COMPOUNDING ENERGY LTD

Working Paper Series • CE-WP-2026-07

Who Cannibalises Whom

Duration Classes at the Storage Equilibrium

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

September 2026

Abstract. Every storage investor asks the same question of a market that already has batteries in it: which duration to build, given what everyone else has built. We answer it at equilibrium. On the unchanged stylised Great Britain machinery of the series, one-, two- and four-hour battery classes are solved jointly with half-hourly price formation across sixty-three fleet mixes, and the revenue of every class is read as a function of every class's capacity. The cross-cannibalisation matrix is negative in every entry: more of any duration lowers the revenue per megawatt of every duration. Its asymmetry has two faces, and we publish both because the pre-registered version was wrong: in pounds per megawatt on the registered step, a gigawatt of one-hour capacity takes slightly more off four-hour revenue than a gigawatt of four-hour capacity takes off one-hour revenue — long duration bears the brunt of a forward build — but as a share of each class's own revenue the ordering reverses, and the shorter class of each pair is the more exposed to the longer's build, on the stylised baselines. Own-effects are ordered by duration, four-hour steepest. Under stylised annuitised costs the entry-stable fleet is mixed, and one-hour capacity survives only where its cost is materially below the longer classes'. Levels are stylised; the signs, the own-effect ordering and the two faces of the asymmetry are the results.

Keywords: battery storage, duration, cannibalisation, equilibrium, fleet mix, Great Britain

Contents

1 Introduction

Clack (2026a) showed that battery revenue in Great Britain is a fixed point: the fleet moves the prices it is valued against, and a forecast that ignores this overstates revenue by a factor that grows with the fleet. That paper treated the fleet as one class. The market does not. Batteries are built at one, two and four hours of duration, they earn different things from the same price shape, and they are financed on revenue cases that take the rest of the fleet as given — including its duration mix. The question every investor asks is therefore a cross-cannibalisation question: which duration to build, given what everyone else has built, and what everyone else's build does to it.

This paper answers it on the series' own machinery, with nothing new in the model. The storage classes were already solved jointly in Clack (2026c) and Clack (2026d); here the fleet mix is the only thing that varies, and the object of interest is the matrix of how each class's revenue responds to each class's capacity. We pre-registered the signs and orderings we expected, ran the grid, and found the matrix's asymmetry runs the opposite way from the one we registered when measured in pounds — and the registered way when measured in shares. Both faces are reported.

2 Framework

The machinery is the case-study model of Clack (2026c), unchanged and imported as is: the stylised single-zone Great Britain of Clack (2026a) in its 2030 background (synthetic demand and renewables, strictly convex inverse supply representation), with battery classes of one, two and four hours' duration dispatched jointly with an aggregate frequency-response product inside one convex welfare program whose duals, served through the representation's stated price floor and cap, are the half-hourly prices. No data-centre load and no flexible cohort are present. Every cell is certified optimal by the solver and passes the primal verification suite at the tolerances recorded in the run manifest.

The grid. Each class's power capacity takes a value in $\{0, 2, 4, 8\}$ GW; every mix with at least one class present is solved, sixty-three mixes in three synthetic weather years, 189 equilibrium cells. A class at zero capacity is absent from the program. Per cell we read each present class's wholesale-plus-response revenue per megawatt, the mean price and the mean daily price spread (the mean over the horizon's days of the daily maximum less minimum of the served price).

Quantities. The *own-effect* of a class is its revenue per megawatt as its own capacity moves along the grid with the other two classes held at 2 GW. The *cross-cannibalisation matrix* is read at the (2, 2, 2) GW mix: the change in class i's revenue per megawatt when two gigawatts of class j are added, per gigawatt added, for every ordered pair, in pounds and as a share of class i's revenue at the reference mix. *Entry-stable mixes* are the grid mixes at which every present class covers a stated annuitised cost within 5% and no absent class would exceed its cost by more than 5% on entry at 2 GW; the band was set after registration, and the strict registered rule is evaluated alongside it. The costs are stylised scenario inputs recorded in the manifest, and this part of the analysis is descriptive.

3 Pre-registration and outturn

Four hypotheses and two bands were written before any cell was solved (`PREREGISTRATION.md` in the bundle). Own-effects negative for every class (H1): confirmed. Cross-effects negative for every ordered pair (H2): confirmed. Own-effects ordered by duration in pounds per megawatt per gigawatt, four-hour steepest, one-hour shallowest (H4): confirmed. The asymmetry hypothesis

(H3) — that per gigawatt added, four-hour capacity reduces one-hour revenue by more than one-hour capacity reduces four-hour revenue, and likewise for the two-hour pairs — was *refuted* in the registered metric: in pounds per megawatt the reverse holds, by a small margin that Section 4.2 shows to be a property of the registered two-gigawatt step. In shares of each class's own revenue the registered direction holds, by a large one. Section 4.2 reports both. The bands passed: at (2, 4, 2) GW, the grid mix nearest the reference build-out of Clack (2026c), every class's revenue lies within 25% of that paper's baseline (the machinery is the same, the mix is not), and every revenue is positive.

4 Results

[Every value in this section is generated by the run matrix of Appendix A at seed 20260901 on the unchanged Clack (2026c) model and passes the physical verification suite; tables are generated from the results files; levels are stylised. Means over three synthetic weather years.]

4.1 Revenue by mix

Table 1: Revenue per megawatt by duration class (£k/MW/yr, wholesale plus frequency response) at selected fleet mixes (P_{1h}, P_{2h}, P_{4h}) in GW; 2030 background; means over three synthetic weather years; stylised levels. A class absent from the mix is shown as –.

Mix (GW)	1 h	2 h	4 h	Mean price (£/MWh)	Mean daily spread (£/MWh)
(2, 0, 0)	75.6	–	–	70.4	161.0
(0, 2, 0)	–	112.3	–	69.5	155.4
(0, 0, 2)	–	–	157.8	68.4	153.8
(2, 2, 2)	33.0	57.2	102.1	66.5	94.0
(2, 4, 2)	25.9	43.9	84.2	65.8	78.8
(4, 2, 2)	25.9	50.3	92.1	66.1	83.2
(2, 2, 4)	23.5	40.2	70.5	65.2	72.4
(8, 2, 2)	19.8	41.4	77.0	65.5	71.7
(2, 8, 2)	18.7	29.7	60.7	64.7	61.0
(2, 2, 8)	14.1	23.2	38.3	63.6	47.6
(8, 8, 8)	6.7	12.2	22.9	62.7	34.0

Table 2: The cross-cannibalisation matrix at the (2, 2, 2) GW mix: change in each class's revenue per megawatt when two gigawatts of the column class are added, expressed per gigawatt added (£k/MW/yr per GW) and as a share of the row class's revenue at (2, 2, 2) (% per GW). Diagonal entries are own-effects.

	£k/MW/yr per GW added			% of own revenue per GW added		
Revenue of / added	1 h	2 h	4 h	1 h	2 h	4 h
1 h	-3.5	-3.5	-4.7	-10.7	-10.7	-14.3
2 h	-3.5	-6.7	-8.5	-6.1	-11.7	-14.8
4 h	-5.0	-8.9	-15.8	-4.9	-8.8	-15.5

Table 1 sets the scene. A class alone in the market at 2 GW earns a great deal — £75.6k, £112.3k and £157.8k per megawatt-year for one, two and four hours, against a mean daily spread above £150/MWh — and the same class at (2, 2, 2) GW earns £33.0k, £57.2k and £102.1k. At (8, 8, 8) GW the spread has fallen to £34/MWh and the classes earn £6.7k, £12.2k and £22.9k. Along its own axis, with the others at 2 GW, going from 2 to 8 GW costs the one-hour class 40% of its revenue per megawatt, the two-hour class 48% and the four-hour class 63%: the longer the

Table 3: Free-entry mixes under two stylised annuitised-cost scenarios (£k/MW/yr, scenario inputs): grid mixes at which every present class covers its cost within 5% and no absent class would exceed its cost by more than 5% on entry at 2 GW (entry-stability only; the band is a runner setting set after registration — the strict registered rule gives six rising-with-energy mixes and the same four flat mixes). Descriptive; no directional hypothesis was registered.

Cost scenario	Cost 1 h / 2 h / 4 h	Entry-stable mixes (GW)	Largest entry-stable fleet
rising with energy	25 / 45 / 85	(0, 4, 2), (2, 0, 4), (2, 2, 2), (2, 4, 0), (2, 4, 2), (4, 2, 2), (4, 4, 0), (8, 0, 2), (8, 2, 0)	10 GW
flat per MW	45 / 45 / 45	(0, 0, 8), (0, 2, 2), (0, 2, 4), (0, 4, 2)	8 GW

duration, the more spread-dependent the revenue, and the steeper its own cannibalisation.

4.2 The matrix

Table 2 is the paper's result. Every entry is negative: whatever is built, every class earns less per megawatt. The diagonal is ordered by duration, as registered — £3.5k, £6.7k and £15.8k per megawatt-year for each gigawatt of own capacity over the registered 2-to-4 GW step. The off-diagonal entries carry the finding we did not register correctly. In pounds, a gigawatt of one-hour capacity takes £5.0k off four-hour revenue while a gigawatt of four-hour capacity takes £4.7k off one-hour revenue; a gigawatt of two-hour takes £8.9k off four-hour while a gigawatt of four-hour takes £8.5k off two-hour. On this step long duration bears the brunt of any build. That pounds ordering is a statement about the registered forward step. On the 2-to-4 GW step it holds for both four-hour pairs at every reference mix where the two classes step by the same amount, on the larger 4-to-8 and 2-to-8 steps by 13–46%, and at the reference mix in each of the three weather years; on the entry step at any three-class reference mix it runs the registered way — at (2, 2, 2) a gigawatt of four-hour capacity entering takes £8.3k off one-hour revenue while a gigawatt of one-hour entering takes £7.6k off four-hour — and the centred difference runs the registered way by under 5%. For the one- and two-hour pair the two pounds entries of Table 2 are within about 1% of each other, the shorter class losing marginally more. The sign of the marginal pounds asymmetry is therefore not identified by the machinery, whose marginal cross-effects on per-megawatt profit are symmetric (an envelope property of the welfare program) and depart from symmetry in gross revenue only through the model's small cost and price-floor terms; what the step reveals is the unequal curvature of the revenue curves. As a share of each class's own revenue, the ordering reverses and the gap widens: a gigawatt of four-hour capacity costs the one-hour class 14.3% of its revenue, while a gigawatt of one-hour capacity costs the four-hour class 4.9%; the two-hour pair is 14.8% against 8.8%. Within each pair, the shorter duration is the more exposed to the longer's build in the terms a financing is written in — share of a revenue case, on these baselines — and the four-hour class is the most exposed in absolute pounds to any forward build. The share face is carried by the revenue-level ratio (3.1 for four-hour to one-hour, 1.8 for four-hour to two-hour) and holds for both four-hour pairs on every step, at every reference mix and in every weather year. The same duality appeared in Clack (2026c), where the four-hour cohort lost the most per megawatt to flexible data-centre load and the one-hour cohort the largest share outright; here the share lost to four-hour build is nearly flat across classes (14.3–15.5%) and the asymmetry is pairwise. It is a property of a market in which revenue per megawatt rises with duration, and it is why a cross-cannibalisation claim must state its denominator and its step.

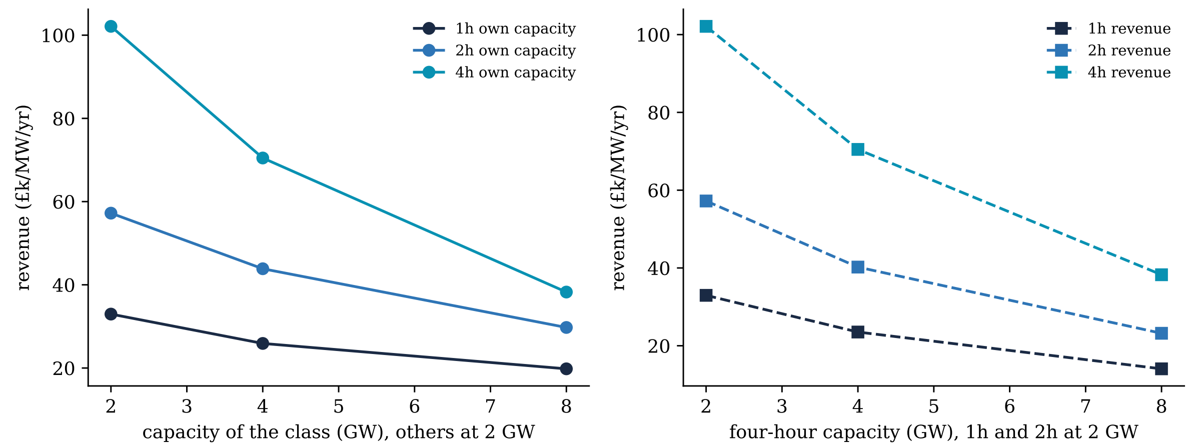

Figure 1: Left: own-effects — each class's revenue per megawatt against its own capacity with the other two classes at 2 GW. Right: every class's revenue against four-hour capacity with one- and two-hour capacity at 2 GW. 2030 background; stylised levels.

4.3 Free entry

Table 3 asks which mixes are entry-stable under stylised costs. With annuitised costs rising with energy (£25k, £45k and £85k per megawatt-year for one, two and four hours), nine grid mixes are entry-stable within the 5% band of Table 3 (six under the strict registered rule, which drops (2, 4, 2), (8, 0, 2) and (0, 4, 2)), and the largest on the grid carries 10 GW either way, with its eight-gigawatt class still profitable at the grid edge; they are mixed, the short-heavy (8, 2, 0) survives, and the reference-like (2, 4, 2) survives within the band with its two-hour and four-hour classes marginally below cost. With a flat £45k per megawatt-year for every class, only mixes without one-hour capacity are entry-stable, under either rule, and the largest on the grid carries 8 GW: at that cost the one-hour class covers itself alongside another class only at (2, 2, 0) and (2, 0, 2), where the absent third class would enter, and nowhere on the grid once all three classes are present. Entry-stability is the whole of the test: present classes cover cost within the band and no absent class would cover cost on a 2 GW entry; whether a present class would profitably expand is not tested, and the entrant is a 2 GW lump. Adding a no-expansion condition at the next grid level and dropping the band leaves three rising-with-energy mixes ((2, 0, 4), (4, 2, 2), (4, 4, 0)), largest 8 GW, and two flat mixes, and (2, 4, 2) is entry-stable only within the band and not in one of the three weather years. The counts are rule-specific; what survives the band, expansion and weather-year variants is that no flat-cost mix contains one-hour capacity and every rising-with-energy set does. The equilibrium duration mix is therefore a statement about relative costs as much as about the matrix, and we make no forecast of it; what the matrix fixes is that whichever classes are built, each one's revenue case must be stressed against the others' build, in both denominators, on these baselines.

5 Discussion

For storage investors. A duration decision made against today's fleet is made against a mix that the decision itself will change, and every column of Table 2 is negative: there is no duration whose build helps another. The practical reading is by denominator. Each class's revenue case should be stressed against every other class's build, in both denominators; which class bears the largest stress in shares depends on the baseline ordering, which the measured fleet inverts. Both stresses are absent from revenue cases that hold the rest of the fleet fixed.

For fleet-level analysis. The single-class fixed point of Clack (2026a) is a special case of this matrix, and the seven-stream stack of Clack (2026b) is where the full frequency-response channel, here reduced to one aggregate product, would enter. The class structure used here — one, two and four hours — is the one the series' storage work has carried since Clack (2026c).

Limitations. A coarse grid (four levels per class; the matrix is read at one reference mix and its forward two-gigawatt steps; the pounds face reverses on the entry step, and per-gigawatt magnitudes fall by roughly half between the 2-to-4 and 4-to-8 steps); a single 2030 background; wholesale-plus-response revenue only, with Balancing Mechanism and Capacity Market excluded; costs illustrative and stated, with no forecast of the equilibrium mix, and entry-stable counts that move with the band, the entry lump and the omitted expansion test; the split of gross revenue between classes identified only to within about one percent, since a megawatt-hour of throughput can migrate between classes at almost no welfare cost (signs, orderings and the registered-step verdicts are unaffected); stylised levels throughout, including a baseline in which revenue per megawatt rises with duration, the reverse of the ordering in the measured CECadence Quarterly GB BESS Index (Compounding Energy Ltd, 2026), so that both faces of the matrix — the pounds face and the proportional face, whose denominators are those same stylised baselines — are properties of the stylised setting; on measured denominators the proportional ordering would reverse. The signs, the own-effect ordering and the existence of two faces are the results. Twenty-seven cells that default solver settings could not certify were certified with the solver's equilibration disabled; every served cell is certified optimal and the settings are recorded per cell.

6 Conclusion

Nobody's build helps anyone else's revenue: the cross-cannibalisation matrix of one-, two- and four-hour batteries at the storage equilibrium is negative in every entry, with own-effects steepest for the longest duration. Its asymmetry depends on the denominator, and we registered the wrong one: long duration loses the most pounds to a forward build, and the shorter class of each pair the larger share to the longer's build, on the stylised baselines. Both faces are now on the record, with the refuted hypothesis published alongside the confirmations, on machinery that has not changed since Clack (2026c), in a series whose record is graded in public (Compounding Energy Ltd, 2026).

A Run matrix and manifest

Scenario `CE-WP-2026-07-stylised-GB`, seed 20260901, code version `2cedb716f388` — the Clack (2026c) case-study model unchanged; the run script only sets the fleet mix — pipeline stamp 2026-09-04 12:27 UTC. Run matrix: $(P_{1h}, P_{2h}, P_{4h}) \in \{0, 2, 4, 8\}^3$ GW excluding the empty fleet, × three synthetic weather years, 189 equilibrium cells; 2030 background (demand scale 1.06, renewables scale 1.45 on the Clack (2026a) generators); no data-centre load. Horizon: 26 representative weeks at half-hourly resolution, annualised. Cost scenarios (stylised scenario inputs, £k/MW/yr): rising with energy 25/45/85, flat 45/45/45; entry-stability band ±5% of cost, a runner setting not in the pre-registration text, with the strict rule recorded alongside; the entry-step matrix at (2, 2, 2) is recorded with the results. Verification tolerances, solver settings per cell (162 default, 27 with equilibration disabled), the pre-registration, the per-cell results and the table generator 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 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. Co-optimising the GB BESS revenue stack: wholesale energy and the NESO dynamic-response suite at the cannibalisation equilibrium. Working Paper CE-WP-2026-02, Compounding Energy Ltd, 2026b. URL https://www.compoundingenergy.com/library/CE-WP-2026-02_GB_BESS_Revenue_Stack.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, 2026c. URL https://www.compoundingenergy.com/library/CE-WP-2026-05_Cannibalisation_of_Flexibility.pdf.

Christopher T. M. Clack. Load growth and the un-cannibalisation ledger: What data-centre demand does to renewable capture rates — and to the battery fleet. Working Paper CE-WP-2026-06, Compounding Energy Ltd, 2026d. URL https://www.compoundingenergy.com/library/CE-WP-2026-06_Load_Growth_Ledger.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.