From Lagging to Leading: The Measured Emergence of Standing Capacitance Behind Consumer Connections, 1997–2025

D. J. Hume

***Abstract*—Both ENTSO-E Final Reports on the 2025 Iberian and North Macedonian blackouts read the excess capacitive reactive power as a network term—the reading the available data supports. The Iberian report asks for what was missing: information on load composition. We build that record from regulatory half-hourly metering at New Zealand's transmission–distribution boundary, spanning 29 years, and find a second source: the demand behind the customer's connection, at mains voltage and in residential, commercial and industrial premises alike, has itself moved capacitive. Overnight reactive power crossed from lagging to leading in 2016, nearly a decade before the 2025 events, and a first-principles charging calculation puts the demand-side share of that drift in the physical band 75–88%. That term behaves as a net standing shunt element—present around the clock, scaling with the connections served, and moving capacitive by about 230 VAr per connection over 2013–2025 and roughly 20 VAr a year at the close. It drives the drift rather than the level: the distribution network's own cable and line charging remains the larger part of the standing leading reactive power, and what changed is that a demand once strongly lagging stopped masking it. The candidate device layer is the electromagnetic-compatibility filter mandated in nearly every mains-connected product (a labelled hypothesis): a few VAr per device, tens of devices per household, millions across a system and perhaps billions on the largest—permanently across the mains and switched by nobody. Its injection rises with the square of mains voltage, so every 230-volt system accumulates about three times more per device than a 120-volt, 60-hertz one. Planning and protection studies still carry a constant lagging load archetype, where the overnight flow now has the opposite sign. Both inputs to the measurement are records any operator already holds.**

***Index Terms*—Reactive power, power factor, overvoltage, load modeling, power distribution, power system measurements.**

I. Introduction

POWER systems internationally are shifting from net-absorbing (lagging) toward net-generating (leading) reactive power at light load, as the connected load fleet modernises toward power-electronic devices and networks add underground cabling. The direction is not in dispute, and this paper does not claim to be the first to observe it. What has been missing is a long *measured* record of the transition at supply-point resolution, and a measured separation of how much of the drift is the load fleet and how much is the network. This paper supplies both for a single national system, connects the result to a failure mode that has recently moved from theory to realised blackouts, and closes with what the reversal implies for the analysis and defences built for the opposite regime. The trail leads somewhere unexpected: to the electromagnetic-compatibility filter inside practically every mains-connected device—capacitance that one part of the regulatory apparatus drives in and no other part counts—advanced here as a labelled hypothesis for one limb of the measured term, with its decisive test identified.

D. J. Hume is with the Electricity Authority, Wellington, New Zealand (e-mail: corresponding author details on submission). ORCID iD: https://orcid.org/0009-0006-6920-0598.

This is descriptive analysis of public regulatory metering and information-disclosure data. The views expressed are the author's. No external funding was received. The analysis informs, but does not advocate, a live reactive-power pricing policy question (Section VI-C).

Reference [1] is the *extended version* of this paper, with the analysis notebooks that compute every number and figure here from public data: the measurement-error analysis, the churn, correspondence and contamination screens, each method's robustness batteries and sensitivities, and the dated regulatory arc. It is cited below at each point its material is used, and not described again.

A. A global reactive transition, already observed

The phenomenon is well attested across jurisdictions: a steep decline in reactive demand at minimum load in Great Britain, with distribution networks beginning to export reactive power to the transmission system [2]; off-peak overvoltage episodes attributed by the French and Queensland transmission operators to distributed generation, underground cabling, and more efficient end-use equipment; rising capacitive export to the extra-high-voltage network in continental Europe [3]; and, at device level, a majority of modern household appliances presenting a leading power factor—the jurisdiction sources are collected in [1]. Expert reviews document the error constant-power-factor load modelling introduces [4], [5]. No consumer-device standard requires a unity fundamental-frequency (displacement) power factor: IEC 61000-3-2 limits harmonic current only [6], and the one device-level power-factor family in force anywhere (the European lighting floors [7]) is a full-load, direction-blind metric that never sees a device's standing draw. The overvoltage-instability literature itself now records the observable: the 2026 review notes, among its causes of voltage rise, that cable substitution has "change[d] the power factor to leading in several urban substations" [8]—an observation offered without citation, and attributed wholly to the network's cables rather than to any change in what the connected demand draws. Every system in this literature runs 230-volt mains; Section V-B returns to that regularity.

The closest measured system-level series is that British one: roughly eight years (about 2005–2013) of reactive demand at grid supply points, projected forward from load-composition modelling—neither a multi-decade measured record nor a quantified demand-versus-network share of the realised drift, and boundary studies elsewhere in Europe likewise characterise or simulate the export without either [3], [1]. The study closest to the present decomposition is Bosisio *et al.* [9] on the Milan distribution network: measured 15-minute boundary data at 37 high-voltage/medium-voltage transformers, naming the same three-term balance used here and, to our knowledge, the only prior study to articulate this split explicitly on measured boundary data. It compares two years across a demand shock, assigns no numerical share, and reads the dominant term as the cables, opposite to ours; we treat it as complementary, and it leaves the national, multi-decade, quantified attribution open [1].

B. The realised risk

In 2025 two power systems suffered major blackouts whose official investigations found a leading-reactive/overvoltage mechanism at their core: the Final Reports of the European Network of Transmission System Operators for Electricity (ENTSO-E), both published 2026, on the Iberian Peninsula event of 28 April 2025 [10] and on the North Macedonia event of 18 May 2025 [11], detailed in Section VI-C. What matters here is what those investigations could not do. Both treat the capacitive surplus as a network-side quantity and low demand as a magnitude rather than as a fleet whose reactive character has itself changed. The Iberian expert panel models the reactive response of load to voltage, but its load model's reactive voltage exponents come from a table of typical values by load class—how the reactive power of demand varies with voltage, not what it was—and the North Macedonian report carries no load power-factor assumption at all; neither states the aggregate reactive character of demand anywhere [1]. The Iberian panel names the gap itself: modelling the transmission–distribution interfaces accurately would need "information on load composition" [10]. These are not criticisms of investigations that reasoned from the data available to them; they identify the space a long measured record can fill. The field's current synthesis has the same shape: the 2026 mechanisms review [8] gives as its causes of voltage rise corridors below surge-impedance loading, cable substitution, synchronous-generator decommissioning, and fixed-power-factor converter operation—network and generation throughout—representing load as lagging or unity power factor in every model, citing no measurement of demand reactive character.

C. Contribution

The contribution is layered. *First, a measured trajectory:* using regulatory half-hourly active and reactive power at New Zealand grid exit points, we measure the system's reactive character at supply-point resolution over 29 years (1997–2025) on a balanced panel—the same grid exit points in every year, so sites entering or leaving the record cannot manufacture the trend—and show the drift is a property of the system rather than of its metering. The computation reuses settlement metering that any operator holds, so the same measurement is available to any system.

Second, a decomposition of the drift into a demand-side load-composition component and a distribution network cable-charging component. The magnitude is carried by a first-principles charging calculation that computes each network's cable charging directly from its disclosed circuit kilometres and so fits no cable coefficient; over the 2013–2025 disclosure window it puts the demand-side (*organic*) share at about 80% (physical band 75–88%), with two statistical methods corroborating the direction but not themselves pinning the share.

Third, the term identified. The deepening is load-independent—as large at the flat evening peak as overnight—so it behaves as a standing shunt element, permanently energised, rather than as a change in how the load draws reactive power while operating; and it lives behind the connections, the 2013–2025 deepening being proportional to the number served with no detectable connection-independent residual, at about 230 VAr per connection (Section V); split by the registry's consumer classes, the residential connections—which hold no power-factor-correction plant—carry the smaller coefficient of the two. We advance the mandated electromagnetic-compatibility filter capacitance as its candidate device layer (a labelled hypothesis) and identify the decisive low-voltage measurement that would settle it.

The significance is what the reversal does to analysis: the direction of the binding voltage-security question reverses; the automatic last-resort defences and the margin tooling of the undervoltage era do not transfer; and the constant-lagging-power-factor load archetypes of planning and protection studies no longer describe the light-load fleet (Section VII)—a qualitative synthesis, not a defence-scheme or load-model redesign.

This work extends the author's EEA 2025 conference paper [12], which described the 2009–2024 aggregate trend qualitatively, by adding the 1997 baseline, grid-exit-point resolution, the artefact test, the decomposition and the connection-scaling identification.

II. Data and Method

A. Sign convention and primary variable

Reactive power carries a direction that is the entire point of this study. We use the convention $Q < 0$ for *leading* (capacitive, voltage-raising) and $Q > 0$ for *lagging* (inductive, voltage-lowering), and lead throughout with the signed reactive power in MVAr, which is additive across supply points and has no discontinuity at unity, and present power factor only as a familiar comparator (load convention; the generator convention reverses both signs). Power factor is non-additive, so where a system power factor is quoted we sum P and Q across the panel first and derive $\mathrm{PF} = \Sigma P / \sqrt{(\Sigma P)^2 + (\Sigma Q)^2}$.

B. The grid-metering record and the balanced panel

The primary data is half-hourly real and reactive power at every grid exit point (GXP), the substations where the national transmission grid hands electricity to the local distribution networks, from 1997 to 2025, published by the New Zealand Electricity Authority on its Electricity Market Information

TABLE I
THE NATIONAL OVERNIGHT TOTALS, AND THE SAMPLE EACH COMES FROM

Sample	Window	Overnight reactive power
Balanced demand panel, 123 GXPs *(the headline)*	1997–2025	$+466 \rightarrow -451$ MVAr
Balanced full panel, 132 GXPs	1997–2025	$+673 \rightarrow -296$ MVAr
Clean and clean-strict cohorts, 119 and 92 GXPs	1997–2025	about -300 MVAr in 2025
Charging footprint, 29 networks	2013–2025	-541 MVAr in 2025, drifting -49.9 MVAr/yr
Connection-scaling panel, 114 GXPs	2013–2025	$+188 \rightarrow -182$ MVAr (medians)

(EMI) platform. Over the record, 238 distinct GXPs appear; a balanced core of 132 is present in all 29 years. The reactive column has been recorded for twenty-nine years because settlement metering records it, not because it was the object of anyone's question: the centralised data set was assembled for demand forecasting and the Grid Investment Test.

We reduce each GXP's year to a small set of features, of which the headline quantities are the average overnight and evening-peak reactive and real power. The overnight window is trading periods 6–10 (approximately 02:30–05:00, the lowest demand) and the evening peak is trading periods 36–38 (approximately 17:30–19:00); we also record the fraction of overnight half-hours that were leading.

The largest way a 29-year trend could be spurious is a change in the set of measured GXPs: new urban, cable-heavy sites appearing and old rural ones dropping out would manufacture a leading trend from bookkeeping alone. The defence is the *balanced panel* of the 132 GXPs present in every one of the 29 years, so any change is a change in the same sites rather than in which sites are measured. The panel is churn-controlled by construction ("balanced" in the panel-data sense, not the three-phase sense), and we follow one rule throughout: *trends on the balanced panel, snapshots on the full panel*. Restricting it to demand networks leaves 123 GXPs; the 132-point and 123-point figures are reported distinctly and kept labelled throughout. The residual bias is signed: the demand grid exit points excluded from the panel are by 2025 on average *more* leading than it and drift faster, so the balanced panel, if anything, understates the national drift [1]. Five national totals follow, each on the sample its method requires, and Table I sets them side by side. They differ by sample, window and summary statistic rather than by measurement; a reader wanting one figure for New Zealand should take the first row.

C. Metering provenance and accuracy

The active and reactive power at each grid exit point is settlement-grade revenue metering in the highest-accuracy category of the regulatory Metering Code [13]: Class 0.2S for active energy under IEC 62053-22 [14], with the reactive channel held to the looser Class 2–3 of IEC 62053-23 [15], and the metered quantity is the fundamental-frequency (*displacement*) reactive power of IEEE Std 1459 [16]. Random error averages down under aggregation and cannot manufacture a near-monotonic 900-MVAr swing through zero; the one systematic error that does not average down—a phase displacement β in the instrument-transformer chain, injecting a spurious term of about $P \sin\beta$ that matters most in exactly the light-load window this paper measures—is bounded, under a common-mode half-degree assumption, at some fifty times smaller than the swing. That bound does not rest on the half degree. A *static* displacement generates apparent drift only through the growth of the load it multiplies, so the most it can produce across the whole record is $\Delta P \sin\beta$; overnight panel load grew by $\Delta P = 280$ MW and $\sin\beta \leq 1$, so a static displacement of any size—90° on every transformer set included—falls three times short of the swing. The stronger defence: it is largest in the loaded window while the deepening is near-equal in both (Section V-A). The full measurement-error analysis is in [1]; we therefore treat the multi-decade *drift* as robust while regarding the standing reactive *level* on any single night as carrying a bounded measurement uncertainty.

D. The asset (parameter) record

To describe each network's physical make-up we assemble a per-company, per-year table from the Commerce Commission's regulated information disclosures: circuit length split into overhead and underground kilometres at each operating voltage (Schedule 9c, disclosed consistently from 2013), capacitor-bank counts (Schedule 9a), demand and connection counts and embedded generation (Schedule 9e), and the regulated asset base (Schedule 4). The disclosures are mapped to the EMI network codes through a 29-year validated correspondence, time-aware because ownership itself changed under the record [1]. Two data-quality choices matter here: only "All"-network roll-up rows are retained, and the capacitor series, counts only and carrying a 2014–2015 definitional break, is used as a direction-only proxy.

E. Confounds in the metered reactive power

A grid exit point meter records the *net* reactive power at the boundary, which is not in general a clean distribution offtake:

$$Q_{\text{metered}} = \underbrace{Q_{\text{load}} + Q_{\text{cable}}}_{\text{the subject of this paper}} - Q_{\text{embedded gen}} \pm Q_{\text{TP plant}},$$

where the last two terms are *confounds*: other reactive sources mixed into the same metered reading. An embedded *synchronous* generator sets its reactive output by the voltage it sees rather than by local load; and at some substations Transpower voltage-regulating plant shares the substation with the offtake, with the revenue-metering boundary determining whether it registers in the metered flow. One behavioural diagnostic catches all of these, whichever party owns the plant: a clean distribution offtake's reactive power tracks its load, so a straight-line fit of Q on P explains an appreciable share of its variation ($R^2 \approx 0.3$–0.5), whereas voltage-regulated reactive is decoupled ($R^2 \to 0$). Computed on three spread-out years (2016, 2019, 2022) so that a flag reflects a stable trait, it isolates six *decoupled*—that is, *contaminated*—grid exit points at Christchurch, Auckland, Bunnythorpe and Tuai [1].

This matters overnight specifically: the minimum removes generation that merely stops, leaving plant that regulates voltage around the clock—synchronous machines, and increasingly inverter-based plant running a night-time reactive mode at zero active power. We flag such sites rather than discard them silently, and Sections III-B and IV show the trajectory and the decomposition robust to their removal: the contamination moves the absolute *level*, not the direction, intensity or attribution of the drift.

F. Analysis methods

The trajectory is summarised by ordinary linear slopes for headline rates and, for robustness, by Theil–Sen slopes [17], which a few anomalous years cannot move; where the two agree—as they do throughout—a trend rests on no small set of years. The metering-artefact test (Section III-B) [1] removes each GXP's own Theil–Sen trend and looks in what remains for a step shared across sites in the same year.

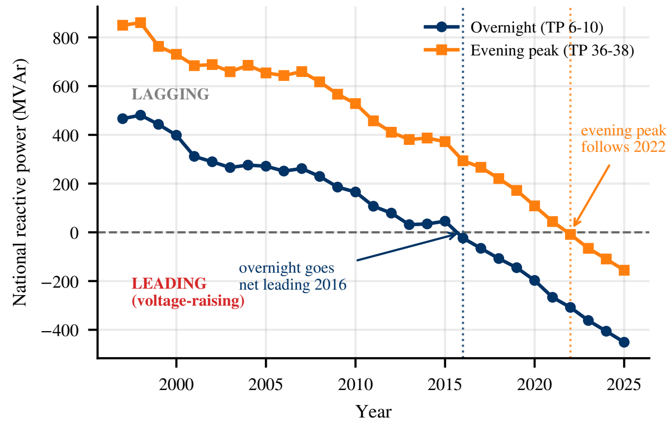

Fig. 1. The 29-year reactive trajectory on the balanced demand panel of 123 grid exit points, overnight against the evening peak ($Q < 0$ is leading).

III. The 29-Year Reactive Trajectory

A. Lagging to leading, overnight

Fig. 1 shows the central empirical result. Summed across the balanced demand panel, mean overnight reactive power fell from $+466$ MVAr (lagging) in 1997 to -451 MVAr (leading) in 2025, a near-monotonic linear drift of -31.1 MVAr/yr (robust Theil–Sen slope -31.3 MVAr/yr, 95% confidence interval (CI) $[-34.0, -28.1]$ [17]—the spread among the pairwise slopes, not a coverage guarantee, since consecutive years are not independent), crossing into net-leading overnight in 2016. On the wider 132-point balanced panel, which also includes non-distribution connection points, the same drift reads $+673 \rightarrow -296$ MVAr at -30.3 MVAr/yr, crossing in 2020. Expressed per megawatt of overnight load, the drift is about -0.018 MVAr/MW/yr. In the comparator of Section II-A, the overnight aggregate power factor moved from 0.964 lagging to 0.975 leading—a small excursion either side of unity concealing a 900-MVAr reversal, which is why signed reactive power is the working variable throughout.

The diurnal contrast is itself a result. The evening peak crossed net-leading only in 2022 and remains far less leading than overnight, which locates the net-leading regime firmly at low load. Expressed as a fraction rather than a sum, the mean share of overnight half-hours that are leading rose from about 1% in 1997 to about 74% in 2025.

B. A measured property, robust to contaminated sites

A measurement artefact cannot reproduce the trend's physical structure. The net-leading crossing is specific to the overnight window—the evening peak stays lagging for years longer (Fig. 1)—while the deepening beneath the crossings is near-equal in the two windows over 2013–2025, with evening-peak load essentially flat (Section V-A): the signature of a standing capacitive term netted against different amounts of lagging load. The drift also tracks each network's underground-cable intensity and grows with each grid exit point's connection count, neither of which a metering error has a reason to do. It does not rest on the older, less homogeneous early data, it is not seasonal, and, since a single system-wide re-base would step every supply point at once, a separate test finds no such event [1]. The lagging-to-leading drift is therefore a measured property of the system, not an artefact of how it was metered.

The trend could also be carried by a few grid exit points whose meters do not see a clean distribution offtake (Section II-E), so we test it by removing them: the six *decoupled* points, leaving the *clean* cohort, and more strictly the steady-generation and net-export points too (*clean-strict*). The absolute total and the MVAr/yr slope *shrink*—the 2025 total from -451 to about -300 MVAr, the slope from -31.1 to -23.4 MVAr/yr—because two removed nodes carry large, load-decoupled leading values. The per-MW drift intensity, by contrast, is essentially unchanged (-0.0175, -0.0178, and -0.0180 MVAr/MW/yr across the full, clean, and clean-strict cohorts), and the share of networks net-leading overnight by 2025 is likewise stable (75%, 74%, 80%). The phenomenon's intensity per unit load and its near-universal reach do not depend on the contaminated nodes; only the headline *size* does.

IV. Decomposition: Demand vs Network

The policy-relevant question is how much of the drift is the network's own cables, which inject leading reactive power as a distributed capacitance, and how much is the changing make-up of demand itself: the retreat of inductive plant and the accumulation of power-electronic equipment (LED lighting, switch-mode supplies, heat pumps and other variable-speed drives). We call the second part *organic*: it is in the load fleet, not the wires. Mostly cables would make it a familiar network-planning problem; mostly the demand fleet makes it structural, nationwide, and outside any single network's control. We approach it three ways; the two statistical methods share a common underground-cable proxy, while the physics-based method fits no cable coefficient and is independent of both.

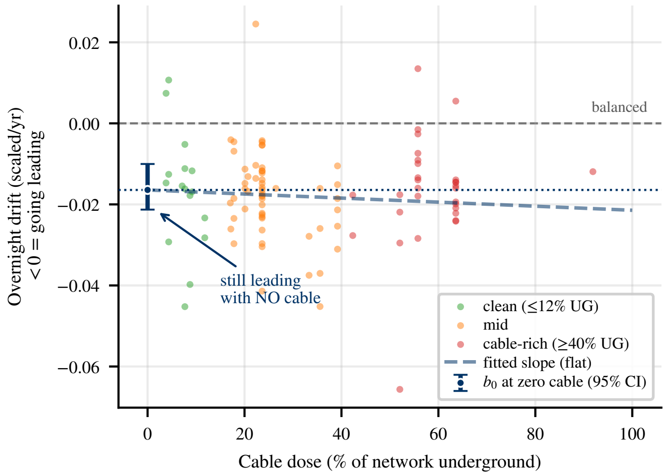

Fig. 2. Method 2, the dose-response: each grid exit point's overnight reactive drift, scaled by its own typical size, against its network's cable dose.

A. Method 1: a clean-cohort contrast

Some distribution businesses are almost entirely overhead line, which injects little leading reactive power, so if their overnight reactive power still drifts leading, that drift can only be the demand fleet. We compare reactive power per MW of load (so network size cancels), load-weighted within each cohort, for networks at or below 12% underground (*clean* in this method's sense) against those at or above 40% (cable-rich).

The clean cohort drifts from $+0.373$ to $+0.065$ MVAr/MW (slope -0.0101/yr), ending near balanced; the cable-rich cohort drifts from $+0.205$ to -0.306 MVAr/MW (slope -0.0134/yr), crossing well into leading. The ratio of slopes puts the demand-side share of the cable-rich cohort's drift at about 75%. The reading rests on two assumptions: that the organic per-MW drift is common across the cohorts (to the extent urban fleets modernise faster, the cable share is overstated), and that the at-most-12%-underground cohort carries no material cable term (to the extent it does, the demand share is overstated); the two pull in opposite directions. The exact ratio also moves with measurement choices, spanning the 64–90% collected in Table II; each variant is worked in [1].

The *direction* is therefore robust while the *exact ratio* is not, which is why we treat Method 1's ratio as a hypothesis and bring two further methods. And although cable is the minority of the drift, it is the marginal factor that tips a network from balanced into net-leading, so cable intensity shapes *where* overvoltage appears first even though demand composition drives the trend.

B. Method 2: a validated dose-response

The second method asks whether the networks with more cable drift faster. Each GXP is reduced to one number—the Theil–Sen slope of its overnight reactive power, scaled by its own typical size—and those per-site drifts are then fitted against the site's *cable dose*, the percentage of its network that runs underground. The line has two coefficients, and only one of them is this paper's result. At zero cable there is no cable charging, so any drift remaining at that end of the line must come from the demand: the intercept b_0 is the organic drift. The slope b_1 is the extra drift per added point of underground—the cable contribution. Fig. 2 shows both.

Two features of the fit follow from how the dose is recorded. Cable dose is a property of the *company*, so every site in a network carries the identical value and the sites are not independent. The model therefore lets each company sit at its own baseline (a random intercept), and the uncertainty comes from a *cluster bootstrap*: the fit is repeated on thousands of datasets drawn at random, with replacement, from the data itself, and the reported interval is the spread of those answers—drawing whole companies rather than single sites, so that the dependence is carried into the interval.

Before any use on the real data the estimator was validated on synthetic panels with planted answers. It recovers b_0 reliably. It recovers the organic *share* only when the cable effect is itself identified, and the reason is arithmetic: that share is $b_0/(b_0+b_1 d)$ at dose d, which puts b_1 in the denominator. On the real data b_1 is not identified, so the share cannot be pinned even though b_0 is—and the magnitude falls to Method 3.

On the real data, the organic drift is $b_0 = -0.016$/yr (95% CI $[-0.021, -0.010]$), leading in 100% of bootstrap resamples. The cable slope is $b_1 = -0.00005$/yr (95% CI $[-0.00037, +0.00012]$), which crosses zero—weakly identified, because cable dose is a company-level trait with only about 23 distinct values—so we deliberately publish no tight interval on the share (at the point estimate it is at least 78%, consistent with 100%, which is an arithmetic consequence of an unpinned denominator rather than a result; Table II). Three further checks hold [1]: adding capacitor intensity as a control moves b_0 about 10%; cable dose and distributed-generation intensity are weakly correlated (-0.22); and restoring the four sites dropped in the mapping moves b_0 under 1%.

A fourth check replaces the dose measure altogether: refitting with each network's computed charging intensity as the dose—Method 3's physical $\omega CV^2\ell$ per MW of maximum demand—moves b_0 only from -0.0164 to -0.0178/yr. The two dose measures are in fact *negatively* correlated across networks ($r = -0.23$), so two nearly opposite dose definitions returning the same organic conclusion is a stronger robustness statement than either alone. No reading here rests on the cable slope under either dose.

C. Method 3: first-principles charging physics

The third method computes the network contribution directly from physics, with no fitted coefficient, so it sidesteps the weak-identification problem entirely. Any energised length of cable or line behaves as a distributed capacitance and injects a fixed leading reactive power $Q = \omega CV^2\ell$, with C the capacitance per km, V the operating voltage, and ℓ the length. Using the published overhead and underground kilometres at each voltage from Schedule 9c (Section II-D), textbook per-km capacitances with an explicit $\pm 35\%$ band carried through to the result, and a band-by-band sum, we compute the physical charging per network per year and subtract it from the measured overnight reactive power; the remainder

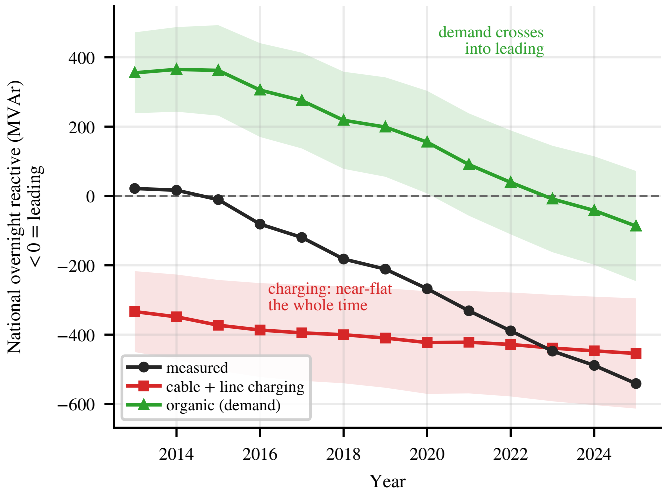

Fig. 3. Measured overnight reactive power across all demand networks (black) against its two components, cable and line charging (red) and the organic demand term (green). Bands are the $\pm35\%$ capacitance range carried through. A different sample from Fig. 1.

is organic. The band voltage is nominal and held across the record; three-quarters of the computed charging sits at or below 33 kV, where the tap changer holds voltage to a target, and a sustained 4% rise across the window would be needed to carry the organic share out of the physical band below [1]. That remainder is a residual, and we name the principal further terms it absorbs: transformer magnetising reactive power, series-reactance absorption, and the distribution capacitor fleet (static-to-shrinking, Section VI-B). Each is either second-order overnight or biases the residual toward lagging, so reading it as the demand-side drift is conservative. The footprint carries no decoupled-node screen, so embedded generation sits in the residual too; re-running the decomposition without those sites moves the organic share *up* rather than down [1].

Over 2013–2025 the measured overnight reactive power drifts at -49.9 MVAr/yr across all demand networks, of which the physical charging contributes only -9.2 MVAr/yr and the organic residual -40.7 MVAr/yr. The organic share of the *drift* is therefore about 80% (physical band 75–88%), and because this method fits no cable coefficient it is independent of the two statistical methods. The window is set by the data: the disclosed circuit kilometres are recorded consistently only from 2013, so extending the $\sim$80% share to the full 29 years is an inference, not a second measurement. This all-network footprint rate is a different sample from the balanced panels of Section III, and the only one of this paper's three drift rates that carries the organic share.

The composition of the leading reactive power *standing on the network today*—the level, as opposed to the drift—is a separate and weaker claim, and on it the physics points the other way: a large, near-static charging baseline has always been present, and what changed is that the organic demand crossed from lagging, which masked it, to leading, which does not. We quote no precise level share, and carry the reading qualitatively into Sections VI and VII.

TABLE II
THREE DECOMPOSITIONS OF THE DRIFT

Method	Establishes	Organic share of drift
1: clean-cohort contrast (cable proxy)	direction (shares the cable proxy with 2)	$\sim$75% (78% with contaminated nodes removed; 64–90% across attribution, cable-intensity, and panel definitions)
2: dose-response (cable proxy)	direction; dominance at typical doses	$\geq$78% at the point estimate, consistent with 100% (cable term not pinned by the data)
3: charging physics (no fitted cable term)	**magnitude**	$\sim$80% (physical band 75–88%)

D. Triangulation

Table II collects the three methods, which play different roles and are not all independent: Methods 1 and 2 both rest on the same underground-cable proxy, whereas Method 3 fits no cable coefficient at all. All three, though, consume the same Schedule 9c circuit-length disclosures, so agreement among them cannot diagnose a systematic kilometre error; under-counted cable would push every method toward the demand side. A cross-sectional check agrees: the most-leading demand grid exit points are not the most-cabled—they sit at roughly half the cable intensity of the most-cabled group—so cable cannot be what drives the deepest-leading behaviour. The three methods share one subject—the *rate* at which the overnight reactive power moved—and Fig. 3 sets it against the standing *level*. The pair carries the practical reading: the distribution network's charging is the larger part of the leading reactive power present tonight, and the demand is the larger part of why that total has moved.

V. The Demand-Side Term Identified: A Standing Shunt Term Behind Connections

The decomposition attributes the drift to the demand fleet; it does not say what in the fleet is doing it. Two measured properties narrow the answer, a specific physical candidate follows from them, and a scaling test probes the candidate's sharpest prediction.

A. Load-independent, around the clock

First, the drift is load-independent. Over the 2013–2025 window we summarise each grid exit point passing the full contamination screen by robust annual medians on wider windows than the Section II features. On the 114 such grid exit points with both endpoints metered, the median per-site change in reactive power is -2.25 MVAr overnight and -2.42 MVAr at the evening peak (median per-site ratio 1.06), while evening-peak load was essentially flat (median change $+2.9\%$). A change in how the connected load draws reactive power while operating would scale with the load operating—far larger at the peak than at 03:00—and so would the load-proportional metering bias bounded in Section II-C. A near-equal deepening around the clock is instead the signature of a *standing shunt term*: always energised, indifferent to what the load is doing. It also reconciles the diurnal structure of Fig. 1, since net-leading

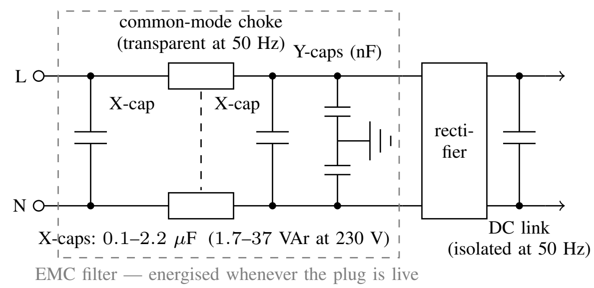

Fig. 4. The mains input stage of a switched-mode device. Values are datasheet-grade typical ranges, per device rather than per capacitor.

surfaces overnight first because the standing component is netted against more lagging load at the peak.

What the test identifies is the *net* standing term. A standing shunt inductance is load-independent and scales with V^2 exactly as a capacitance does, so the record returns the difference of two limbs and never either one—and both moved across the window. Filter-carrying electronics accumulated; the permanently energised magnetics retreated, in linear transformer supplies, overnight magnetic ballasts, always-on shaded-pole motors, and the distribution transformer's own magnetising reactive power, itself standing, connection-scaling and inductive [1]. A retirement moves the net exactly as an addition does. We report the net term throughout, and read the device layer below as a candidate for its capacitive limb alone.

B. A physical candidate: the mandated EMC filter

Second, the term must live behind essentially every connection, because the drift's reach is near-universal across grid exit points. We therefore advance a specific physical candidate, labelled plainly as a hypothesis: the electromagnetic-compatibility (EMC) filter that practically every mains-connected electronic product is required to carry. A switched-mode device chops current at tens to hundreds of kilohertz, and conducted-emission limits (the international CISPR 32 limits [18], a condition of sale in most jurisdictions, New Zealand included) oblige its manufacturer to stop that noise at the appliance's own terminals; the universal compliance solution is a passive filter directly across the mains input (Fig. 4). Its elements carry safety-class names from where they connect: *X-capacitors* bridge line and neutral, the one position where a capacitor is permitted to be large (typically 0.1–2.2 μF), because a short circuit there is an overcurrent event the fuse clears rather than a shock hazard; *Y-capacitors* run to earth and are held to nanofarads by touch-current limits.

At 50 Hz the filter's only material injection is the X-capacitance, which sits across 230 V drawing pure leading current whenever the plug is live—ahead of the rectifier and of any standby logic. The arithmetic is $Q = \omega C V^2$, and the fleet is one worldwide design, so the capacitance is chosen once and the mains it lands on sets the draw: the same microfarad that injects 16.6 VAr at 230 V injects 5.4 VAr on North America's 120 V, 60 Hz mains and 3.1–3.8 VAr at Japan's 100 V—roughly three times the standing draw per device here as in North America, and four to five times Japan's. The accumulation is one-way: one instrument drives the capacitance into practically every product, while none bounds its fundamental-frequency injection at any scale, as far as we are aware [1].

The hypothesis is narrow: the large DC-link capacitor behind the rectifier makes no standing injection (Fig. 4); motor-run capacitors sit in series with the motor's auxiliary winding and are disconnected with it; and rooftop solar inverters contribute almost none overnight, the isolation relay stranding the converter-side filter capacitors—a bound rather than an assumption [1].

A bottom-up order-of-magnitude estimate—roughly two million households, each with 20–40 permanently connected filter-carrying devices at 2–4 VAr (roughly 80–320 MVAr in aggregate), plus a commercial and industrial fleet of the same order—puts the national standing total at roughly 150–600 MVAr, or 100–200 VAr per household as the central band. Heat pumps, an adoption jump inside the study window, work out at roughly 37–69 MVAr [1], of order a quarter of the residential envelope, so the bulk belongs to the far more numerous small supplies. The per-device values are datasheet-grade reference designs rather than teardowns, and the device counts are estimate arithmetic anchored by intrusive standby surveys rather than a census [1]. The estimate makes two testable predictions: the national scale of the drift, and—because the fleet lives behind consumer connections—that the drift should scale with the number of connections served, with nothing left at zero connections.

The first prediction is met at the order-of-magnitude grade the estimate carries. Summed across the same 114 grid exit points, the median overnight reactive power moved from +188 MVAr (net absorbing) in 2013 to −182 MVAr in 2025. Those exit points serve about 1.5 million connections—two-thirds of the national connection count—so the 369 MVAr rise normalises to about 210–280 VAr per connection [1]—above the estimate's 100–200 VAr per household centre rather than inside it, because a connection is not a household: Section V-C separates the consumer classes and puts the residential rate back inside the band.

The dose arithmetic adds a third, cross-country prediction: if the filter is the mechanism, the per-connection drift should order across systems as ωV^2—Japan, then North America, then the 230-volt world—an ordering cables, heat pumps and correction plant have no reason to reproduce. It is a test and not a conclusion: a 120-volt supply draws twice the current for the same power and may need *more* X-capacitance, so the dose ordering holds exactly only for the universal-input supply designed once for every market [1]. Every system in Section I-A runs 230-volt mains; we are not aware of an equivalent North American measured report, an absence of measurement rather than evidence of absence. Of the two domestic predictions the second is the more discriminating; Section V-C tests it.

C. A connection-scaling test

Connection counts are public: the retail market-structure dataset records the number of installation control points

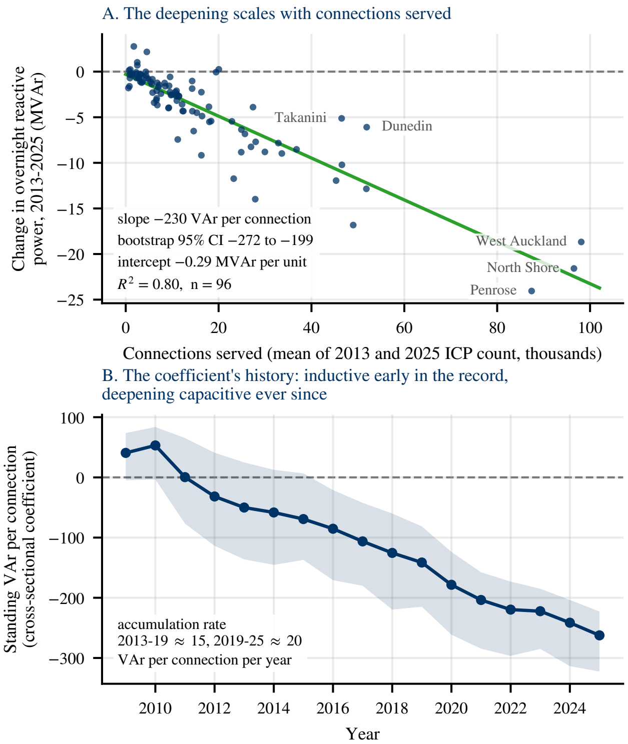

Fig. 5. The connection-scaling test: (A) each panel unit's 2013–2025 change in overnight reactive power against the connections it serves; (B) the same coefficient refitted on each year's standing level.

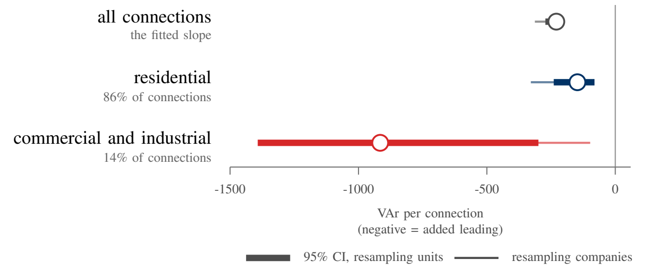

Fig. 6. The 2013–2025 connection-scaling coefficient split by registry consumer class.

(ICPs—metered connections) behind each network supply point, monthly, from December 2003 [19]. We sum them to each grid exit point and regress each panel unit's 2013–2025 change in median overnight reactive power on the connections it serves, on the $n = 96$ panel units a correspondence audit leaves, grid exit points grouped where the registry's supply-point attribution and the metering boundaries disagree. Section V-B's estimate sets the test in advance: a slope clear of zero, inside a per-connection band around the bottom-up figure, and explaining a substantial share of the cross-site variance. A grid-exit-point-level metering artefact predicts a mean-only model instead. The criteria, the audit, and the site dispositions are in [1].

The result is Fig. 5: the deepening scales linearly with connections served, at −230 VAr per connection, and a through-origin fit agrees (−237). All three criteria are met, the mean-only artefact model loses decisively on the standard model-selection score, and the intercept is indistinguishable from zero under every inference scheme—so the bottom-up estimate's prediction of nothing at zero connections is met on both limbs. Connections also beat size, the null that matters here since both scale with it: fitted jointly against each unit's overnight or peak real power, the connection term keeps an interval clear of zero and the load term does not, and dropping the five highest-leverage units *raises* the coefficient to −261 [1].

The coefficient also has a history (Fig. 5B). Refitted year by year on each year's standing level, so that it carries the whole connection-scaling class, the per-connection term starts on the *inductive* side of the axis (+41 VAr per connection in 2009), crosses zero early in the 2010s, and reaches −262 VAr per connection by 2025. The yearly intervals leave the ramp and the post-2016 deepening as what the data establish, rather than the crossing date. The 2013–2025 movement, −212 VAr per connection, reproduces the two-endpoint slope within 10%.

Read honestly, the slope is the whole connection-scaling shunt class, not the device fleet alone: each connection also brings its share of medium-voltage cable—a standing 190 VAr on Section IV-C's physics, low-voltage cable being negligible since charging scales with V^2—and commercial premises hold the one rival standing capacitance at scale—power-factor-correction banks left in service overnight wherever the control is fixed or coarsely stepped. The registry's industry classification separates these populations, and each takes its own coefficient (Fig. 6). The residential coefficient is −147 VAr per connection, measured on the 86% of connections that hold no correction plant at all, where a device fleet is the only standing capacitance a connection brings. That is a twelve-year *change*, not the standing *level* the bottom-up estimate predicts; the level, from the same split, is −310 VAr per residential connection in 2025, *above* the 100–200 VAr band because it carries that cable share as well as its devices. The commercial-industrial coefficient is −915, larger per connection as commercial device fleets and shared cable predict, and the two classes carry about half the panel's deepening each on the point estimates, despite the six-to-one difference in their numbers. That commercial-industrial term stays composite: commercial device fleets, any stranded correction, and shared medium-voltage cable all sit inside it, and because the two counts are collinear across units the coefficients are bounds rather than a sharp allocation. Two marks bear on the correction reading without settling it: the population able to hold a bank is 1.7% of recent connection growth, so the stock that would have to be growing is small; and the commercial-industrial class remains net lagging overnight in every registry vintage—falling, and with an interval spanning zero—so its standing position has not turned capacitive. Either way the term stays demand-side and connection-scaling: correction plant behind a connection is not network cable, so its presence would qualify this section's device-layer attribution rather than the decomposition of Section IV.

The $P \sin\beta$ error class of Section II-C would mimic the scaling, but is bounded far below the measured deepening and would concentrate it in the loaded window rather than equally in both.

One rival survives in principle: a load-independent metering-registration drift phased in with meter replacement. It has no mechanism to produce the connection-scaling term, and the connection-independent residual it could sit in is indistinguishable from zero, so it is heavily disfavoured while remaining formally excludable only by a measurement that bypasses the grid-exit metering chain—night-time reactive flow at a residential distribution transformer with a known house count. That is the stated next step.

VI. From Leading Reactive Power to Overvoltage

A. Why leading reactive power raises voltage

The voltage change along a line is $\Delta V \approx (RP + XQ)/V$, with the XQ term dominant on the grid: lagging reactive power ($Q > 0$) drops voltage toward the load; leading ($Q < 0$, the new overnight reality) pushes it up. Light load is the dangerous condition—demand is lowest, so little load absorbs the reactive power, while the cable charging is always on—and the mechanism can run away: charging rises with the square of voltage, and a protection trip leaves the remaining network carrying more charging per unit of load. Two distinct feedback loops run through that description, usually described as one. In the *demand-side* loop, disconnected load takes with it whatever absorption it was providing, so whether shedding raises voltage or lowers it turns on the power factor of what comes off. In the *plant-side* loop, rising voltage disconnects plant that was absorbing and voltage rises further; it runs on protection's response to voltage, behind a demand of any power factor. In isolated and autonomous systems the feedback is well documented [20]; in neither 2025 blackout did the lost absorption leave with the demand (Section VI-C).

B. The system is already responding

The New Zealand system is visibly responding from both ends. Distribution capacitors are coming out: the national bank count fell from 298 in 2015 to 254 in 2025 (counts, not MVAr)—a static-to-shrinking capacitor fleet cannot be causing a rising leading trend; a count cannot see switching practice [1]. Transmission absorption is going in: shunt reactors and STATCOMs installed in response to operator security forecasts that record overnight high voltages being managed by switching transmission circuits out of service [21], the follow-up study then finding no significant overvoltage issues for current and future light-load conditions once the reactors and a STATCOM were in service [22]. And the operator has named the demand-side term: Transpower's 2023 transmission planning report expects the high-voltage issue to worsen and lists "load characteristics becoming generally more capacitive" among the factors [23]—the asset owner attributing part of the standing pressure to the electrical character of the load, the quantity this paper measures.

C. The realised risk, abroad

The Iberian event of 28 April 2025 was found to be a voltage and reactive collapse—the Final Report identifies "non-effectiveness of voltage control" as a root cause, and the absorption the cascade removed left with the generation, tripped on overvoltage protection, not with the demand [10]. The North Macedonia event of 18 May 2025 was a cascade of transformer trips from sustained overnight overvoltage, driven by the high capacitive reactive power of light-load networks with no compensation plant and no load shed at all, so it ran with the demand-side loop never operating [11]. No investigated event attributes a blackout to the demand-side loop. We read them as corroboration that the leading-reactive regime is a realised blackout mechanism, New Zealand's distinctive contribution being the 29-year measurement of how a system drifts into it.

The provisions written against the lagging era do not reach this window, and neither cover nor price the light-load surplus; whether that should change is outside this paper's descriptive scope [1].

VII. Implications for Power-System Analysis

The measured reversal relocates the binding voltage-security question. For the lagging era's failure mode—voltage collapse under heavy load—power-system practice holds a mature, layered stack: automatic under-voltage load shedding, tap-changer blocking, PV- and QV-margin analysis, all built on load models of an inductive fleet. In the light-load regime—overnight here and in North Macedonia, midday where distributed photovoltaics hollow out net transmission demand—the equivalent layers are thinner: the *automatic* last-resort layer acts on voltage falling, not rising, and the margin tooling reports distance to collapse, not to the upper limit.

The *drift* into the regime is dominantly the demand fleet, but the standing light-load surplus the mechanism feeds on is substantially long-standing distribution network charging the fleet's shift has uncovered (Section IV)—distributed, always energised, and not itself dispatchable short of spending N-1 security on voltage margin. And the cascade arithmetic is inverted relative to collapse: in a collapse each element lost reduces the demand depressing voltage, so the disturbance partly self-relieves, whereas in the plant-side loop each protection operation removes absorption and strengthens the driver, minutes in Iberia, hours in North Macedonia [1].

The analysis stack under-represents the regime it is entering. Planning and protection studies still commonly represent demand with constant, lagging power-factor archetypes calibrated to the motor-dominated fleet [4], [5]; the measured net overnight flow now carries the opposite sign. That reaches past the archetype: since the demand-side loop's sign turns on the power factor of what comes off (Section VI-A), a shedding or load-rejection study calibrated on lagging demand carries an unexamined sign at light load. The gap is calibration rather than structure: the standard composite load model does carry a distribution-side shunt term, but back-solves it from the power-flow Q [24], so at light load it reproduces whatever archetype the case was built on. Shedding whole feeders takes that capacitance with them, but shedding is a late event; at initiation the demand contributes a standing level, not a trip. Overvoltage protection, automatic shunt switching and the QV curve's upper limb each address part of the question, and the 2026 mechanisms review proposes containment countermeasures—but none is a routine margin quantity [8].

The measurement gives that correction a concrete, portable form: a net standing shunt term, capacitive by about 230 VAr per connection over 2013–2025 (Section V-C), alongside whatever power-factor archetype represents the operating draw. It arrives behind existing connections rather than with demand growth—which alone would leave Fig. 5B's coefficient flat—and archetypes calibrated even a decade ago predate it entirely.

VIII. Conclusion

The defences, margins, and load models of power-system practice were built for lagging load and the undervoltage direction; the measured overnight flow has crossed to the other side, where no last-resort scheme is addressed to voltage rising and, in the plant-side loop, each protection operation reinforces the disturbance (Section VII). New Zealand's record—an advanced-state, small, isolated, heavily cabled analogue whose net-leading crossing preceded the 2025 European events by nearly a decade—is, to our knowledge, the longest direct measurement of that crossing. If the standing term is carried by the globally standardised product fleet, every system connected to the same supply chain is accumulating it on a similar schedule, at a dose set by its supply voltage, and that dose ordering is itself the decisive comparative test. Supply voltage sets the dose; transmission voltage sets its visibility. The term is created at the mains, indifferent to the network above it, while in any system-wide reactive account the boundary flow that would reveal it is dwarfed by the transmission network's own charging, which grows with the square of transmission voltage—so New Zealand's 220-kV maximum is an unusually transparent window, and in a 400-kV system that contrast is several times sharper, where a network-side reading is the natural one (the reading both 2025 investigations took). The claim is therefore priority of measurement, not of exposure: the failures arrived first where the network term is largest, not where the demand-side term is clearest.

The drift shows no sign of slowing: the per-connection coefficient was running at about 20 VAr per connection per year at the 2025 end, and the next device waves—an electric-vehicle charging and home-battery fleet expected to carry the largest input filters yet—are still arriving. The prediction also runs the other way: unlike the passive filter fleet, that arriving wave is controllable, and interconnection standards already define the voltage–reactive-power response modes to steer it (IEEE Std 1547) [25], so large inverters required to absorb at light load are the natural first counter to the leading surplus. Tracking the coefficient with a renewal model of device generations, and the decisive attribution measurement at a residential distribution transformer, are the nearest work [1].

The limitations are stated throughout: the balanced panel controls churn but does not eliminate it; the decomposition attributes the drift, not a causal mechanism in any single load; the exact organic share and the level composition are not point-identified; the measurement gives the net standing term and not its inductive and capacitive limbs separately, so the device-level attribution remains a labelled hypothesis for the capacitive limb, with its decisive low-voltage measurement identified; and a single national system is one observation, offset not cured by its position as a leading-edge analogue. What this paper offers is the measured trajectory, its decomposition, and the term it identifies.

Reproducibility

Every number and figure is computed from public data (Electricity Authority grid metering and Commerce Commission information disclosures) in an accompanying set of analysis notebooks. Both input sets are public under CC BY 4.0—grid metering on the Electricity Authority's Electricity Market Information platform, network parameters in the Commerce Commission's information disclosures—and ship with the archived package [26].

Acknowledgment

The author thanks colleagues at the Electricity Authority and Transpower for discussion of the reactive-power and voltage-management context, and the authors of the Rhodes studies for sharing corrected source material.

The author used Anthropic's Claude models as a research assistant throughout this work: literature searching, analysis and replication code, manuscript drafting and editing, and verification passes over numbers and citations. The research questions, analytical design, interpretation, and conclusions are the author's own. Every printed number is reproduced from public data by the accompanying notebooks and every citation was checked against its primary source; the author is responsible for all content.

References

[1] D. J. Hume, "From lagging to leading: the measured emergence of standing capacitance behind consumer connections, 1997–2025 (extended version)," Zenodo, 2026, doi: 10.5281/zenodo.22238685.

[2] C. G. Kaloudas, L. F. Ochoa, B. Marshall, S. Majithia, and I. Fletcher, "Assessing the future trends of reactive power demand of distribution networks," *IEEE Transactions on Power Systems*, vol. 32, no. 6, pp. 4278–4288, Nov. 2017.

[3] M. Tazky, M. Regula, and A. Otcenasova, "Impact of changes in a distribution network nature on the capacitive reactive power flow into the transmission network in Slovakia," *Energies*, vol. 14, no. 17, p. 5321, 2021.

[4] CIGRE Joint Working Group C4.24/CIRED, "Power quality and EMC issues with future electricity networks," Technical Brochure 719, 2018.

[5] Electric Power Research Institute, "Enhanced load modeling: Improved reactive power load modeling," Palo Alto, CA, Technical Update 3002022354, Aug. 2021.

[6] "IEC 61000-3-2: Limits for harmonic current emissions (equipment input current $\leq$16 A per phase)," IEC, Tech. Rep., 2024.

[7] European Commission, "Commission regulation (EU) 2019/2020: ecodesign requirements for light sources and separate control gears," Official Journal of the European Union, L 315, Tech. Rep., 2019.

[8] C. Vournas and T. Van Cutsem, "A review of overvoltage instability mechanisms," *Sustainable Energy, Grids and Networks*, vol. 47, p. 102475, 2026.

[9] A. Bosisio, F. Soldan, A. Morotti, G. Iannarelli, E. Bionda, and S. Grillo, "Lessons learned from Milan electric power distribution networks data analysis during COVID-19 pandemic," *Sustainable Energy, Grids and Networks*, vol. 31, p. 100755, 2022.

[10] ENTSO-E ICS Investigation Expert Panel, "Grid incident in Spain and Portugal on 28 April 2025 — final report," Brussels, Belgium, Tech. Rep., Mar. 2026.

[11] ——, "Grid incident in North Macedonia on 18 May 2025 — final report," Brussels, Belgium, Tech. Rep., Apr. 2026.

[12] D. Hume, "Power factors in the New Zealand power system," in *Electricity Engineers' Association (EEA) Conference*, Christchurch, New Zealand, Sep. 2025.

[13] Electricity Authority, "Electricity industry participation code 2010, part 10: Metering," Tech. Rep., 2025, cl 10.37(2)(a) and Schedule 10.1, Tables 1, 5, 6; consolidation as at 1 April 2025.

[14] "IEC 62053-22: Static meters for AC active energy (classes 0,1S, 0,2S and 0,5S)," IEC, Tech. Rep., 2020.

[15] "IEC 62053-23: Static meters for reactive energy (classes 2 and 3)," IEC, Tech. Rep., 2020.

[16] "IEEE Std 1459-2010: Definitions for the measurement of electric power quantities," IEEE, Tech. Rep., 2010.

[17] P. K. Sen, "Estimates of the regression coefficient based on Kendall's tau," *Journal of the American Statistical Association*, vol. 63, no. 324, pp. 1379–1389, 1968.

[18] "CISPR 32: Electromagnetic compatibility of multimedia equipment—emission requirements," IEC, Tech. Rep., 2019.

[19] Electricity Authority, "Market share trends by root network supply point (NSP)," Electricity Market Information (EMI) market-structure dataset, 2026. [Online]. Available: https://www.emi.ea.govt.nz/Retail/Datasets/MarketStructure

[20] C. D. Vournas, V. C. Nikolaidis, and G. I. Tsourakis, "Coordinated countermeasures against overvoltage instability in autonomous power systems," *IEEE Transactions on Power Delivery*, vol. 36, no. 6, pp. 3329–3338, Dec. 2021.

[21] Transpower New Zealand (System Operator), "System security forecast 2022, part D: Power system security analysis — grid zone 2 (Auckland)," Tech. Rep., Dec. 2022.

[22] ——, "N-1 thermal and voltage study — system security forecast 2024," Tech. Rep., Jul. 2024.

[23] Transpower New Zealand, "Transmission planning report 2023," Tech. Rep., 2023.

[24] WECC Modeling and Validation Subcommittee, "WECC composite load model specification," Tech. Rep., 2021.

[25] "IEEE Std 1547-2018: Interconnection and interoperability of distributed energy resources," IEEE, Tech. Rep., 2018.

[26] D. J. Hume, "From lagging to leading — data and code archive," Zenodo, 2026, doi: 10.5281/zenodo.21927098.