand term becomes $D/(12N)$, a constant independent of utilization. A no-demand EV rate is the special case $D = 0$. Because the demand term is billed-demand-invariant per unit throughput, $c(LF)$ is essentially independent of site size (only the small service term scales); we therefore report $c(LF)$ per 350 kW dispenser and note where site size enters (Section 4.3).

4.2 From utilization to load factor

Utilization is reported by Paren as the share of a station's operating time spent in charging sessions, counting both successful and unsuccessful charging activity; for the full year 2025 this averaged 16.4% across the United States, from data covering more than 95% of US DCFC ports [6]. We assume 24-hour operation so this maps onto the 8,760-hour year. Because this session-time measure includes occupied-but-not-delivering time, the true energy-delivered load factor is if anything lower than our estimate, which would only deepen the finding. Utilization is not load factor: during an active session the average delivered power is well below a 350 kW nameplate, because batteries taper as they fill and the vehicle mix rarely sustains peak power. Reported average delivered power on high-power hardware sits well under nameplate, implying an average-to-peak (taper) ratio near 0.3; because this ratio is uncertain we sweep it from 0.25 to 0.40 (Section 5.6) [1, 6]. Then

$$LF = u \cdot (P_{avg} / P_{peak}) \approx 0.164 \cdot 0.31 \approx 0.05 \quad (3)$$

We treat a 5% load factor as the real-world average operating point, 10% as a genuinely busy corridor site (about double the national average), and 15% as exceptional. Because the taper ratio is uncertain, Section 5.6 reports the operating point and headline results across a taper band of 0.25 to 0.40 (load factors of about 4% to 7%); the qualitative findings depend on the low-load-factor regime, not on the precise conversion.

4.3 Full pro-forma

To move from the electricity bill to station viability we add capital and operating costs for a single, explicitly defined station: three 350 kW dispensers, a station billed peak of 1,050 kW, and annual throughput of $1,050 \cdot H \cdot LF$ kilowatt-hours. Because $c(LF)$ is billed-demand-invariant per unit throughput (Section 4.1), the per-kWh cost is unchanged from the single-dispenser figures; only the fixed service charge, counted once for the station, differs negligibly. Installed capital is taken as \$675,000, within reported hardware-plus-installation ranges for 350 kW equipment [3, 4], and is understood to exclude major utility interconnection and transformer upgrades, which can be large and are treated in Limitations. Fixed operating cost (network, warranty, maintenance, insurance) is \$36,000 per year, and the base retail price charged to drivers is the Paren full-year 2025 national average of \$0.49/kWh [6]. A fraction ϕ of capital is covered by grants; the remainder is financed at 4% over ten years. The station clears its costs when

$$1,050 \cdot H \cdot LF \cdot (R - c(LF)) - O - A(\phi) > 0 \quad (4)$$

where R is the retail price, O is annual operating cost, and $A(\phi)$ is the annual debt service on the unfunded capital. Equation (4) is evaluated for every schedule at each load factor, funding level, and sensitivity case.




Figure 1. Effective all-in electricity cost versus station load factor for four representative tariff types. A high-demand-charge schedule and a moderate wires-demand schedule are ruinous at real-world utilization and only approach the retail band as the station fills. An EV demand-cap rider is flat below its cap threshold, and a no-demand EV rate is flat everywhere. The vertical line marks the real-world average near 5% load factor; the shaded band is the typical retail charging price. Territories are shown by type, not name.

5. Results

5.1 The cost of electricity at real-world utilization

At a 5% load factor the median all-in electricity cost across the 241 schedules is \$0.35/kWh, with a mean of \$0.39/kWh and an interquartile range of \$0.25 to \$0.51/kWh. The distribution is not symmetric. It has a long right tail (Figure 2): 27% of schedules exceed the national average retail price of \$0.49/kWh, and 14% exceed \$0.60/kWh, on electricity alone and before a single dollar of capital or operating cost. A fast charger placed in the wrong territory can pay more for raw electricity than it is able to charge its customers for a finished fast-charge session.




Figure 2. Distribution of effective all-in electricity cost across the 241 modeled schedules at 5% load factor. The median (\$0.35) sits below the retail band, but a heavy right tail of demand-charge-driven schedules pushes 14% of territories above the top of the band.

5.2 Tariff design, not energy price, drives the tail

Sorting schedules by structure shows where the tail comes from (Table 1). Schedules that levy a conventional demand charge (191 of 241) have a median cost of \$0.40/kWh at a 5% load factor. Schedules that remove the demand charge through a no-demand EV rate (45 schedules) have a median of \$0.23/kWh, and the demand-cap riders (5 schedules) sit at \$0.23/kWh. The gap between a conventional and a no-demand rate, about \$0.17/kWh at real-world utilization, is a structural penalty imposed by rate design, not by the cost of generating electricity in that region. The same physical kilowatt-hour is priced about 70% higher, from \$0.234 to \$0.399 per kWh, purely because of the shape of the tariff it flows through.

A between-group comparison like this can be confounded: utilities that offer no-demand EV rates may sit in structurally cheaper territories to begin with, so part of the gap could reflect where the rates are offered rather than the rates themselves. A cleaner test holds the schedule fixed and removes only its demand charge. Setting the demand charge to zero on each of the 191 conventional schedules and re-pricing, the median cost falls by \$0.30 per kilowatt-hour at real-world utilization (interquartile range \$0.18 to \$0.44). This within-schedule figure carries no selection bias, and it is larger than the between-group gap, which confirms that the penalty is the demand charge itself rather than a difference between the territories that adopt each kind of rate.

Tariff structure	Schedules	Median \$/kWh at 5% LF
Conventional demand charge	191	\$0.40
EV rate, no demand charge	45	\$0.23
EV demand-cap rider	5	\$0.23

Table 1. All-in electricity cost by tariff structure at real-world (5% load factor) utilization.

5.3 The rate wall, territory by territory

Figure 3 and Table 2 name the worst walls. At real-world utilization the most expensive territories, led by Glendale Water & Power in California, reach \$0.72 to \$1.10/kWh of all-in electricity cost, from well above to more than double the retail price they could plausibly charge, even allowing for the higher session prices (up to roughly

\$0.62/kWh) seen in supply-constrained states such as New Jersey. They are a mix of high-cost municipal utilities, deregulated Northeast and Mid-Atlantic territories where a delivery demand charge stacks on top of default supply, and conventional Midwestern and Mountain-West schedules. What unites them is not geography or fuel mix: it is a demand charge with no EV relief. Because the library takes the EV-relief rate wherever a utility offers one, every territory on the wall is, by construction, one where no EV-relief rate was found. The wall is a map of missing rate reform.




Figure 3. The fourteen most expensive territories in the library at real-world (5% load factor) utilization. Hatched bars remain above the retail price even at a busy 10% load factor (a deeper, energy-price problem); solid bars clear retail once busy, meaning their wall is built almost entirely of demand charges that amortize with throughput. Colors are colorblind-safe and the two categories differ by fill pattern as well as hue.

Utility (state, type)	Schedule	\$/kWh 5% LF	\$/kWh 10% LF
Glendale Water & Power (CA, municipal)	PC-1-A Large Business	\$1.10	\$0.63
Riverside Public Utilities (CA, municipal)	Schedule TOU Large General & Industrial	\$0.87	\$0.50
Fort Collins Utilities (CO, municipal)	GS50 General Service 50-749 kW	\$0.86	\$0.46
Alliant Energy / Interstate Power & Light (IA, investor-owned)	Large General Service	\$0.83	\$0.44
Pasadena Water & Power (CA, municipal)	L-2 Large Secondary	\$0.80	\$0.48
PSE&G (NJ, deregulated IOU)	LPL-S delivery + BGS default supply	\$0.78	\$0.46
Burbank Water & Power (CA, municipal)	L-TOU Large General Service	\$0.77	\$0.48
Rhode Island Energy (RI, deregulated IOU)	G-32 Large Demand delivery + LRS	\$0.75	\$0.48
City of Naperville (IL, municipal)	GS2 General Service Level 2	\$0.75	\$0.45
AES Indiana (IN, investor-owned)	Rate SL Secondary Large	\$0.74	\$0.41
Great Lakes Energy (MI, cooperative)	LP Large Power Service	\$0.74	\$0.43
Potomac Edison (MD, deregulated IOU)	Schedule PH delivery + SOS supply	\$0.73	\$0.40
Black Hills Energy (Cheyenne LF&P, WY) (WY, investor-owned)	Schedule SG Secondary General	\$0.73	\$0.41
Upper Peninsula Power (MI, investor-owned)	Cp-U Large Commercial & Industrial TOU	\$0.72	\$0.42

Table 2. Highest all-in electricity cost at real-world utilization, with the same station's cost once busy (10% load factor). Every listed territory lacks an EV-relief rate in the library. Costs computed from each utility's own published tariff sheet as of mid-2026.

5.4 Utilization is the master variable

Raising utilization dissolves most, though not all, of the wall. At a busy 10% load factor the median cost falls to \$0.25/kWh and only 2 schedules (about 1% of the library) still exceed the retail price on electricity alone. That residual set is different in kind: it is California municipal utilities where a high underlying energy price, rather than the demand charge, drives the residual gap. Even there the energy price alone stays below retail (Section 5.5), so these territories too clear at high enough utilization; they are simply the slowest to get there. For the large demand-charge tail, throughput is the cure, because a demand charge spread over twice the energy is half the burden per kilowatt-hour. This is the sense in which the utilization narrative is correct. The wall it must climb, however, is set by the tariff, and for a station stuck near a 5% load factor for years while adoption grows, the tariff is the difference between survival and closure.

5.5 Adding capital and operating cost

Electricity is necessary but not sufficient. Table 3 evaluates the full pro-forma of Equation (4) for the 1,050 kW station. Under typical funding, in which grants cover about a fifth of capital, an average-traffic station (5% load factor) clears its costs in 34% of territories; a busy station clears in 88%, and an exceptional station in 97%. Even under strong funding covering three-fifths of capital, the average-traffic station pencils out in 50% of territories, while busy and exceptional stations reach 92% and 98%. These base cases are optimistic in that they exclude interconnection cost; under the stress cases of Section 5.6 the average-traffic share falls to roughly a fifth. The policy reading is that capital subsidy, the main instrument of federal charging programs, moves the average-traffic outcome only modestly, because the binding constraint at low utilization is the recurring electricity bill, not the up-front hardware.

This ordering does not depend on a low cost of capital, which would be the obvious objection to a pro-forma financed at 4%. Re-run at a realistic hurdle rate of 8 to 10%, the average-traffic viability share under typical funding falls to about a fifth (21% to 18%), and tripling the grant to three-fifths of capital lifts it only to about 47%; moving the same territories to a no-demand EV rate instead lifts it to about 89%. Rate design beats capital subsidy at every financing cost we tested, and the margin widens as the cost of capital rises, so the conclusion is a finding rather than an artifact of the financing rate.

Framed as required utilization, the spread is stark. About 12% of territories need a load factor above 10% (roughly 32% utilization, the busiest US state average) to break even at typical funding, and about 3% need a load factor above 15% (roughly 48% utilization, beyond any US market average); the worst require a utilization no US market has approached. No territory, however, is unprofitable at every utilization: because no tariff's energy price alone exceeds the retail price (0 of 241 schedules), sufficient throughput eventually clears the electricity bill everywhere. The barrier for the worst territories is that the required utilization is physically out of reach at today's demand, not that profit is mathematically impossible. This is the sense in which some fast chargers cannot be made to pay while others clear their costs at ordinary traffic.

Funding scenario	Avg (5% LF)	Busy (10% LF)	Exceptional (15% LF)
Typical (20% of capex)	34%	88%	97%
Strong (60% of capex)	50%	92%	98%

Table 3. Share of modeled territories where a three-dispenser, 1,050 kW station clears capital, operating, and electricity cost, by funding level and traffic. Retail \$0.49/kWh; capital \$675,000 (excluding interconnection); operating \$36,000/yr; financing 4% over ten years.




Figure 4. The same pro-forma, drawn. Subsidy shifts the average-traffic bar but cannot lift it to a majority; only utilization moves a station reliably into the black.

5.6 Sensitivity and robustness

The central finding is robust to the main assumptions (Table 4). Moving the retail price across its observed range (\$0.45 to \$0.55/kWh) changes the share of tariffs above retail from 33% to 20% and the average-traffic viability share from 19% to 48%. Varying the taper ratio from 0.25 to 0.40 (real-world load factor 4% to 7%) moves the median cost from \$0.40 to \$0.30/kWh. Stressing the pro-forma (capital up 30%, operating up 50%, or financing at

8%) cuts the average-traffic viability share to 18%, 19%, and 21% respectively. Taking billed demand at 0.8 times nameplate, which reflects the vehicle-limited monthly peaks (roughly 250 to 320 kW on a 350 kW port) that real sites often set rather than the full nameplate, lowers the share of tariffs above retail at a 5% load factor from 27% to 17% (Table 4). Across all cases the average-traffic station remains viable in a minority of territories, and the demand-charge penalty of Section 5.2 is unaffected.

Weighting sharpens the picture, and the comparison is made like for like. Unweighted, 27% of all tariffs and 26% of the investor-owned tariffs alone exceed retail at a 5% load factor; the two are close, so the tail is not merely a small-utility artifact. Weighting the investor-owned utilities by retail customers (Energy Information Administration data [8], each utility counted once), the share falls to 13%, because the most expensive tariffs are concentrated in smaller systems while the largest investor-owned utilities are, on average, less punishing; the customer-weighted figure therefore understates the number of distinct territories affected. The results also reproduce the established direction of the literature: demand charges dominate DCFC cost at low utilization, as in Borlaug et al. (2020) and Borlaug and Bennett (2022) [1, 2], and the marginal, subsidy-dependent viability we find echoes recent US assessments [10, 11]. The pricing itself applies each utility's own published tariff formulas directly and identically across schedules, with no fitted parameters, and the pricer is strict: a malformed rate value halts the run rather than defaulting to zero, so no schedule can be silently reclassified as a no-demand rate, and on the full library no value required coercion. To guard against encoding error, each schedule was transcribed from the utility's published tariff sheet with its source recorded, and any value that could not be confirmed against the primary document was flagged in the library as a best estimate. Excluding all 70 flagged schedules leaves the finding intact, if anything slightly stronger: the share of tariffs above retail at a 5% load factor moves from 27% to 29% and the median from \$0.35 to \$0.36, because the flagged schedules are on average cheaper than the rest. The released reproducibility sample (Data availability) lets a reader re-price a subset directly from the encoded parameters and check it against the source tariff.

Assumption	Case	Result
Retail price	\$0.45 / \$0.49 / \$0.55 per kWh	tariffs above retail 33% / 27% / 20%
Taper ratio	0.25 / 0.31 / 0.40 (LF 4% / 5% / 7%)	median cost \$0.40 / \$0.35 / \$0.30
Capital +30%	avg traffic, typical funding	viability 18% (base 34%)
Operating +50%	avg traffic, typical funding	viability 19%
Financing 8%	avg traffic, typical funding	viability 21%
Billed demand 0.8x nameplate	vehicle-limited monthly peak	tariffs above retail 17% (base 27%)

Table 4. Sensitivity of the headline results to the main assumptions. The average-traffic station is viable in a minority of territories in every case.

6. Discussion

Three implications follow. First, siting is a rate-map problem, not only a traffic-map problem. Two sites with identical vehicle counts can differ by more than \$0.50/kWh in electricity cost purely because of the tariff on each side of a utility border. Any model that projects fast-charging profitability from utilization while holding electricity cost fixed at a national average will be optimistic in the demand-charge tail and pessimistic in the EV-rate territories, and it will misrank sites.

Second, the most effective near-term lever for charging economics may not be more capital subsidy but tariff reform. The no-demand EV rates and demand-cap riders modeled here already exist; where they are offered, they flatten the cost curve and remove the wall (Figure 1). Extending well-designed EV commercial rates to the demand-charge territories in Table 2 would do more for a struggling charger's survival, dollar for public dollar, than an equivalent grant toward hardware it cannot afford to operate. This is consistent with the demand-charge literature [2, 5] and gives it a concrete, current, national target list; the most exposed territories are named in Table 2. Some of those territories already run temporary demand-charge rebates rather than reformed tariffs: PSE&G in New Jersey, for

instance, refunds three-quarters of a fast charger's demand charges, but only through the end of 2026, after which the underlying schedule and its full wall return. That is exactly the difference this paper draws between a subsidy, which lapses, and rate design, which does not.

Third, a residual problem is not fixable by rate design. The territories that remain above retail even at a busy 10% load factor are high energy-price regions, chiefly California municipal utilities, where the underlying cost of a kilowatt-hour, not its tariff wrapper, exceeds what drivers pay. There, fast-charging economics wait on generation cost, on-site solar and storage, or higher retail prices, and no EV rider will close the gap.

Set against the larger picture, none of this says fast charging is doomed. It says the path to viability runs through utilization, and that the steepness of that path is set by a policy variable that is currently underweighted in most siting and most subsidy design. The optimistic reading (chargers fill up, costs fall) and the pessimistic reading (chargers lose money, the model is broken) are both too coarse. The same charger can be a sound investment on one side of a rate boundary and an unrecoverable loss on the other, and which one it is can be known in advance.

7. Limitations

Several caveats bound these results. The mapping from utilization to load factor (Equation 3) depends on an assumed taper ratio; Section 5.6 reports the sensitivity, and load factors are stated directly so readers can re-map. A single national session shape blends time-of-use energy rates, which understates sites that can shift charging into off-peak windows and overstates those that cannot. Billed demand is taken at nameplate, an upper bound for non-ratcheted schedules. The station is modeled as three independent 350 kW dispensers; a power-shared cabinet would lower both billed demand and throughput, with a smaller effect on cost per kilowatt-hour than on absolute margin. Capital cost excludes utility interconnection and transformer upgrades, which can be substantial for a megawatt-class site and would lower the viability shares; taxes, incentives beyond the modeled capital grant, and site rent are also excluded. The pro-forma is single-site and ignores portfolio effects across a network. Retail and cost inputs are national averages, so the pro-forma is best read as a comparative sorting of territories, not a site-level forecast; the retail benchmark is a single national average, and because published session prices vary by state (roughly \$0.45 to \$0.62 per kWh), a state-specific retail comparison would shift the per-territory margin and trim the exact share above retail. The pro-forma station draws 1,050 kW, which in some territories would qualify for a primary-voltage schedule different from the secondary schedule a single 350 kW dispenser takes; about 15% of the 241 schedules carry a stated eligibility ceiling below 1,050 kW, well below the share in the deliberately large-schedule-weighted public sample. We hold the schedule fixed; for the few no-demand EV rates so capped this makes the pro-forma slightly optimistic, while for a conventional schedule a larger site usually moves to a primary-voltage rate with a lower per-kilowatt demand charge, which works the other way. Deregulated supply prices are point-in-time and drift with quarterly resets. Utilization is Paren's session-time measure, mapped to a calendar year under a 24-hour-operation assumption. Finally, each schedule is one observation; the customer-weighted check in Section 5.6 shifts the exact percentages but not the central finding that tariff design produces a wide, policy-driven dispersion in cost.

8. Conclusion

Across 241 utility schedules priced from their own tariffs, the cost of electricity for a fast charger at real-world utilization ranges from routine to prohibitive, and the range is governed by demand-charge design more than by the price of energy. At a 5% load factor, 27% of territories charge more for raw electricity than a fast charger can sell a session for, and removing the demand charge from a conventional schedule lowers its cost by a median of \$0.30/kWh. Utilization remains the master variable, and rising adoption will relieve most demand-charge sites in time. But the wall those sites must climb, and whether some can climb it at all, is a choice embedded in utility rate design. Extending EV-relief rates to the territories that lack them is the most direct public-policy lever available, and one that current capital-subsidy programs do not pull.

CRediT authorship contribution statement

James Lightfoot: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing - original draft, Writing - review and editing, Visualization.

Declaration of competing interest

The author operates RightSite, a commercial EV-charging site-selection service that uses the tariff library and pricing methodology described in this paper, which constitutes a competing financial interest in the subject matter of this article.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors; the work was self-funded by RightSite.

Data availability

The analysis code and a public reproducibility sample, comprising the 14 named schedules of Table 2 plus 25 randomly selected others (39 schedules with full parameters), are openly available in the GitHub repository <https://github.com/rightsitera-arch/rate-wall-reproducibility>, archived at Zenodo (DOI 10.5281/zenodo.21667465), and reproduce the reported per-schedule costs and the pricing method. The full 241-schedule library is a proprietary dataset compiled from utilities' publicly published tariff sheets and is available for verification under a confidentiality agreement; all source tariffs are public documents obtainable from the respective utilities.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used a generative AI assistant (Claude, Anthropic) to help write the analysis code, generate figures, conduct literature searches, and draft and edit the manuscript. After using this tool, the author reviewed, verified, and edited all content, independently checked the quantitative results and cited sources, and takes full responsibility for the content of the publication.

References

- [1] Borlaug, B., Salisbury, S., Gerdes, M., & Muratori, M. (2020). Levelized Cost of Charging Electric Vehicles in the United States. *Joule*, 4(7), 1470–1485. <https://doi.org/10.1016/j.joule.2020.05.013>
- [2] Borlaug, B., & Bennett, J. (2022). *EV Charging and the Impacts of Electricity Demand Charges*. National Renewable Energy Laboratory, NREL/PR-5400-82738.
- [3] Idaho National Laboratory (2022). *Breakdown of Electric Vehicle Supply Equipment Costs*. INL/RPT-22-68598.
- [4] Wood, E., et al. (2023). *The 2030 National Charging Network: Estimating U.S. Light-Duty Demand for Electric Vehicle Charging Infrastructure*. National Renewable Energy Laboratory, NREL/TP-5400-85654.
- [5] Fitzgerald, G., & Nelder, C. (2017). *EVgo Fleet and Tariff Analysis: Phase 1, California*. Rocky Mountain Institute.
- [6] Paren, Inc. (2026). *State of the Industry Report: U.S. EV Fast Charging, Full-Year 2025*. Paren, Inc. (national average utilization 16.4%, average price \$0.49/kWh, coverage >95% of US DCFC ports).
- [7] Edison Electric Institute (2023). *Electric Company Industry Statistics*. EEI, Washington, DC.
- [8] U.S. Energy Information Administration (2024). *Electric Power Annual* and Form EIA-861 (2023 data): utility sales and customer counts.
- [9] Bauer, G., Hsu, C.-W., Nicholas, M., & Lutsey, N. (2021). *Charging Up America: Assessing the Growing Need for U.S. Charging Infrastructure through 2030*. International Council on Clean Transportation.
- [10] Bernal, D., Raheem, A. A., Inti, S., & Wang, H. (2024). Assessment of Economic Viability of Direct Current Fast Charging Infrastructure Investments for Electric Vehicles in the United States. *Sustainability*, 16(15), 6701. <https://doi.org/10.3390/su16156701>

[11] Sawant, V., & Zambare, P. (2024). DC fast charging stations for electric vehicles: A review. *Energy Conversion and Economics*, 5(1), 54–71. <https://doi.org/10.1049/enc2.12111>

Appendix A. Model and assumptions

Pricing. Effective cost follows Equations (1) and (2). Energy is time-of-use, blended over each tariff's peak/mid/off windows using the session profile of Appendix B. Deregulated territories add a competitive generation or default-supply price. Demand is a flat dollar-per-kilowatt charge on billed demand, capped where an EV demand-relief rider applies, and taken at nameplate in every month otherwise.

Operating points. Load factors of 5%, 10%, and 15% correspond to real-world average, busy, and exceptional utilization under Equation (3), from a 16.4% utilization anchor and a 0.31 taper ratio. **Retail band.** Retail reference band \$0.43 to \$0.60/kWh, bracketing the \$0.49 national average; published state-level session prices span roughly \$0.45 to \$0.62 [6]. **Station.** Three 350 kW dispensers, 1,050 kW billed peak; capital \$675,000 (equipment plus installation, excluding interconnection); operating \$36,000 per year; financing 4% over ten years; funding fractions 20% and 60% of capital.

Library. 241 commercial rate schedules encoded from utilities' published tariffs: 151 investor-owned, 56 municipal, 34 cooperative, each set to the schedule a 350 kW fast charger would take, including EV-specific and demand-cap rates where offered. Identities are reported by state and ownership type.

Appendix B. Session profile

The time-of-use energy blend uses the following stylized public-charging session weights by hour of day, a daytime-peaked shape consistent with published US public-DCFC session-start distributions and the DCFC load profiles of Borlaug et al. (2020) [1]. Weights are relative and normalized within each tariff's time-of-use windows.

Hour	0	1	2	3	4	5	6	7	8	9	10	11
Weight	0.4	0.4	0.4	0.4	0.4	0.5	0.8	1.1	1.3	1.5	1.7	1.8
Hour	12	13	14	15	16	17	18	19	20	21	22	23
Weight	1.9	2.0	2.0	2.0	1.9	1.8	1.6	1.4	1.2	0.9	0.7	0.5