COMPOUNDING ENERGY LTD
*Working Paper Series • CE-WP-2026-06*

# Load Growth and the Un-cannibalisation Ledger

*What Data-Centre Demand Does to Renewable Capture Rates — and to the Battery Fleet*

**Dr Christopher T. M. Clack**
*Founder, Compounding Energy Ltd*
christopher@compoundingenergy.com

September 2026

**Abstract.** Gigawatt-scale data-centre load is arriving on a system whose wind and solar fleets are already cannibalising their own capture rates. This paper prices the load-growth side of that arrival, the complement to the flexibility valuation of Clack (2026b), on the same stylised Great Britain equilibrium: data-centre load is added to a 2030 background at 2–15 GW connected, a stated share of it responds flexibly, and storage dispatch, prices and the flexible response are solved jointly. Two pre-registered arms bracket the answer, the second registered after the first tripped its own scarcity band. When the background is planned to accommodate the load, inflexible data-centre demand un-cannibalises renewables: the solar capture ratio rises from 0.36 to 0.58 and negative-price half-hours fall from about 1,800 to under 100 per year at 10 GW, worth £1.7 billion a year to the renewables fleet, while the same load removes the surplus hours the battery fleet charges in and cuts its wholesale-plus-response revenue by 26%; making half the load flexible deepens the battery loss to 38%. Data-centre load is therefore the renewables fleet's customer and the battery fleet's competitor at every penetration and every flexible share we ran — the ledger has opposite signs everywhere. When the load is not accommodated, scarcity dominates: prices triple at 10 GW, every asset class gains, and the pre-registered band that flags a stylised stack driving the answer trips, as designed. The sign structure, not the levels, is the result: a second accommodation device roughly halves the ledger levels and preserves every sign and ordering. All results are reproducible from a stated (scenario, code version, timestamp) triple; the case-study bundle is available on request.

**Keywords:** data-centre load growth, renewable capture rates, cannibalisation, battery storage, equilibrium, Great Britain

# Contents

# 1 Introduction

Two conversations about data centres and the grid are running in parallel and rarely meet. One is about *flexibility*: whether large computing loads can shift or shed in response to grid conditions, and what that responsiveness is worth. We priced it at equilibrium in Clack (2026b) and found that the value erodes with its own uptake. The other is about *load growth*: whether the system can absorb gigawatts of new demand at all. This paper prices the second conversation on the same machinery, and asks the question the two conversations share but neither has answered: who, among the assets already on the system, does data-centre load help, and whom does it hurt?

The candidates are the renewables fleet and the battery fleet. Wind and solar in Great Britain increasingly generate into hours of their own making — surplus half-hours where the price sits at or below zero — and their capture ratio, the generation-weighted price relative to the time-average price, falls as their capacity grows. That is cannibalisation in its original sense (Hirth, 2013; López Prol et al., 2020). A large, near-flat new load consumes in those surplus hours as much as in any other, so it should lift the price floor renewables sell into and *un*-cannibalise them. But those same surplus hours are precisely the hours in which batteries charge; the spread between them and the evening peak is the storage business. A load that fills the troughs raises the battery's buy price. And if the load is flexible, it moves its own consumption into the troughs deliberately, competing with the battery for the same spread — the cross-cannibalisation of Clack (2026b), now in a background where the load itself is growing.

Whether the answer is "customer" or "competitor" is not a matter of opinion; it is a matter of price formation, and it depends on one planning assumption that is usually left implicit: whether the system has been built to accommodate the load. We make that assumption explicit and run both cases.

# 2 Framework

The machinery is that of Clack (2026b), itself the equilibrium framework of Clack (2026a) with flexible-demand cohorts: a stylised single-zone Great Britain with synthetic demand and renewables, a strictly convex inverse supply representation, a reference storage fleet of 1 h, 2 h and 4 h cohorts (3, 5 and 2 GW in the 2030 background), an aggregate frequency-response product, and a flexible-demand cohort with deferral, curtailment, backlog, payback-window and connection-headroom constraints anchored to public demonstrations. Storage dispatch, the flexible response and half-hourly prices are the joint minimiser of one convex welfare program; prices are the duals of power balance served through the supply representation. Every solution is certified optimal by the solver and passes the primal verification suite at solver-consistent tolerances recorded in the run manifest; nothing in the program has changed.

Three switches turn the flexibility case study into a load-growth case study, and the paper's case-study code carries them as such. First, the data-centre load $L_G(t)$ — a near-flat profile at load factor 0.90 on $G$ GW of connected capacity — is *added* to the background residual rather than netted at its mean: in Clack (2026b) the background was taken to embed the projected data-centre demand so that only the load's shape and its flexible response entered residual demand; here the load's level is the question. Second, a flexible share $s \in \{0, 0.25, 0.50\}$ of the connected load responds through the flexible cohort of Clack (2026b) at its central service-level parameters; the remainder is inflexible. Third, wind and solar are carried separately so that each fleet's capture ratio can be read at every equilibrium.

**Two arms.** The *accommodated* arm is the central case: the background is planned for the load it hosts, which we represent by moving the supply representation's mid-merit reference and scarcity knee out by the added mean load $\mathrm{LF} \cdot G$ and leaving its oversupply term unchanged. The device is not a pure translation: the knee moves, but the mid-merit segment beneath it is rescaled rather than shifted, so prices in the former surplus hours rise by more than a translation would give and the peak segment cheapens as load is accommodated (the 95th-percentile price falls from £182 to £169/MWh between 0 and 15 GW). A pure-translation device is run as a post-hoc sensitivity (Section 5). The *unaccommodated* arm adds the load to a fixed supply stack. The unaccommodated arm was run first; its pre-registered band on the mean price tripped, and the accommodated arm was pre-registered and run in response (Section 3). Both arms are reported.

**Quantities.** For each equilibrium we read the wind and solar capture ratios, the number of negative-price half-hours, the mean price, the renewables fleet's annual revenue, the storage fleet's wholesale-plus-response revenue per megawatt (fleet-weighted and by duration), and the flexible cohort's value per kilowatt of flexible capacity. The *ledger* differences the renewables' revenue and the storage fleet's revenue against the same weather year at $G=0$, in £m per year and per gigawatt of connected load. Levels are stylised and badged as such; the signs, orderings and shapes are the results.

## 3 Pre-registration and outturn

The hypotheses, bands and run matrix were written before any production cell was solved, and the addendum registering the accommodated arm was written after the unaccommodated arm's outturn and before any accommodated cell (`PREREGISTRATION.md` in the bundle). We report them as they landed.

*Unaccommodated arm.* The mean-price band — the $G=15$ mean below twice the $G=0$ mean, else scarcity from a fixed stylised stack is driving the answer — tripped at 5.9× (Table 4). In that regime the hypothesis that capture ratios rise monotonically in $G$ was refuted (the solar and combined ratios rise to 5 GW and the wind ratio only to 2 GW, and all then fall as scarcity lifts the time-average price faster than renewables' weighted price), storage revenue rises with $G$ as predicted, and every asset class gains against $G=0$, so the ledger's signs do not oppose. We retain the arm as the scarcity bound and draw no capture-rate conclusion from it beyond 5 GW. A second band, solar capture at $G=0$ between 0.5 and 1.0, was missed at 0.36 in both arms: it is structural to the stylised 2030 background, which carries about 1,800 negative-price half-hours a year, and the $G=0$ cells are the same equilibrium in both arms; no new band was set. A third band, that the $G=0$ storage baseline reproduces the Clack (2026b) inflexible baseline at the same seeds (its 0.25 GW netted cell), passed: fleet-weighted £39.8k/MW/yr in both, within 0.02% per duration and weather year (the check, `band_wp05_baseline.json`, ships in the bundle).

*Accommodated arm.* All five registered hypotheses were confirmed (the evaluation, `hypotheses_acc.json`, ships in the bundle): capture ratios rise monotonically in $G$ at every flexible share and negative-price half-hours fall; the mean price stays within ±25% of its $G=0$ value at every $G$ (it rises 13% at 10 GW and 16% at 15 GW); the storage fleet's revenue *falls* with accommodated load; a larger flexible share deepens both effects; and the ledger has opposite signs at every $(G, s)$.

## 4 Results

*[Every value in this section is generated by the run matrix of Appendix A at seed 20260901 and passes the physical verification suite; tables are generated from the results files; levels are stylised. Means over three synthetic weather years.]*

### 4.1 Un-cannibalisation

Table 1 is the renewables' side of the ledger. Accommodated inflexible load fills the surplus hours: at 10 GW the solar capture ratio rises from 0.36 to 0.58 and the wind capture ratio from 0.79 to

Table 1: De-cannibalisation under accommodated load growth: renewables capture ratios (generation-weighted price over time-average price), negative-price half-hours and mean price against connected data-centre load $G$, inflexible ($s$=0); 2030 background; means over three synthetic weather years; stylised levels.

| | $G$=0 | $G$=2 | $G$=5 | $G$=10 | $G$=15 |
|---|---|---|---|---|---|
| Wind capture ratio | 0.79 | 0.81 | 0.83 | 0.85 | 0.87 |
| Solar capture ratio | 0.36 | 0.43 | 0.50 | 0.58 | 0.63 |
| Negative-price half-hours per year | 1,794 | 1,370 | 636 | 84 | 9 |
| Mean price (£/MWh) | 65.3 | 67.9 | 70.9 | 74.1 | 76.1 |
| Renewables revenue vs $G$=0 (£m/yr) | – | +493 | +1,082 | +1,731 | +2,123 |

Table 2: Storage fleet wholesale-plus-response revenue under accommodated load growth, £k per MW per year, by connected data-centre load $G$ and flexible share $s$ (fleet-weighted over the 1 h/2 h/4 h reference build-out); the last block gives the $s$=0 split by duration. Baseline ($G$=0) in the first column; changes are relative to it.

| | $G$=0 | $G$=2 | $G$=5 | $G$=10 | $G$=15 |
|---|---|---|---|---|---|
| Fleet, $s$=0 | 39.8 | 37.1 | 33.6 | 29.5 | 27.1 |
| Fleet, $s$=0.25 | 39.8 | 36.5 | 32.2 | 27.0 | 23.6 |
| Fleet, $s$=0.5 | 39.8 | 35.9 | 30.9 | 24.7 | 21.3 |
| Change vs $s$=0 at same $G$, $s$=0.25 | – | -0.6 | -1.4 | -2.5 | -3.5 |
| Change vs $s$=0 at same $G$, $s$=0.5 | – | -1.2 | -2.7 | -4.8 | -5.7 |
| 1 h cohort, $s$=0 | 21.0 | 19.7 | 18.0 | 16.2 | 14.7 |
| 2 h cohort, $s$=0 | 37.4 | 34.8 | 31.5 | 27.3 | 25.2 |
| 4 h cohort, $s$=0 | 73.9 | 68.8 | 62.3 | 55.1 | 50.1 |

0.85, and negative-price half-hours fall from about 1,800 a year to 84. The renewables fleet earns £1.7 billion a year more than at $G$=0 — £246 million per gigawatt over the first two gigawatts and £142 million per gigawatt averaged over fifteen, the marginal gigawatt between 10 and 15 GW being worth about £78 million, as the surplus hours are used up. This is un-cannibalisation: the load pays renewables for the energy the system would otherwise have priced at the floor. The flexible share adds little to it (Table 3); the level of the load does the work, and flexibility moves a small further amount of consumption into the troughs.

## 4.2 The battery fleet

Table 2 is the other side. The same accommodated load cuts the storage fleet's wholesale-plus-response revenue by 26% at 10 GW when none of it is flexible, because the troughs the fleet charges in have been filled and, under the device, the peak it discharges into has cheapened, so the spread is compressed from both ends; the four-hour cohort loses the most per megawatt (£73.9k to £55.1k), as in Clack (2026b). When half the load is flexible the loss deepens to 38% at 10 GW and 46% at 15 GW: the flexible half is a demand-side battery competing for what remains of the spread, and its own value per kilowatt of flexible capacity erodes with penetration in the same way (£10.3 to £7.3/kW/yr from 2 to 15 GW at $s$=0.25). Read against the measured record, the levels do not transfer — the stylised baseline inverts the measured duration ordering, exactly as Clack (2026b) discloses — but the direction and its growth with $G$ are the result.

Table 3: The two-sided ledger under accommodated load growth: change in annual renewables revenue and in storage-fleet revenue relative to $G$=0 (£m per year; storage fleet 10 GW), and the same per gigawatt of connected data-centre load, by flexible share $s$.

| | $G$=2 | $G$=5 | $G$=10 | $G$=15 |
|---|---|---|---|---|
| Renewables, $s$=0 | +493 | +1,082 | +1,731 | +2,123 |
| Renewables, $s$=0.25 | +496 | +1,088 | +1,737 | +2,126 |
| Renewables, $s$=0.5 | +499 | +1,094 | +1,742 | +2,129 |
| Storage fleet, $s$=0 | -27 | -62 | -103 | -127 |
| Storage fleet, $s$=0.25 | -33 | -76 | -128 | -162 |
| Storage fleet, $s$=0.5 | -39 | -89 | -151 | -185 |
| Renewables per GW, $s$=0 | +246 | +216 | +173 | +142 |
| Storage per GW, $s$=0 | -13.6 | -12.4 | -10.3 | -8.5 |
| Storage per GW, $s$=0.5 | -19.4 | -17.8 | -15.1 | -12.3 |

Table 4: The scarcity bound: the same load added to a fixed supply stack (unaccommodated arm, $s$=0). Pre-registered band: mean price at $G$=15 below twice the $G$=0 mean — tripped, as the band was designed to detect.

| | $G$=0 | $G$=2 | $G$=5 | $G$=10 | $G$=15 |
|---|---|---|---|---|---|
| Mean price (£/MWh) | 65 | 82 | 113 | 204 | 383 |
| Solar capture ratio | 0.36 | 0.42 | 0.47 | 0.45 | 0.41 |
| Storage fleet (£k/MW/yr) | 40 | 46 | 62 | 126 | 243 |

## 4.3 The ledger

Table 3 puts the two sides together. At every penetration and every flexible share the renewables fleet gains and the storage fleet loses: at 10 GW, £1.7 billion a year against £103 million a year with inflexible load, and £151 million a year with half of it flexible. Per gigawatt the renewables' gain is an order of magnitude larger than the storage fleet's loss, so the two-fleet ledger is strongly positive — much of the renewables' gain is paid by consumers through the higher mean price, and the ledger says nothing of the thermal fleet or the data centre's own energy cost — but the loss lands on a 10 GW fleet whose revenue cases assume the troughs stay empty. Data-centre load is the renewables fleet's best customer and the battery fleet's newest competitor, and the flexible share decides only how much of a competitor.

## 4.4 The scarcity bound

Table 4 shows what happens when the load is not accommodated. Prices triple at 10 GW and rise nearly sixfold at 15 GW, storage revenue rises with load rather than falling, and the solar capture ratio peaks at 5 GW and then declines. Nothing here is a forecast; it is what a fixed stylised stack does with unaccommodated load, and the pre-registered band caught it; prices in this arm reach the representation's cap in some periods, so its levels are cap-dependent as well. Its use is as a bound: between accommodation and its absence lies every real planning outcome, and the sign of the battery fleet's exposure flips across that range. In the paper's device family the battery fleet's loss shrinks as the scarcity knee is shifted by less than the added load and becomes a gain once the knee is shifted by about half of it, while the capture-ratio gain survives any partial shift of the knee: the storage-fleet sign requires peak adequacy to be maintained close to fully, and every device that passed the pre-registered mean-price band kept it. That is the same conditionality Clack (2026b) discloses for the price-taking bias, seen from the load's side.

Table 5: Device sensitivity at $G$=10: the paper's accommodation device (reference and knee moved out by the added mean load) against a post-hoc pure translation of the base and scarcity terms by the same load (not pre-registered; run through the same program and shipped in the bundle). Three-year means; stylised levels.

| | Paper's device | Pure translation |
|---|---|---|
| Solar capture ratio, $s$=0 | 0.58 | 0.49 |
| Wind capture ratio, $s$=0 | 0.85 | 0.82 |
| Negative-price half-hours per year, $s$=0 | 84 | 403 |
| Mean price vs $G$=0, $s$=0 | +13% | +7% |
| Renewables revenue vs $G$=0 (£m/yr), $s$=0 | +1,731 | +914 |
| Storage fleet vs $G$=0, $s$=0 | -26% | -13% |
| Storage fleet vs $G$=0, $s$=0.5 | -38% | -26% |
| Storage ledger (£m/yr), $s$=0 / $s$=0.5 | -103 / -151 | -53 / -105 |

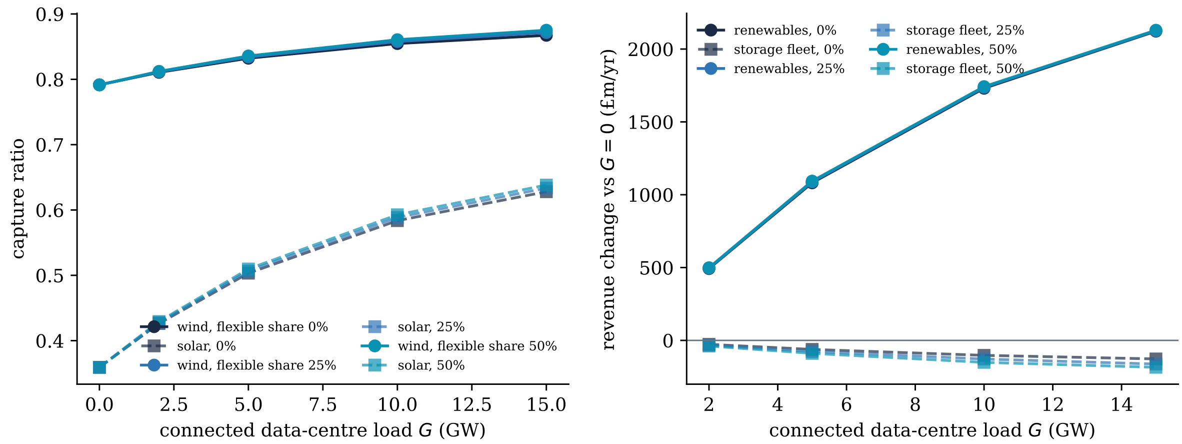


Figure 1: Accommodated load growth, 2030 background, stylised levels. Left: wind and solar capture ratios against connected data-centre load by flexible share. Right: the ledger — change in renewables revenue and in storage-fleet revenue against $G$=0, £m per year.

## 5 Discussion

**For renewables owners.** Accommodated data-centre load is a capture-rate hedge for the renewables fleet, whether or not any power purchase agreement yet prices it: the un-cannibalisation is worth more per gigawatt of load in the first gigawatts than in the last, so the earliest connections are the most valuable to the renewables fleet.

**For storage financing.** The downside scenario proposed in Clack (2026b) — gigawatt-scale flexible data-centre load — is not the only one. Inflexible data-centre load, accommodated by the planner, compresses the arbitrage spread on its own, provided the load is accommodated at the peak as well as in the mid-merit; the flexible share compounds it. A storage credit case should carry data-centre penetration and its accommodation path as scenario axes, not treat load growth as a tailwind.

**For planners.** The accommodation assumption is the switch. Load planned for is a customer of the renewables fleet and a competitor of the storage fleet; load not planned for is scarcity, in which everyone earns more and nobody should be reassured. The equilibrium framework prices both, and the difference between them is the value of the planning decision.

**Limitations.** National pricing resolution; a single 2030 background; a near-flat load profile with no locational structure; accommodation represented by one stylised device (a shift of the scarcity knee and a rescaling of the mid-merit reference by the added mean load), on which the levels, though not the signs, depend: shifting the mid-merit reference also flattens the mid-merit slope, which carries roughly half of the printed spread compression, and a post-hoc pure translation of the same stack by the same load — not pre-registered, run through the same program and shipped with the bundle — passes the same mean-price band and gives every sign and ordering at roughly half the levels (Table 5: at 10 GW the renewables fleet gains about £0.9 billion a year rather than £1.7 billion, the storage fleet loses 13% at $s$=0 and 26% at $s$=0.5 rather than 26% and 38%, solar capture reaches 0.49 rather than 0.58, about 400 negative-price half-hours remain rather than under 100, and the mean price rises 7% rather than 13%); wholesale-plus-response revenue only, with Balancing Mechanism and Capacity Market excluded; stylised levels throughout, with the duration-ordering inversion against the measured Index carried over from Clack (2026b). The solar capture band was missed at $G$=0 for the structural reason given in Section 3. The curtailment-cap check carries a 0.05 GWh absolute floor that the Clack (2026b) suite did not; it was added after two unaccommodated cells certified by the second solver overshot their annual cap by 31 and 39 MWh (2.9%), and the re-run under the new code version reproduced every cell bit-for-bit (the pre-change cells ship in the bundle). Six accommodated-arm cells that default solver settings could not certify were certified with the solver's equilibration disabled (five) or with a second solver (one), and seven unaccommodated-arm cells with the second solver; every served cell is certified optimal and the settings are recorded per cell.

## 6 Conclusion

Data-centre load growth is good for renewables and bad for batteries, at every penetration and every flexible share we ran, provided the system is planned to host it. The renewables fleet's gain is large — un-cannibalisation of the surplus hours — and diminishes as those hours are used up; the battery fleet's loss is smaller in aggregate but lands on revenue cases that assume the opposite, and a flexible share of the load deepens it. Where the load is not accommodated, scarcity makes every asset earn more, and the framework's own band says not to believe it. As everywhere in this series, the claim to be trusted on is not the mechanism but the record: pre-registered, published as it landed, with the refuted hypothesis and the tripped band in this paper alongside the confirmations (Compounding Energy Ltd, 2026).

## A Run matrix and manifest

Scenario `CE-WP-2026-06-stylised-GB`, seed 20260901, code version `bd905ab555a3` (derived from the Clack (2026b) case study, code version `2cedb716f388`, with the three switches of Section 2); pipeline stamps (carried by every artefact of the arm): accommodated arm 2026-09-03 22:03 UTC, unaccommodated arm 21:43 UTC, post-hoc translation sensitivity 23:09 UTC. Run matrix per arm: $G \in \{0, 2, 5, 10, 15\}$ GW $\times$ $s \in \{0, 0.25, 0.50\}$ $\times$ three synthetic weather years, 39 equilibrium cells (117 in total across the three arms); 2030 background; reference storage build-out 3/5/2 GW (1/2/4 h); central flexible cohort of Clack (2026b) (deferrable depth 0.25, curtailable depth 0.05, backlog 4 h, overhead 0.95, 24 h tumbling payback, load factor 0.90). Horizon: 26 representative weeks at half-hourly resolution, annualised. Verification tolerances, solver settings per cell, the pre-registration with its addendum and dated post-registration record, the per-cell results, the hypothesis and band evaluations, and the table generator (every table in this paper is generated from the results files) ship with the case-study bundle (available on request). The frequency-response clearing curve and the supply representation's parameters are stylised scenario inputs recorded in the run manifest and deliberately not reproduced here.

## References

Christopher T. M. Clack. Cannibalisation as a fixed point: revenue forecasting for battery energy storage in liberalised electricity markets. Working Paper CE-WP-2026-01, Compounding Energy Ltd, 2026a. URL https://www.compoundingenergy.com/library/CE-WP-2026-01_Cannibalisation_Fixed_Point.pdf.

Christopher T. M. Clack. The cannibalisation of flexibility: Data-centre demand response and the equilibrium value of grid flexibility in Great Britain. Working Paper CE-WP-2026-05, Compounding Energy Ltd, 2026b. URL https://www.compoundingenergy.com/library/CE-WP-2026-05_Cannibalisation_of_Flexibility.pdf.

Compounding Energy Ltd. CECadence quarterly GB BESS index, 2025Q3–2026Q2. Technical report, Compounding Energy Ltd, 2026. URL https://www.compoundingenergy.com/gb-bess-index-2026q2.html. Released 22 July 2026; computed from Elexon settlement cashflows and NESO auction records for the full 168-BMU battery registry.

Lion Hirth. The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy Economics*, 38:218–236, 2013. doi: 10.1016/j.eneco.2013.02.004.

Javier López Prol, Karl W. Steininger, and David Zilberman. The cannibalization effect of wind and solar in the California wholesale electricity market. *Energy Economics*, 85:104552, 2020. doi: 10.1016/j.eneco.2019.104552.