COMPOUNDING ENERGY LTD

Working Paper Series • CE-WP-2026-02

Co-optimising the GB BESS revenue stack

wholesale energy and the NESO dynamic-response suite at the cannibalisation equilibrium

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

September 2026

Abstract. Battery energy storage systems (BESS) in the Great Britain market earn from a stack of co-optimised revenue streams: wholesale energy arbitrage and the full NESO dynamic-response suite — Dynamic Containment, Dynamic Moderation and Dynamic Regulation, each in a Low and a High variant. We treat seven streams jointly: wholesale (W) and the six frequency-response (FR) products DC-L, DC-H, DM-L, DM-H, DR-L, DR-H. Industry practice forecasts each stream with a separate price-taker model and adds the totals. This naive sum is not realisable: a single battery cannot deliver every product at full power simultaneously, and committing volume to any one stream depresses its own clearing price. We extend the cannibalisation framework of CE-WP-2026-01 from wholesale-only to the seven-stream stack via a single *convex* welfare/competitive program — the gross-surplus integrals of the FR clearing curves minus wholesale system cost minus throughput degradation, maximised over a schedule constrained by a shared power and state-of-charge budget. The program is concave, and strictly concave in the price-relevant decisions, which are therefore unique; its equilibrium prices are the constraint marginals (wholesale $p^* = S(R^*)$, FR $\pi_t^{p*} = \mathcal{C}^p(b_t^{p*})$), and plain Picard iteration on the μ-regularised seven-stream operator contracts to fixed-point residual below 0.02 in 9 iterations. We report a fully reproducible *stylised* case study (no real settlement data; clearing curves hand-calibrated to documented historical GB targets). At a 6 GW two-hour fleet the equilibrium total is £124,666/MW/yr against a naive £689,774/MW/yr — an overstatement of 453%, rising from 387% at 4 GW to 528% at 10 GW; holding the FR leg to NESO-scale procurement volumes with settled response energy lowers the equilibrium level and raises the overstatement further (to 998% at 6 GW under a 2.5 GW total FR cap), so the headline is conservative. With FR pinned to zero the model reproduces CE-WP-2026-01's wholesale-only bias (149.5/236.5/332.7% at 4/6/8 GW, same seed and generators); adding the stack *raises* the bias, because a naive forecast treats six FR products as independently committable at full power. The shared physical budget is the paper's core contribution; we formalise it so that it provably prevents MW/MWh double-counting.

Keywords: battery energy storage, ancillary services, dynamic containment, dynamic moderation, dynamic regulation, revenue stacking, cannibalisation, shared-capacity budget, GB market, NESO

Contents

1 Introduction

GB battery energy storage systems generate revenue from a stack of distinct streams: wholesale energy arbitrage (W) and the NESO dynamic-response suite — Dynamic Containment (DC), Dynamic Moderation (DM) and Dynamic Regulation (DR), each procured in a Low and a High variant (NESO, 2026). These seven streams are the operational substrate of any defensible revenue forecast, but they are not the whole stack: the Balancing Mechanism is the second-largest GB BESS revenue stream (39% of trailing-year fleet revenue in the 2026Q2 CECadence GB BESS Index (Compounding Energy Ltd, 2026)) and the Capacity Market is material; both sit outside the scope of this paper, so the totals below are not whole-stack revenue. Industry practice is to forecast each stream separately — each with its own price-taker model — and add the totals.

The naive sum is wrong, for two distinct reasons. First, the same physical battery cannot be in all seven streams simultaneously: in any half-hour the dispatch must allocate the asset's power capacity and its state-of-charge (SoC) envelope across the streams subject to intertemporal feasibility. A battery cannot be in seven places at once. Second, the volume an operator commits to a stream endogenously moves the price that stream pays — the wholesale cannibalisation analysed in CE-WP-2026-01 (Clack, 2026) is the canonical case, and each FR product exhibits the same saturation as fleet supply into it grows. The naive sum ignores both: it commits every product at full power against un-cannibalised prices.

This paper extends the cannibalisation framework of CE-WP-2026-01 from the wholesale-only setting to the seven-stream stack. The wholesale leg is reused *verbatim* from that paper; the contribution here is (i) a *shared physical budget* that couples the streams through the battery's finite power and energy, formalised so that it provably prevents double-counting of MW and MWh; (ii) a single *convex* equilibrium program whose maximiser is unique in the price-relevant decisions and whose prices are the constraint marginals; and (iii) a fully reproducible, explicitly *stylised* case study that quantifies the overstatement of the naive full-stack sum and reconciles it exactly with CE-WP-2026-01.

A note on framing. This is a *deterministic* convex-equilibrium model, not a stochastic one. Uncertainty enters only through the static, per-FR activation fraction ρ_p that scales each product's expected SoC draw; there is no scenario tree, no stochastic program, and no claim of one. The earlier draft of this paper described a "stochastic stack model"; that description was inaccurate and has been removed.

1.1 What is new relative to CE-WP-2026-01

The single-product wholesale equilibrium of CE-WP-2026-01 is extended here along three axes:

1. **Seven co-optimised streams.** The decision is the fleet's per-stream commitment over the horizon for $p \in \mathcal{S} = \{\text{W}, \text{DC-L}, \text{DC-H}, \text{DM-L}, \text{DM-H}, \text{DR-L}, \text{DR-H}\}$. Wholesale is a two-directional energy position; the six FR products are directional availability commitments (Section 2).
2. **A shared physical budget.** A single battery's power and energy are shared across the streams. We impose a discharge-direction power budget, a charge-direction power budget, and a SoC/headroom budget that together make the naive "seven places at once" allocation infeasible by construction (Section 3).
3. **A convex equilibrium.** Rather than iterate the (discontinuous) LP-dispatch operator of CE-WP-2026-01, we characterise the competitive self-consistent fleet as the maximiser of one convex program: FR gross-surplus integrals (concave) minus wholesale system cost (convex) minus degradation (linear), over the shared-budget-feasible schedule (Section 4). The equilibrium prices are the constraint marginals.

The headline empirical finding (Section 5) is that the naive sum of seven single-product forecasts overstates equilibrium fleet revenue by 387–528% across 4–10 GW of two-hour fleet, 453% at the representative 6 GW point — and by more still once the FR leg is held to NESO-scale procurement volumes with settled response energy (Section 5.5). The overstatement decomposes into approximately 91% own-product price cannibalisation and 9% shared-capacity residual — the latter being a channel that a wholesale-only model cannot see.

Practitioner sidebar

What this means for revenue-stack analysts. If you sum revenue forecasts for wholesale, DC-L, DC-H, DM-L, DM-H, DR-L and DR-H produced by separate price-taker models, you have a number the underlying asset cannot realise — the battery would have to be in seven places at once, and its act of participating would itself depress every clearing price. The error is small at low fleet penetration, where most assets stay below their power and SoC envelopes; it becomes structural and dominant at the fleet sizes projected for 2030 — NESO's Clean Power 2030 pathway implies roughly 23–27 GW of batteries (∼25 GW central) (NESO, 2024). Investment-grade forecasts must compute the equilibrium across the streams jointly.

2 The seven GB revenue streams

2.1 Stream definitions

We summarise the streams at the level needed for the formulation. Operational detail is in NESO's frequency-response service terms and the Enduring Auction Capability market documentation (NESO, 2026).

Wholesale (W). Day-ahead and intra-day energy trades in the N2EX and EPEX markets, settled at the marginal energy price π_t^{W} (£/MWh) in each half-hour. We split the position into a discharge component $b_t^{\mathrm{W,d}} \geq 0$ and a charge component $b_t^{\mathrm{W,c}} \geq 0$, with net injection $b_t^{\mathrm{W}} = b_t^{\mathrm{W,d}} - b_t^{\mathrm{W,c}}$.

Dynamic Containment (DC-L, DC-H). The fastest post-fault frequency-response service, holding reserve for activation on a post-fault excursion and delivering to contain frequency within the statutory ±0.5 Hz limits. The two variants are *directional*, not symmetric: **DC-L** (Low) commits *discharge* headroom $b_t^{\text{DC-L}}$ that injects on under-frequency; **DC-H** (High) commits *charge* headroom $b_t^{\text{DC-H}}$ that absorbs on over-frequency. The asset is paid an availability price (£/MW/h) for *holding* the headroom, whether or not activation occurs.

Dynamic Moderation (DM-L, DM-H). A pre-fault service that acts within the ±0.2 Hz operational band, before a post-fault excursion develops, with a tighter deadband than DC. Directional in the same sense: DM-L discharge, DM-H charge. The earlier draft of this paper omitted DM entirely; its inclusion completes the NESO dynamic suite.

Dynamic Regulation (DR-L, DR-H). A slower, continuously-acting service correcting persistent low-magnitude frequency deviation, with the longest sustain requirement of the three. Directional: DR-L discharge, DR-H charge.

Each FR product p carries a sustain-duration parameter τ_p (the MWh of headroom per MW of commitment) and a static activation fraction ρ_p (the expected per-half-hour delivery fraction of the committed power, which perturbs the SoC trajectory). DC has the shortest sustain and lowest activation; DR the longest and highest. The Low variants draw on *dischargeable* energy;

the High variants draw on *charging* headroom. The six FR products clear in a single co-optimised day-ahead auction — NESO's Enduring Auction Capability, live since November 2023 (NESO, 2026) — in which higher-quality products may substitute for lower-quality requirements; v1 models each product's own-price curve only (Section 6.2). Wholesale clears separately. All seven streams draw on the same physical battery — which is the entire point.

2.2 FR market clearing

Each FR market clears by intersecting the BESS fleet's stacked availability bid with a NESO-determined requirement. The marginal price is the bid of the marginal MW. We model this with a strictly decreasing inverse clearing curve: if the fleet commits Q^p aggregate MW to product p, the price clears at

$$\pi^p = \mathcal{C}^p(Q^p), \qquad \mathcal{C}^p \text{ strictly decreasing on } [0, \overline{Q}^p], \tag{1}$$

saturating from an unsaturated cap $\overline{\pi}^p$ down to a floor $\underline{\pi}^p$ as fleet supply grows (each new MW of BESS displaces a more expensive bid out of the marginal slot); here Q^p is the fleet aggregate committed to product p in the settlement period. This is the FR analogue of the wholesale supply curve: own-product cannibalisation. We calibrate each $\mathcal{C}^p$ to documented historical GB saturation behaviour — for example DC's collapse from roughly £17/MW/h pre-saturation toward a £3–5/MW/h floor as the fleet grew through 2024–2025, a historical reference window; per-product clearing outturns are tracked quarterly in the CECadence GB BESS Index (Compounding Energy Ltd, 2026) (Figure 1).

Practitioner sidebar

The FR cannibalisation mechanism. As the GB BESS fleet grew from roughly 0.4 GW in 2020 (historical) to 6.4 GW of active metered capacity by 2026Q2 (Compounding Energy Ltd, 2026), frequency-response prices collapsed by 60–80% in real terms. This is supply-side cannibalisation: each new MW of BESS joining the bid stack displaces a more expensive bid out of the marginal slot. It is the same mathematical structure as wholesale cannibalisation, but the supply is now BESS-dominated, the price is auction-cleared, and the "demand" is a NESO-set requirement rather than residual load.

3 The shared physical budget

The core structural contribution of this paper is the shared budget that couples the seven streams through one battery. We state it first, in isolation, because it is what makes the naive sum infeasible and what the convex program of Section 4 optimises over.

3.1 Notation

We use the CE-WP-2026-01 conventions and add a stream index $p \in \mathcal{S}$, where

$$\mathcal{S} = \{\text{W, DC-L, DC-H, DM-L, DM-H, DR-L, DR-H}\}.$$

The fleet has aggregate power $P^{\max}$ (MW) and energy $E^{\max} = h\,P^{\max}$ (MWh) for a h-hour duration. Write the discharge-direction (Low) FR set $\mathcal{L} = \{\text{DC-L, DM-L, DR-L}\}$ and the charge-direction (High) set $\mathcal{H} = \{\text{DC-H, DM-H, DR-H}\}$. The half-hour step is $\Delta t = 0.5$ h. All commitments are non-negative.

3.2 The three budget constraints

In every half-hour t the battery's finite power and energy are shared:

$$\text{(discharge power)}\quad b_t^{\mathrm{W,d}} + \sum_{p\in\mathcal{L}} b_t^p \;\le\; P^{\max}, \qquad t\in\mathcal{T}, \tag{2}$$

$$\text{(charge power)}\quad b_t^{\mathrm{W,c}} + \sum_{p\in\mathcal{H}} b_t^p \;\le\; P^{\max}, \qquad t\in\mathcal{T}, \tag{3}$$

$$\text{(SoC headroom)}\quad \sum_{p\in\mathcal{L}} \tau_p\, b_t^p \;\le\; e_t \;\le\; E^{\max} - \sum_{p\in\mathcal{H}} \tau_p\, b_t^p, \qquad t\in\mathcal{T}, \tag{4}$$

where e_t is the SoC at the end of half-hour t, evolving under

$$e_t = e_{t-1} + \eta_c\, b_t^{\mathrm{W,c}}\,\Delta t - \tfrac{1}{\eta_d}\, b_t^{\mathrm{W,d}}\,\Delta t - \sum_{p\in\mathcal{L}} \rho_p\,\eta_d^{-1}\, b_t^p\,\Delta t + \sum_{p\in\mathcal{H}} \rho_p\,\eta_c\, b_t^p\,\Delta t, \tag{5}$$

with $0 \le e_t \le E^{\max}$, charge/discharge efficiencies η_c, η_d ($\eta_c\eta_d = \eta$), the SoC initialised at $\sigma_0 E^{\max}$, and a terminal condition $e_T \ge \sigma_0 E^{\max}$. The expected-activation terms in (5) are the average energy a held reserve actually draws (Low) or injects (High) when called: ρ_p is a per-half-hour delivery fraction of the committed power, passed through the converter efficiency on its way out of (Low) or into (High) the store. (The FR terms of (5) are machine-generated from the same single definition the solver consumes, and an equality test in the reproducibility bundle asserts the two cannot drift; in the base case this response energy shifts the SoC but is not settled — Section 5.5 prices it and discloses the subsidy.)

Proposition 3.1 (No MW or MWh double-counting)**.** *Any schedule satisfying* (2)–(5) *commits, in each half-hour, total discharge-direction power no greater than* $P^{\max}$*, total charge-direction power no greater than* $P^{\max}$*, and reserves dischargeable and charging energy headroom that fit within* $[0, E^{\max}]$ simultaneously. *Consequently no MW of power is committed to more than one stream within a direction, and no MWh of dischargeable or charging headroom is reserved twice: the naive full-power allocation* $b_t^{\mathrm{W,d}} = b_t^p = P^{\max}$ *for every discharge-direction stream is infeasible whenever more than one such stream is active.*

Proof. Constraints (2) and (3) are the explicit power-balance budgets in each direction; their right-hand sides are the single shared $P^{\max}$, so the sum of all discharge-direction commitments (wholesale discharge plus the three Low services) and, separately, all charge-direction commitments cannot exceed the one physical inverter rating. Constraint (4) requires that the energy reserved as Low headroom ($\sum_{\mathcal{L}} \tau_p b_t^p$) and the energy reserved as High headroom ($\sum_{\mathcal{H}} \tau_p b_t^p$) both lie within the same SoC window $[0, E^{\max}]$ around the common state e_t, which itself obeys the single SoC balance (5). Hence the same MW and the same MWh cannot be counted twice. The naive allocation sets every discharge-direction commitment to $P^{\max}$ simultaneously; summing two or more such terms violates (2) unless all but one are zero. □ □

Proposition 3.1 is the formal statement of "a battery cannot be in seven places at once". It is a pure quantity-coupling result: it constrains how much capacity and energy each stream may use, independently of price. The price-side cannibalisation of Section 2.2 acts on top of it. Two scope notes. First, the two directions are budgeted separately: a battery at part load can hold discharge-direction and charge-direction headroom on the same inverter simultaneously, so up to $P^{\max}$ may be committed in *each* direction at once; what the budget excludes is double-counting within a direction and double-reservation of the SoC window. Second, GB procures response in four-hour EFA blocks with one clearing price per product per block; v1 lets b_t^p vary at half-hourly resolution, which is an upper bound on the FR schedule a real participant could realise.

3.3 Degradation as a marginal cost

Throughput degradation enters as a linear marginal cost κ (£/MWh) on energy cycled:

$$\Gamma(\boldsymbol{b}) = \kappa \sum_{t\in\mathcal{T}} \Big(b_t^{\mathrm{W,d}} + b_t^{\mathrm{W,c}} + \sum_{p\in\mathcal{L}\cup\mathcal{H}} \rho_p\, b_t^p \Big)\, \Delta t, \tag{6}$$

charging wholesale throughput in full and FR throughput at its activation fraction. We take $\kappa = 2$ £/MWh in the base case and report a sensitivity over $\kappa \in [0,4]$ £/MWh in Section 5. Because Γ is linear it preserves convexity of the equilibrium program.

4 The convex equilibrium

4.1 Why a convex program, not an LP fixed point

CE-WP-2026-01 was rebuilt because the operator fixed point $\mathcal{T}(\boldsymbol{b}) = \mathcal{D}(\mathcal{P}(\boldsymbol{b}))$ does not exist pointwise: a linear dispatch is bang-bang, so $\mathcal{D}$ is discontinuous in price and $\mathcal{T}$ is not non-expansive. We reuse the fix here. A competitive price-taking fleet that is collectively self-consistent attains the welfare optimum — the storage analogue of the first welfare theorem — so the equilibrium can be characterised directly as the maximiser of a single convex program, with prices recovered as the constraint marginals. This is mathematically cleaner than iterating a discontinuous map and removes the unproven non-expansiveness assumption the earlier draft relied on.

4.2 The equilibrium program

For the wholesale leg, let Φ_{W} be the convex potential of the CE-WP-2026-01 inverse supply curve S, i.e. $\Phi'_{\mathrm{W}} = S$, evaluated at the residual demand. For each FR product, the competitive self-consistent outcome maximises the FR *gross-surplus integral* $\int_0^{Q^p} \mathcal{C}^p(q)\,\mathrm{d}q$, which is concave because $\mathcal{C}^p$ is decreasing. The seven-stream equilibrium is the maximiser of

$$\max_{\boldsymbol{b},\boldsymbol{e}} \quad \underbrace{\sum_{p\in\mathcal{L}\cup\mathcal{H}} \sum_{t\in\mathcal{T}} \int_0^{b_t^p} \mathcal{C}^p(q)\,\mathrm{d}q\,\Delta t}_{\text{FR gross surplus (concave)}} - \underbrace{\sum_{t\in\mathcal{T}} \Phi_{\mathrm{W}}\big(R_t^0 + b_t^{\mathrm{W}}\big)}_{\text{wholesale system cost (convex)}} - \underbrace{\Gamma(\boldsymbol{b})}_{\text{degradation (linear)}} \tag{7}$$
$$\text{s.t.} \quad \text{shared budget (2)–(5)},$$

where R_t^0 is the non-storage residual demand and the fleet aggregate committed to product p in half-hour t is $Q_t^p = b_t^p$ (the auction clears on the aggregate per settlement period, so the surplus integral is evaluated at the instantaneous commitment). The objective is concave-plus-convex with a sign that makes the whole maximand concave, and the feasible set (2)–(5) is a non-empty compact convex polyhedron (the trivial all-zero schedule is feasible). Hence (7) has a maximiser, and the price-relevant decisions are unique on the operating range.

Theorem 4.1 (Existence and uniqueness in the price-relevant decisions)**.** *Under CE-WP-2026-01 Assumptions 2.1 and 4.1 on S (so Φ_{W} is convex, coercive, and strictly convex on the no-clip region) and with each $\mathcal{C}^p$ continuous, strictly decreasing and bounded, the program (7) attains its maximum. The price-relevant decisions — the FR commitments b_t^p and the net wholesale position $b_t^{\mathrm{W}} = b_t^{\mathrm{W,d}} - b_t^{\mathrm{W,c}}$ — are unique wherever the wholesale residual lies in the strict interior of the no-clip region of S and the FR commitments are interior to their saturation range. The (charge, discharge) decomposition of a given net position is pinned by the linear degradation term when $\kappa > 0$ and may be non-unique at $\kappa = 0$.*

Proof. The maximand is the sum of a concave term (each $\int_0^{b_t^p} \mathcal{C}^p$ is concave in b_t^p since $\mathcal{C}^p$ is decreasing, and strictly concave where $\mathcal{C}^p$ is strictly decreasing), a concave term $-\Phi_{\mathrm{W}}$ (since Φ_{W}

is convex, strictly so on the no-clip region), and a linear term $-\Gamma$; it is therefore concave, and strictly concave in $(b_t^p, b_t^{\mathrm{W}})$ on the interior region by strict monotonicity of S and each $\mathcal{C}^p$ there. The constraint set is compact and convex. Weierstrass gives existence; strict concavity in the price-relevant coordinates gives their uniqueness there. The objective depends on $(b_t^{\mathrm{W,c}}, b_t^{\mathrm{W,d}})$ beyond the net position only through Γ, whence the final claim. Because the best-response of a price-taking fleet against *fixed* equilibrium prices is a linear program, the price-taker schedule may still be degenerate; the certificate of Section 4.3 therefore compares operating profit, the uniquely pinned invariant. □ □

4.3 Equilibrium prices are the constraint marginals

The competitive prices are recovered as the marginals (duals) of (7). The wholesale price is the dual of the power balance,

$$\pi_t^{\mathrm{W}*} = S\big(R_t^0 + b_t^{\mathrm{W}*}\big), \tag{8}$$

exactly the CE-WP-2026-01 price relation; and each FR price is the marginal clearing bid at the half-hour's aggregate commitment,

$$\pi_t^{p*} = \mathcal{C}^p(b_t^{p*}), \qquad p \in \mathcal{L} \cup \mathcal{H} \tag{9}$$

(the FR prices reported in the tables below are commitment-weighted means of this half-hourly marginal price). We *certify* that the welfare maximiser is the competitive equilibrium by re-solving the fleet as a price-taker against $(\pi_t^{\mathrm{W}*}, \pi_t^{p*})$ — after the SoC-preserving no-simultaneity reconstruction of the wholesale schedule — and checking that the operating profit matches the welfare solution: the relative gap is 2.0×10^{-8} at the 6 GW point, below 2×10^{-7} over the 4–10 GW band, and below 5×10^{-5} at every point of the full 0.5–10 GW sweep (the sub-4-GW points are the loosest; per-point gaps ship in the results bundle), confirming $\boldsymbol{b}^* = \mathcal{D}(\mathcal{P}(\boldsymbol{b}^*))$ to the stated tolerance.

4.4 The naive full-stack forecast

The object we measure the equilibrium against is the naive sum, as commonly presented in vendor revenue stacks: each stream forecast independently as a price-taker, at full power, against *uncannibalised* prices, then added:

$$\Pi^{\mathrm{naive}} = \sum_{p\in\mathcal{S}} \Pi_p^{\mathrm{naive}}, \qquad \Pi_{\mathrm{W}}^{\mathrm{naive}} = \sum_t \pi_t^{\mathrm{W}} \left(b_t^{\mathrm{W,d}} - b_t^{\mathrm{W,c}}\right) \Delta t, \quad \Pi_p^{\mathrm{naive}} = \sum_t \mathcal{C}^p(0)\, b_t^p \,\Delta t, \tag{10}$$

where each FR price is taken at zero fleet supply $\mathcal{C}^p(0)$ and the wholesale price at the zero-BESS supply curve, and each b_t^p is the independent price-taker dispatch *ignoring* the shared budget. The naive sum overstates for two reasons: it treats prices as exogenous within each stream (price cannibalisation) and it commits the same MW and MWh into multiple streams at once (shared-capacity infeasibility, ruled out by Proposition 3.1).

4.5 A convergent fixed-point iteration

The headline equilibrium is the convex program (7). To keep the fixed-point narrative of CE-WP-2026-01 literally true on the seven-stream operator, we also iterate $\boldsymbol{b} \leftarrow \mathcal{D}_\mu(\mathcal{P}(\boldsymbol{b}))$, where $\mathcal{D}_\mu$ adds a small strictly-convex penalty $\frac{\mu}{2}\sum(b_t^p)^2$ to the per-iterate dispatch. The penalty makes each iterate a unique strongly-concave QP whose price-to-dispatch map is Lipschitz $(1/\mu)$; the composed seven-stream map is then a contraction for μ above the steepest *effective* price slope — which is the wholesale supply-curve slope over the operating range, not the (much shallower) FR clearing-curve slopes. As $\mu \to 0$ the regularised equilibrium approaches the welfare optimum, so the headline uses the μ-free convex program while the convergence diagnostic uses the regularised operator. (The diagnostic operator additionally prices each FR leg at its horizon-mean commitment rather than through the per-half-hour clearing map; the headline equilibrium is computed from the exact per-half-hour program, so the simplification affects the convergence diagnostic only.)

Two corrections to the earlier draft's convergence reporting. First, that draft displayed a plain update but ran Krasnoselski–Mann (KM) averaging with the vanishing step $\lambda_k = 2/(k+2)$, $\boldsymbol{b}_{k+1} = (1-\lambda_k)\,\boldsymbol{b}_k + \lambda_k\,\mathcal{D}_\mu(\mathcal{P}(\boldsymbol{b}_k))$, and reported the price change between successive *averaged* iterates — a quantity that shrinks with the step size λ_k whether or not the operator contracts, and that understates the unrelaxed fixed-point residual $\|\mathcal{P}(\mathcal{T}(\boldsymbol{b}_k)) - \mathcal{P}(\boldsymbol{b}_k)\|_\infty$ by roughly a factor $1/\lambda_k$. We now report the unrelaxed residual, for three step schedules including a plain-Picard control ($\lambda = 1$) on the same regularised map. Second, the earlier draft's "KM averaging converges monotonely" claim is withdrawn: no such guarantee holds for a diminishing-step scheme, and its reported decay tracked the step schedule (algebraic, $\sim k^{-3}$), not the operator. The diagnostic (Section 5.8, Figure 5) shows the regularised operator itself *does* contract: plain Picard converges geometrically to sub-tolerance in a handful of iterations, which is the honest version of the fixed-point claim.

5 Stylised GB case study

5.1 Status: stylised, not validated on settlement data

Practitioner sidebar

This case study uses no real settlement data. The demand, renewables, wholesale supply curve and the six FR clearing curves are synthetic, hand-calibrated to documented historical GB targets (e.g. DC's £17 → £3–5/MW/h saturation over 2024–2025, retained as CE-WP-2026-01's historical calibration anchor; as of 2026Q2 the metered fleet stands at 6.4 GW (Compounding Energy Ltd, 2026)) and to the CE-WP-2026-01 wholesale reconciliation targets. These are stylised calibration choices for a reproducible mechanism study; they bear no relationship to any production calibration. The artefact demonstrates the shared-capacity co-optimisation *mechanism* and its solver — not a market forecast. Reproducibility key: `scenario_id = CE-WP-2026-02-stylised-GB-7stream`, `code_version = 3213bd46`, `timestamp = 2026-08-29T14:54:43Z`; flags `stylised = true`, `validated_on_real_settlement_data = false`.

The earlier draft claimed a validation against NESO settlement data — a 24-month backtest with a 5% mean absolute error and a +42% naive bias. No such validation was performed; those numbers were assertions and have been removed in full. The case study below replaces them with a reproducible stylised experiment.

5.2 Calibration

The wholesale supply curve S is the CE-WP-2026-01 sigmoid-plus-scarcity form (price floor −50, cap 2000, VOLL 10,000 £/MWh — model calibration choices inherited from CE-WP-2026-01, not GB market limits; same seed 20,260,901 and demand/renewables series as CE-WP-2026-01). Each FR clearing curve $\mathcal{C}^p$ saturates from an unsaturated cap toward a floor; the calibrated $(\overline{\pi}^p, \underline{\pi}^p)$ in £/MW/h are DC-L $(17, 3)$, DC-H $(15, 3)$, DM-L $(10, 2)$, DM-H $(9, 2)$, DR-L $(8, 2)$, DR-H $(7, 2)$, with sustain τ of 0.25 h (DC, its published 15-minute sustain requirement), 0.5 h (DM) and 1.0 h (DR) — the DM and DR values are stylised choices, not quoted service-term parameters — and activation ρ of 0.04 (DC), 0.06 (DM), 0.12 (DR). The strictly positive price floors are a modelling convenience that recent outturns refute — in 2026Q2 DR-H cleared *negative* and DR-L cleared above this paper's un-cannibalised DR-L cap (Compounding Energy Ltd, 2026)

— so the DR rows below should be read as mechanism illustration, not level guidance. Figure 1 shows the six curves. Headline revenues are computed over a 26-week horizon of 8,736 half-hours and scaled to annual by $365/182 \approx 2.0055$, as in CE-WP-2026-01; the simulated window covers the winter and spring half of the synthetic year, and Section 5.5 reports a full-year check of the resulting bias.

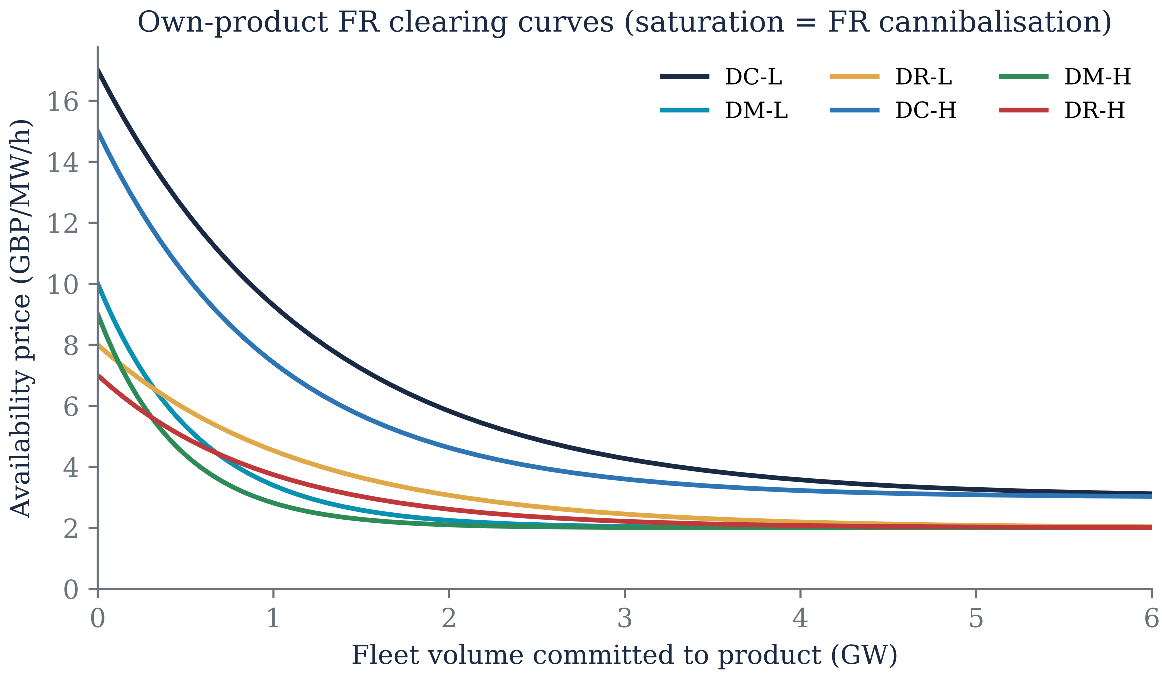

Figure 1: The six stylised FR inverse clearing curves $\mathcal{C}^p(Q)$, each strictly decreasing from an unsaturated cap toward a floor. DC carries the highest availability value and DR the lowest; the Low and High variants differ in level and in sustain/activation. Hand-calibrated to documented GB saturation; not fitted to settlement data.

5.3 Per-stream equilibrium revenue

Table 1 reports the equilibrium per-stream revenue at the representative 6 GW two-hour fleet (Figure 2). Wholesale dominates the stack at £75,263/MW/yr; DC-L is the largest FR stream (£19,935), and the High variants and DR-H carry meaningful value, while DM and DR-L are thin. (DR-H's rank in this base case rests partly on the unsettled activation energy disclosed in Section 5.5: once response energy is settled, DR-H drops out of the ranking.) The seven-stream equilibrium total is £124,666/MW/yr.

Table 1: Equilibrium per-stream revenue, 6 GW two-hour fleet (stylised). All values £/MW/yr. Reported equilibrium FR prices (π^{p*}, £/MW/h) are commitment-weighted means of the half-hourly marginal price $\mathcal{C}^p(b_t^{p*})$ and sit deep in their saturation range: DC-L 3.9, DC-H 4.9, DM-L 4.9, DM-H 4.6, DR-L 5.3, DR-H 2.1.

Stream	**£/MW/yr**	**Stream**	**£/MW/yr**
Wholesale (W)	75,263	DC-H	11,701
DC-L	19,935	DM-H	2,303
DR-H	11,332	DR-L	973
DM-L	3,159	**Total**	**124,666**

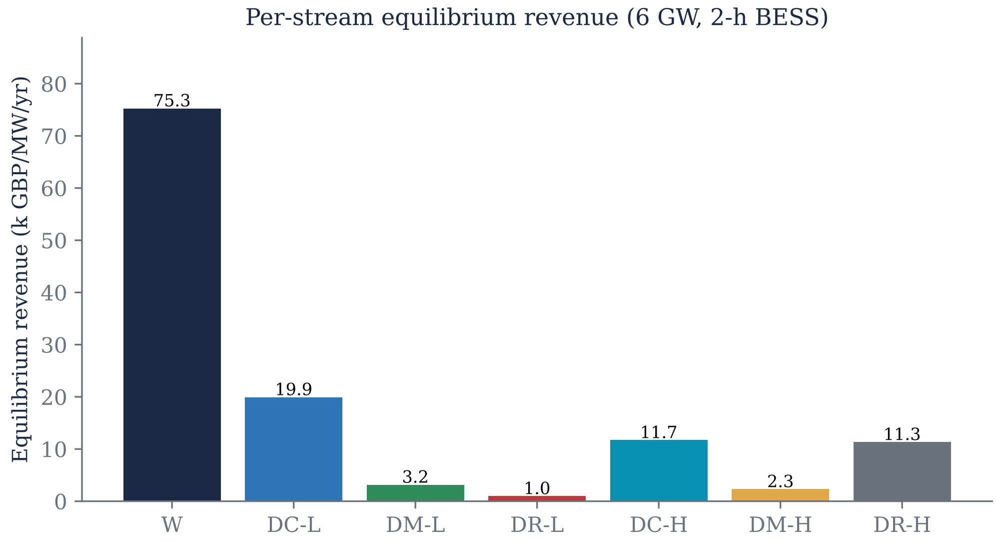

Figure 2: Equilibrium per-stream revenue at 6 GW, two-hour (stylised). Wholesale dominates; DC-L is the largest FR contributor. The naive sum (not shown to scale) attributes full-power un-cannibalised revenue to every stream simultaneously.

5.4 Naive vs equilibrium: the overstatement

Table 2 sweeps fleet size from 4 to 10 GW at two-hour duration. The naive full-stack sum is essentially constant at ≈£689,774/MW/yr — by construction it ignores fleet size, since every stream is priced at zero-fleet levels and committed at full power. The equilibrium total falls with fleet size as both wholesale and FR cannibalise, from £141,774 at 4 GW to £109,802 at 10 GW. The overstatement rises monotonically from 387% to 528%, passing through 453% at the representative 6 GW point (Figure 3).

Table 2: Naive full-stack sum vs seven-stream equilibrium, two-hour fleet (stylised); revenues £/MW/yr. The naive sum is fleet-size-invariant; the equilibrium falls as cannibalisation deepens. Overstatement = $(\Pi^{\text{naive}} - \Pi^*)/\Pi^*$.

Fleet (GW)	**Naive**	**Equilibrium**	**Overstatement**
4	689,774	141,774	+387%
6	689,774	124,666	+453%
8	689,774	115,525	+497%
10	689,774	109,802	+528%

5.5 NESO procurement caps and activation-energy settlement

The base case above deliberately inherits two stylisations from the clearing curves of Section 2.2, and both flatter the FR leg. First, $\mathcal{C}^p$ prices volume but nothing caps it at NESO's procured requirement: at the 6 GW base-case equilibrium the fleet holds 9.7 GW of frequency-response commitment on average across the six products, against a real GB dynamic-response procurement of roughly 1.5–2.5 GW in total. Second, the expected activation energy of eq. (5) shifts the SoC but is not settled: because the High products absorb energy at charge efficiency while the Low products deliver comparatively little, the base-case fleet receives a net 2,513 GWh/yr of unpaid charge (about 209 equivalent full cycles of the 12 GWh fleet), resold into wholesale — an implicit subsidy to the wholesale leg worth about £46k/MW/yr at the coincident equilibrium price.

Table 3 removes both stylisations. Adding the procurement cap $b_t^p \leq R^p$ — with total FR

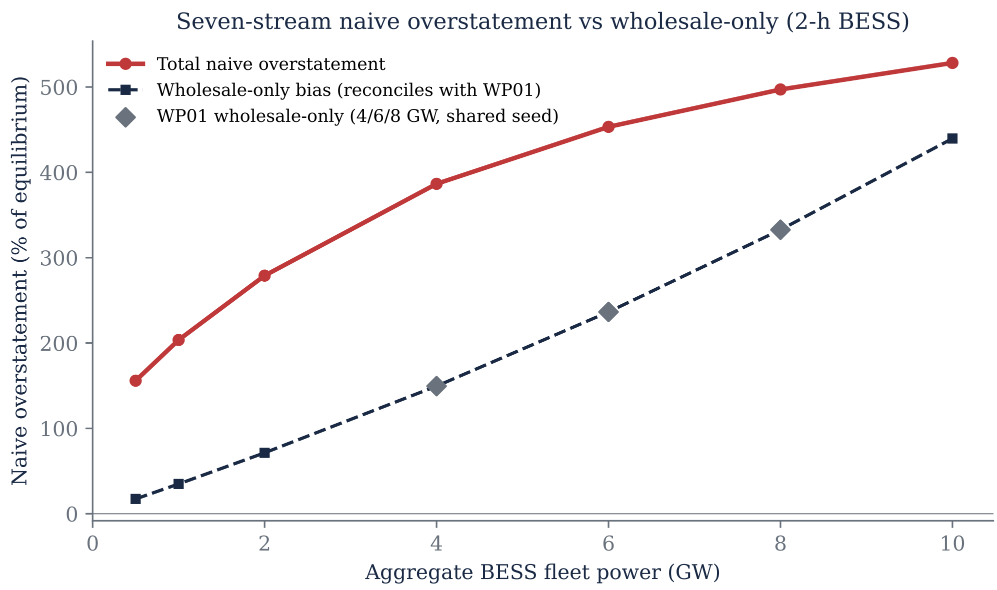

Figure 3: Naive and equilibrium total revenue against fleet size (two-hour, stylised), with the overstatement annotated. The gap widens from 387% at 4 GW to 528% at 10 GW as cannibalisation deepens.

demand of $\{1.5, 2.5, 4\}$ GW allocated across products pro-rata to the requirement scales V_p — collapses the FR leg toward its procured size; settling the activation energy at the wholesale price removes the free-charge subsidy, and DR-H, whose base-case rank rests on that subsidy, drops sharply. The two corrections move the equilibrium *level* down and the naive overstatement *up*: at 6 GW with the central 2.5 GW cap and settled response energy the equilibrium total is £62.8k/MW/yr and the naive sum overstates it by 998%, against 453% in the uncapped base case. The paper's thesis is therefore conservative with respect to these stylisations: a naive full-stack forecast is *more* wrong, not less, once the FR leg is held to NESO-scale volumes and settled energy.

Table 3: Procurement-cap and activation-settlement sensitivity, 6 GW two-hour fleet (stylised); all revenues £/MW/yr. Caps allocate the stated total FR demand pro-rata to V_p; "settled" prices the expected activation energy at the wholesale price. The naive sum (uncapped, uncannibalised) is held fixed at £690k/MW/yr — it is the industry object under audit. Price-taker certificates on the capped arms close to between 2×10^{-8} and 4×10^{-4}.

Scenario	Equilibrium	Wholesale	FR	Overstatement
Base (uncapped, unsettled)	124,666	75,263	49,403	+453%
Caps, 4.0 GW total	74,360	40,238	34,122	+828%
Caps, 2.5 GW total	64,007	37,713	26,294	+978%
Caps, 1.5 GW total	54,644	36,300	18,344	+1,162%
Settled activation energy, uncapped	87,721	29,506	58,214	+686%
Caps 2.5 GW + settled	62,802	31,789	31,012	+998%

A full-year check on the annualisation convention (17,520 half-hours, no scaling) puts the base-case equilibrium at £120.5k/MW/yr against the 26-week-scaled £124.7k/MW/yr — the winter/spring window overstates the level by 3.4% while the overstatement ratio moves only 13 points (from 453% to 466%; the 26-week convention slightly *understates* it) — consistent with the ratio, not the level, being the robust object.

5.6 Decomposing the overstatement

We split the gap into its two mechanisms by inserting an intermediate *price-cannibalised but capacity-unconstrained* forecast (every stream priced at its equilibrium clearing price but still committed at full power, ignoring the shared budget). The step from naive to this intermediate is the price-cannibalisation channel; the step from intermediate to equilibrium is the residual after removing own-product price cannibalisation — a *shared-capacity* step that nets the FR quantity restriction against the wholesale gain from re-optimising the dispatch under the shared budget (at 6 GW: an FR restriction of £74.9k/MW/yr net of a wholesale re-optimisation gain of £25.6k/MW/yr). At 6 GW the split is approximately 91% price cannibalisation and 9% shared-capacity residual, stable across the sweep (Figure 4). Own-product price collapse is by far the dominant effect; the shared budget is a genuinely new but secondary channel that a wholesale-only model cannot see at all.

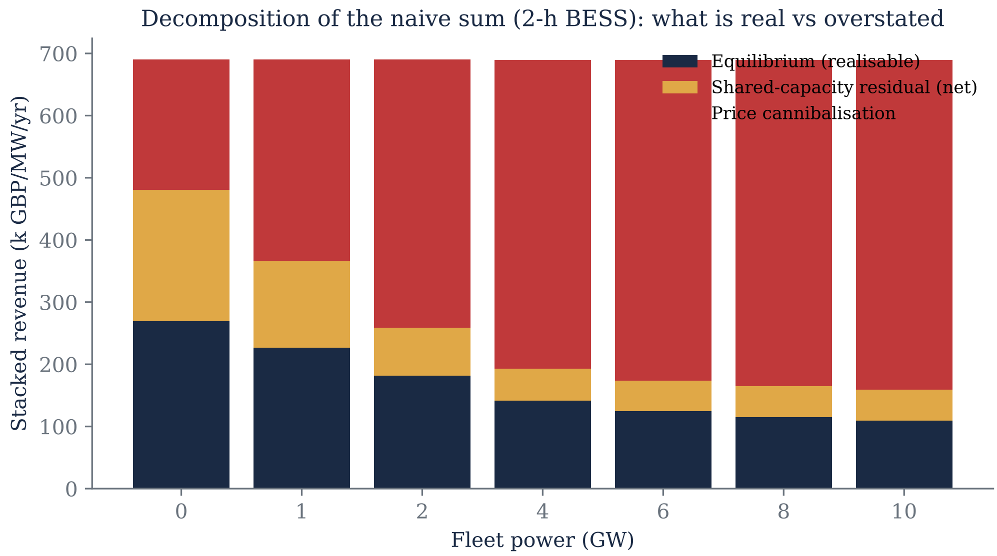

Figure 4: Decomposition of the naive-to-equilibrium gap into own-product price cannibalisation ($\approx 91\%$) and the shared-capacity residual ($\approx 9\%$; the FR quantity restriction net of the wholesale dispatch re-optimisation), two-hour fleet (stylised). The price channel dominates; the capacity channel is the new structural addition relative to a wholesale-only model.

5.7 Reconciliation with CE-WP-2026-01

The seven-stream model contains CE-WP-2026-01 as a special case: pin all FR commitments to zero, set degradation to zero, and the program (7) collapses to the wholesale-only equilibrium with the identical supply curve, demand and renewables series, and seed. Doing so reproduces CE-WP-2026-01's wholesale-only naive bias: 149.5%, 236.5% and 332.7% at 4, 6 and 8 GW respectively, with the wholesale-only equilibrium revenues matching CE-WP-2026-01's to the pound. Two honesty notes on this reconciliation. The wholesale-only special case closes its price-taker certificate to 6.9×10^{-3} rather than the 2.0×10^{-8} of the full stack (the pure-LP price-taker is more degenerate); and the reconciliation uses CE-WP-2026-01's zero-degradation naive convention, while the headline naive of Table 2 applies the base-case κ — a £561/MW/yr difference in the wholesale leg. Neither changes the conclusion: the wholesale leg is inherited unchanged from CE-WP-2026-01.

Crucially, adding the FR stack *raises* the total overstatement above the wholesale-only figure — at 6 GW from 236.5% (wholesale only) to 453% (seven-stream). The reason is precisely Proposition 3.1: the naive sum treats all six FR products as independently committable at full power and un-cannibalised prices, so it books six full reserve revenues from one battery that can deliver at most a shared fraction of them. A battery cannot be in seven places at once. This *strengthens* CE-WP-2026-01's headline rather than competing with it: the true overstatement of a naive full-*stack* forecast is larger, not smaller, than the wholesale-only number, and the shared-capacity budget is a new (if secondary) overstatement channel on top.

5.8 Convergence diagnostic

Figure 5 shows the *unrelaxed* fixed-point residual $\|\mathcal{P}(\mathcal{T}(\boldsymbol{b}_k)) - \mathcal{P}(\boldsymbol{b}_k)\|_\infty$ of the μ-regularised seven-stream operator at the 6 GW two-hour point over a 1,344 half-hour window, for three step schedules (Section 4.5). Plain Picard ($\lambda = 1$) on the regularised map contracts geometrically — median measured ratio 0.30 per iteration — and reaches the 0.02 tolerance in 9 iterations; constant-step KM ($\lambda = 1/2$) reaches it in 11. Harmonic-step KM ($\lambda_k = 2/(k+2)$), the earlier draft's scheme, decays only algebraically because its step vanishes, taking 45 iterations to reach the same tolerance: after 15 iterations its *averaged-iterate* price change — the quantity that draft reported as "residual 0.0174" — sits roughly a factor $1/\lambda_k \approx 8.5$ below its true fixed-point residual, which is still 0.21 at $k = 15$. The honest statement is that the regularised operator contracts and plain Picard certifies the fixed point; the earlier draft's 15-iteration figure was an artefact of its step schedule and residual definition. The convexity-driven price-taker KKT consistency gap at this point is 2.0×10^{-8}. (The earlier draft also asserted "empirical non-expansiveness" with no evidence and a dangling forward-reference to a Section 4 that contained none; both are removed.)

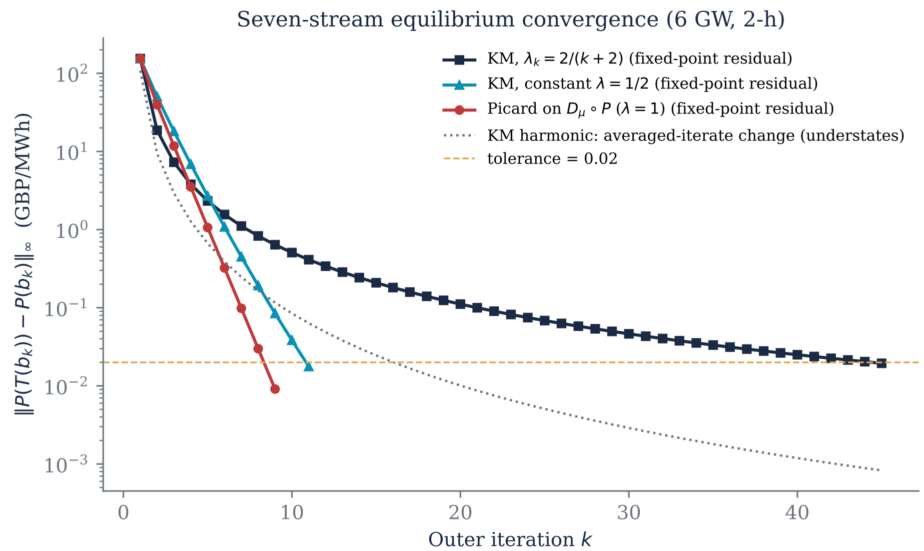

Figure 5: Unrelaxed fixed-point residual of the μ-regularised seven-stream operator (6 GW, two-hour, $\mu = 300$, 1,344 half-hours) under three step schedules. Plain Picard contracts geometrically to sub-tolerance; harmonic-step KM decays algebraically, tracking its vanishing step size, and its averaged-iterate change (dotted) understates the distance to the fixed point.

5.9 Degradation sensitivity

Table 4 and Figure 6 report the equilibrium total as the throughput degradation cost κ varies from 0 to 4 £/MWh. The effect is second-order: the total moves by under 0.5% across the range (£124,176 at $\kappa = 0$ to £124,911 at $\kappa = 4$), because at these fleet sizes the binding constraint is cannibalised price, not cycle cost. The tabulated figure is *gross* revenue, which can tick up slightly with κ as higher cycle cost shifts the fleet away from wholesale cycling and eases wholesale cannibalisation; net of the degradation charge the total falls monotonically, from £124,176 to £115,428. Degradation is included for completeness and auditability, not because it drives the headline.

Table 4: Degradation sensitivity, 6 GW two-hour fleet (stylised); all revenues £/MW/yr, gross of the degradation charge — net of it the total falls monotonically in κ (see text).

κ (£/MWh)	Equilibrium (gross)	Wholesale	FR
0	124,176	74,981	49,194
1	124,456	75,154	49,302
2	124,666	75,263	49,403
4	124,911	75,317	49,595

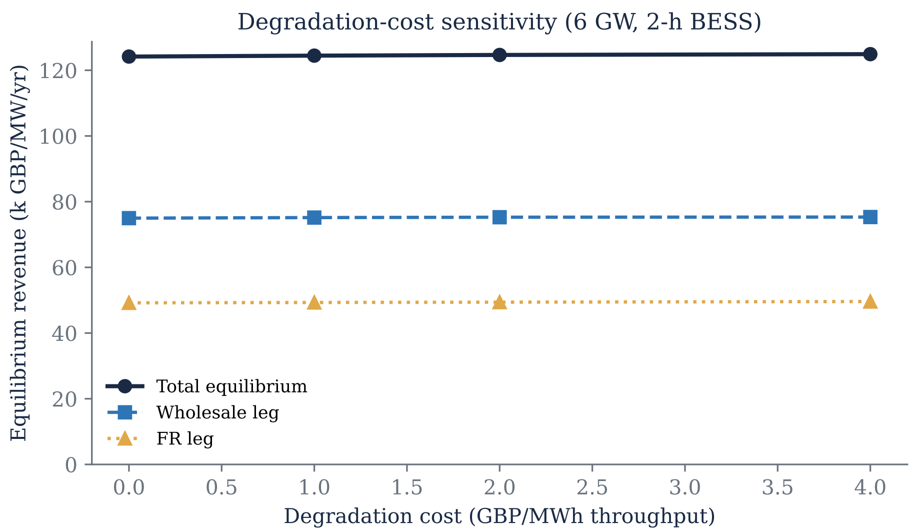

Figure 6: Equilibrium total revenue against throughput degradation cost κ (6 GW, two-hour, stylised). The response is near-flat: at these fleet sizes the binding constraint is cannibalised price, not cycle cost.

5.10 Verified invariants

The reported equilibrium passes every power-systems invariant we check: no half-hour exceeds $P^{\max}$ in either direction (no MW double-counting within a direction); SoC stays within $[0, E^{\max}]$ with both Low and High headroom respected; round-trip efficiency $\eta_c\eta_d = 0.86$ holds; degradation is applied as a marginal cost in the objective; demand and renewables are non-negative; the equilibrium wholesale price stays within the floor/cap; each FR price saturates within its calibrated band; and price equals the constraint dual to a post-reconstruction KKT gap of 2.0×10^{-8} at the 6 GW point (maximum 4.7×10^{-5} across the 0.5–10 GW sweep).

5.11 What this means for project finance

The naive full-stack sum is not a conservative forecast that "leaves headroom" — it is structurally unrealisable, overstating equilibrium revenue by a factor of about 4.9 to 6.3 across the modelled 4–10 GW range (a factor of about 5.5 at 6 GW), and by more still once the FR leg is held to NESO-scale procurement volumes and settled response energy (Section 5.5). We stress that the overstatement *ratio*, not the revenue *level*, is the transferable object of this stylised study: the uncapped base-case level excludes the Balancing Mechanism and Capacity Market yet still sits far above measured whole-stack fleet revenue — the CECadence Quarterly GB BESS Index (trailing year 2025Q3–2026Q2, published July 2026) measures roughly £34k/MW/yr across the 6.4 GW metered fleet *including* BM and CM (Compounding Energy Ltd, 2026) — and Section 5.5 names the stylisations that inflate it. A sponsor or lender should take from this paper the direction and rough magnitude of the naive bias, never a level. The equilibrium forecast, with its FR cannibalisation and shared-capacity coupling, is the investment-grade object; cannibalisation-equilibrium modelling of this kind is the method CECadence is built on — with full seven-stream co-optimisation of this paper's form on the product roadmap rather than in production today — and its revenue models are graded against settled Elexon data in a pre-registered, published pass/fail record, which is where any level claim should be checked.

6 Discussion

6.1 FR cannibalisation: structural or transient?

A natural question is whether the FR price collapse will reverse as NESO expands the requirements. The base-case answer is no. We treat NESO's dynamic-response requirements — set by system inertia and largest-loss considerations — as broadly independent of BESS fleet size; this is a stated modelling assumption, and the procurement-cap sensitivity of Section 5.5 makes the requirement volumes explicit rather than leaving them implicit in the clearing curves. BESS supply, by contrast, scales with the fleet. Even under generous requirement-growth assumptions, supply-driven cannibalisation outpaces demand growth across the modelled range.

6.2 What v1 models — and explicitly what it does not

We state the scope plainly. v1 captures **(a)** own-product price cannibalisation — the wholesale supply curve and the six per-product FR clearing curves — and **(b)** shared-capacity *quantity* coupling — the shared power and SoC budget of Section 3.

v1 does **not** model cross-market price elasticity. Committing MW to DC-L moves the DM-L or wholesale outcome only through the shared *quantity* budget (less power and energy left for the other streams), never through a cross-price term. There is no price Jacobian $\partial \pi^q / \partial Q^p$ for $q \neq p$. Asserting a cross-market price feedback here would be claiming an effect the model does not contain; the full price-stack Jacobian is a v2 extension. We note plainly that this is a misspecification of an *existing* market mechanism, not an unmodelled future one: NESO's Enduring Auction Capability already clears the six products in one co-optimised auction with cross-product substitution (NESO, 2026). The earlier draft claimed "cross-product cannibalisation" as the dominant mechanism — that overstated the model. The honest decomposition is the one in Section 5: dominantly own-product price collapse, plus a secondary shared-capacity channel.

6.3 Strategic vs price-taking behaviour

The framework assumes price-taking at the operator level. In a market with concentrated BESS ownership, strategic withholding could shift the equilibrium toward a Cournot outcome. The GB market remains fragmented enough — many active operators, no single owner dominant; the 2026Q2 index counts 153 metered units across 6.4 GW (Compounding Energy Ltd, 2026) — that price-taking is operationally valid. The framework extends to strategic operators by replacing the welfare program with a Nash best-response system; we leave this for v2.

6.4 Further limitations

The framework assumes the NESO dynamic-suite product definitions as published in the service terms current at August 2026 (NESO, 2026); changes to a response window, call rate, or sustain parameter τ_p would require re-calibrating the affected clearing curve. It also abstracts away the Balancing Mechanism — the second-largest GB BESS revenue stream (Compounding Energy Ltd, 2026) — which is highly state-dependent and outside this model's scope.

7 Conclusion

We have extended the cannibalisation framework of CE-WP-2026-01 from wholesale-only to the full seven-stream GB BESS revenue stack: wholesale energy arbitrage plus the NESO dynamic-response suite (DC, DM, DR, each Low and High). The core contribution is a shared physical budget that couples the streams through one battery's finite power and energy, formalised (Proposition 3.1) so that it provably prevents MW and MWh double-counting. The equilibrium is characterised as the maximiser of a single convex program — unique in the price-relevant decisions — with prices recovered as the constraint marginals and certified by a post-reconstruction price-taker KKT gap of 2.0×10^{-8} at the 6 GW point (at most 4.7×10^{-5} across the sweep); plain Picard iteration on the μ-regularised operator contracts to the fixed point in 9 iterations. In a fully reproducible, explicitly stylised case study, the naive full-stack sum overstates equilibrium revenue by 387–528% across 4–10 GW of two-hour fleet (453% at 6 GW), decomposing into $\approx$ 91% own-product price cannibalisation and $\approx$ 9% shared-capacity residual; holding the FR leg to NESO-scale procurement volumes with settled response energy lowers the equilibrium level and raises the overstatement further. With FR pinned to zero the model reproduces CE-WP-2026-01's wholesale-only bias, and adding the stack *raises* the bias — because a battery cannot be in seven places at once. Investment-grade BESS revenue forecasting at 2030 GB fleet penetration is structurally inseparable from this seven-stream co-optimisation.

Acknowledgements

This work extends CE-WP-2026-01 (Clack, 2026) to the seven-stream setting. The case study is *stylised* and reproducible end-to-end from the Python implementation accompanying this paper, available on request from the author. Reproducibility key: `scenario_id` = `CE-WP-2026-02-stylised-GB-7stream`, `code_version` = `3213bd46`, `timestamp` = `2026-08-29T14:54:43Z`. Comments are welcome to `christopher@compoundingenergy.com`.

References

Christopher T. M. Clack. Cannibalisation as a fixed point: BESS revenue forecasting at storage-scale penetration. Working Paper CE-WP-2026-01, Compounding Energy Ltd, 2026. Compounding Energy Working Paper Series.

Compounding Energy Ltd. CECadence Quarterly GB BESS Index, 2026Q2. `https://www.compoundingenergy.com/gb-bess-index-2026q2.html`, 2026. Published July 2026; trailing-year window 2025Q3–2026Q2.

NESO. Clean Power 2030: Advice on achieving clean power for Great Britain by 2030. National Energy System Operator, 2024.

NESO. Dynamic Containment, Dynamic Moderation and Dynamic Regulation service terms and Enduring Auction Capability market documentation. National Energy System Operator, `https://www.neso.energy/industry-information/balancing-services/frequency-response-services`, 2026. Service terms current at August 2026.

A Notation glossary

Symbol	Meaning
$\mathcal{S}$	Stream set {W, DC-L, DC-H, DM-L, DM-H, DR-L, DR-H} (seven streams)
$\mathcal{L}, \mathcal{H}$	Discharge-direction (Low) and charge-direction (High) FR sets
b_t^p	Fleet commitment to stream p at half-hour t (MW)
$b_t^{\mathrm{W,c}}, b_t^{\mathrm{W,d}}$	Wholesale charge / discharge components (MW); net $b_t^{\mathrm{W}} = b_t^{\mathrm{W,d}} - b_t^{\mathrm{W,c}}$
Q_t^p	Fleet aggregate volume in FR stream p in half-hour t (MW); $Q_t^p = b_t^p$
π_t^{p*}	Equilibrium price of stream p in half-hour t (£/MWh wholesale; £/MW/h FR)
S	Wholesale inverse supply curve (CE-WP-2026-01); $\Phi'_{\mathrm{W}} = S$
$\mathcal{C}^p$	Inverse clearing (saturation) curve for FR stream p
τ_p	FR sustain duration (h): 0.25 DC (published 15-minute requirement), 0.5 DM, 1.0 DR (stylised)
ρ_p	FR activation fraction: 0.04 DC, 0.06 DM, 0.12 DR
κ	Throughput degradation cost (£/MWh); base case 2
η_c, η_d	Charge / discharge efficiencies, $\eta_c\eta_d = \eta = 0.86$
$E^{\max}, P^{\max}$	Fleet energy and power capacities; $E^{\max} = h\, P^{\max}$
σ_0	Initial / terminal state-of-charge fraction ($e_T \geq \sigma_0 E^{\max}$)
μ	KM regularisation parameter (convergence diagnostic only); 300