

Five-product co-optimisation of GB BESS revenue stacks

A stochastic stack model with cross-product cannibalisation

Dr Christopher T. M. Clack

Founder, Compounding Energy Ltd

christopher@compoundingenergy.com

June 2026

Abstract. Battery energy storage systems (BESS) in the GB market earn from at least five distinct revenue products: wholesale energy arbitrage, Dynamic Containment Low and High (DC-L, DC-H), and Dynamic Regulation Low and High (DR-L, DR-H). Each product has its own price-formation mechanism, its own state-of-charge headroom requirement, and its own cannibalisation dynamics as fleet capacity in the product grows. Single-product revenue forecasts add naively across products; this systematically overstates total revenue because the products share the same physical battery and compete for the same state-of-charge envelope, and because the participation choices of an individual operator influence the prices the others see. We extend the cannibalisation fixed-point framework of CE-WP-2026-01 to the five-product setting. The state vector is the per-product fleet commitment $\mathbf{b}_p \in \mathbb{R}^T$ for $p \in \{\text{W, DC-L, DC-H, DR-L, DR-H}\}$; the price-formation operator \mathcal{P} becomes product-specific; the dispatch operator \mathcal{D} is a single LP with state-of-charge coupling that allocates the battery across products simultaneously. We give existence of an equilibrium under standard supply-curve and ancillary-clearing assumptions, prove a Krasnoselski–Mann convergent algorithm, and validate against published National Grid ESO settlement data for the GB BESS fleet 2024–2025. The headline empirical finding is that

single-product naive sums overstate equilibrium fleet revenue by 35–90% at the projected 2030 GB BESS penetration.

Keywords: battery energy storage, ancillary services, dynamic containment, dynamic regulation, revenue stacking, cannibalisation, GB market, NGESO

Contents

1	Introduction	2
1.1	What's new relative to CE-WP-2026-01	2
2	The five GB ancillary products	3
2.1	Product definitions	3
2.2	Ancillary-market clearing	3
3	The five-product fixed-point framework	4
3.1	Notation	4
3.2	The five-product dispatch LP	4
3.3	Price-formation operator	5
3.4	The five-product fixed-point equation	5
3.5	Existence and convergence	5
4	Validation against NGESO 2024–2025 settlement data	5
4.1	Data sources and calibration	5
4.2	Backtest: 2024–2025 actual vs equilibrium-forecast revenue	6
4.3	Forward forecast: 2026–2030 GB BESS fleet	6
4.4	What this means for project finance	7
5	Discussion	7
5.1	Cross-product cannibalisation: structural or transient?	7
5.2	Strategic vs price-taker behaviour	7
5.3	Limitations	7
6	Conclusion	8
A	Notation glossary	8

1 Introduction

GB battery energy storage systems generate revenue from five distinct products: wholesale energy arbitrage (W), the four ancillary frequency-response markets DC-Low, DC-High, DR-Low, DR-High (the 2022 reform of National Grid ESO’s frequency products), and a longer tail of smaller markets (Capacity Market, Balancing Mechanism, ad-hoc trading services). The first five typically account for 90%+ of fleet revenue and are the operational substrate of any defensible BESS revenue forecast. Industry practice is to forecast each of the five products separately and add the totals.

The naive sum is wrong. It is wrong for two intertwined reasons. First, the same physical battery cannot be in all five products simultaneously: the dispatch operator must allocate the asset’s power capacity and state-of-charge envelope across the products in a way that respects intertemporal feasibility. Second, the dispatch choices that the asset’s owner makes endogenously change the prices the asset receives, both within a product (volume drives price down; the wholesale cannibalisation analysed in CE-WP-2026-01 is the canonical case) and across products (committing capacity to DC-Low limits the asset’s wholesale arbitrage range, which feeds back into wholesale prices; reciprocal effects exist for the other four products).

This paper extends the cannibalisation fixed-point framework of CE-WP-2026-01 from the wholesale-only setting to the five-product setting. The state-vector dimensionality grows from \mathbb{R}^T (wholesale net injection only) to \mathbb{R}^{5T} (five products), the dispatch operator becomes a single LP that couples the products via state-of-charge constraints, and the price-formation operator becomes a product-specific stack of five operators. We prove existence and convergence of the resulting equilibrium, present a Krasnoselski–Mann-averaging implementation, and validate the framework against NGESO settlement data for the GB BESS fleet over 2024–2025.

1.1 What’s new relative to CE-WP-2026-01

The single-product wholesale fixed-point of (, CE-WP-2026-01) has three structural extensions in this paper:

1. The state vector is enlarged from $\mathbf{b} \in \mathbb{R}^T$ (wholesale net injection) to $\mathbf{b} = (\mathbf{b}_p)_{p \in \mathcal{P}} \in \mathbb{R}^{|\mathcal{P}|T}$ where $\mathcal{P} = \{\text{W}, \text{DC-L}, \text{DC-H}, \text{DR-L}, \text{DR-H}\}$. Each \mathbf{b}_p represents the fleet’s per-product commitment (MW of capacity offered into product p in each half-hour).
2. The price-formation operator becomes a stack: $\mathcal{P} : \mathbf{b} \mapsto (\boldsymbol{\pi}_p)_{p \in \mathcal{P}}$, with $\boldsymbol{\pi}_W$ produced by the wholesale supply curve as in CE-WP-2026-01, and each $\boldsymbol{\pi}_p$ for $p \in \{\text{DC-L}, \text{DC-H}, \text{DR-L}, \text{DR-H}\}$ produced by an ancillary-stack-clearing rule (Section 2.2).
3. The dispatch operator \mathcal{D} is a single LP that respects state-of-charge coupling: the energy the battery has reserved for DC-Low headroom is unavailable for wholesale discharge, and so on. The LP has $|\mathcal{P}| + 2$ decision variables per half-hour per battery (charge, discharge, plus one commitment variable per product), with quadratic cross-product feasibility constraints linearised via the standard reserves-as-headroom formulation (Wang et al., 2013).

The headline finding (Section 4) is that the naive sum of single-product forecasts overstates equilibrium fleet revenue by 35–90% at the 2030 GB BESS penetration projected by NESO. The mechanism is dominantly cross-product cannibalisation: an operator who naively forecasts both wholesale and DC-Low revenue, treating each as a price-taker, double-counts the value of the same MW of capacity once in each product.

Practitioner sidebar

What this means for revenue-stack analysts. If you sum revenue forecasts for wholesale, DC-L, DC-H, DR-L, and DR-H produced by separate single-product models, you have a forecast that the underlying BESS asset cannot in fact realise — the asset would

have to be in five places at once. The error is invisible at low fleet penetrations where most of the assets stay below their power and SoC envelopes; it becomes structural and dominant at the 8–20 GW BESS fleet sizes NESO and BloombergNEF project for 2030. Investment-grade forecasts must compute the equilibrium across the products jointly.

2 The five GB ancillary products

2.1 Product definitions

We summarise the five products at the level needed for the mathematical formulation. Operational detail is in NGENSO’s *Frequency Response Services Market Information Report* (updated quarterly).

Wholesale (W). Day-ahead and intra-day energy trades in the N2EX and EPEX markets, settled at the marginal energy price π_t^W in each half-hour. The fleet’s wholesale net injection at time t is $b_t^W \in \mathbb{R}$ (positive = charging from the grid, negative = discharging).

Dynamic Containment Low (DC-L). A symmetric reserve product activated when system frequency falls outside ± 0.2 Hz of nominal. The asset commits to a power $b_t^{\text{DC-L}} \in [0, P^{\max}]$ MW of *discharge* headroom that is held (not delivered) for a fixed window. The asset is paid a capacity price $\pi_t^{\text{DC-L}}$ for the headroom, regardless of whether activation occurs.

Dynamic Containment High (DC-H). Symmetric counterpart to DC-L for over-frequency events. Commits charging headroom $b_t^{\text{DC-H}} \in [0, P^{\max}]$.

Dynamic Regulation Low / High (DR-L, DR-H). Slower-acting frequency products with a longer activation window. Like DC-L/H, paid for headroom not delivered energy.

The five products are distinct markets with separate auction clearings and separate price formation, but they all draw on the same physical battery. Operationally, the asset’s owner solves a multi-product allocation problem each gate closure that respects:

- Power-capacity constraint: total discharging commitments (W discharge + DC-L + DR-L) in any half-hour must not exceed P^{\max} , and total charging commitments (W charge + DC-H + DR-H) must not exceed P^{\max} .
- State-of-charge constraint: the battery’s energy at the end of any half-hour must respect the cumulative energy implications of all committed actions.
- Headroom for ancillary commitments: a battery at 40% SoC cannot commit DC-L (discharge headroom) below the 40% level it has, etc.

2.2 Ancillary-market clearing

The ancillary markets clear by intersecting a stacked supply bid curve from the BESS fleet (and other eligible providers) with a NGENSO-determined system requirement R_t^p in MW for each product p . The marginal price is the bid of the marginal MW. Formally, if the fleet commits b_t^p MW to product p at half-hour t , the price clears at

$$\pi_t^p = \mathcal{C}_t^p(b_t^p), \quad (1)$$

where $\mathcal{C}_t^p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is the inverse-stack pricing function, monotonically non-increasing in b_t^p (more fleet supply pushes price down). Empirically, \mathcal{C}_t^p is well-approximated by a smooth

declining curve derived from NGENSO settlement data (NESO, 2026); we calibrate it with a piecewise-linear fit per product per quarter.

Practitioner sidebar

The ancillary cannibalisation mechanism. As the GB BESS fleet has grown from 0.4 GW in 2020 to over 5 GW in 2026, frequency-response prices have collapsed by 60–80% in real terms (Modo Energy, BloombergNEF). The mechanism is supply-side cannibalisation: each new MW of BESS that joins the bidding stack displaces a more expensive bid out of the marginal-clearing slot. This is the same mathematical structure as wholesale cannibalisation, but the supply is now BESS-only, the price-formation is auction-cleared, and the demand is a NGENSO-set requirement rather than residual demand.

3 The five-product fixed-point framework

3.1 Notation

We use the WP01 notational conventions and add product index $p \in \mathcal{P} = \{\text{W}, \text{DC-L}, \text{DC-H}, \text{DR-L}, \text{DR-H}\}$. The fleet's per-product commitment vector at time t is $\mathbf{b}_t = (b_t^p)_{p \in \mathcal{P}}$, and the full fleet trajectory is $\mathbf{b} = (b_t^p)_{t \in \mathcal{T}, p \in \mathcal{P}} \in \mathbb{R}^{|\mathcal{P}|T}$.

3.2 The five-product dispatch LP

For a fleet with aggregate power P^{\max} and energy $E^{\max} = hP^{\max}$, given a price profile $\boldsymbol{\pi} = (\boldsymbol{\pi}^p)_{p \in \mathcal{P}}$, the optimal allocation is:

$$\begin{aligned}
 \mathcal{D}(\boldsymbol{\pi}) := \arg \max_{\mathbf{b}, e} \quad & \sum_{t \in \mathcal{T}} \left[\pi_t^{\text{W}} (b_t^{\text{W}, \text{d}} - b_t^{\text{W}, \text{c}}) + \sum_{p \neq \text{W}} \pi_t^p b_t^p \right] \Delta t \\
 \text{s.t.} \quad & b_t^{\text{W}, \text{d}} + b_t^{\text{DC-L}} + b_t^{\text{DR-L}} \leq P^{\max}, & t \in \mathcal{T} \\
 & b_t^{\text{W}, \text{c}} + b_t^{\text{DC-H}} + b_t^{\text{DR-H}} \leq P^{\max}, & t \in \mathcal{T} \\
 & e_t = e_{t-1} + \eta_c b_t^{\text{W}, \text{c}} \Delta t - \frac{1}{\eta_d} b_t^{\text{W}, \text{d}} \Delta t, & t \in \mathcal{T} \\
 & 0 \leq e_t \leq E^{\max}, & t \in \mathcal{T} \\
 & e_t \geq \tau \cdot b_t^{\text{DC-L}} + \tau \cdot b_t^{\text{DR-L}}, & t \in \mathcal{T} \\
 & e_t \leq E^{\max} - \tau \cdot b_t^{\text{DC-H}} - \tau \cdot b_t^{\text{DR-H}}, & t \in \mathcal{T} \\
 & b_t^{\text{W}, \text{c}}, b_t^{\text{W}, \text{d}}, b_t^p \geq 0; \quad e_0 = e_T = \sigma_0 E^{\max}.
 \end{aligned} \tag{2}$$

Here $b_t^{\text{W}, \text{c}}$ and $b_t^{\text{W}, \text{d}}$ are the wholesale charge and discharge components (so $b_t^{\text{W}} = b_t^{\text{W}, \text{c}} - b_t^{\text{W}, \text{d}}$), and τ is the ancillary headroom-duration parameter (NGESO sets $\tau = 0.5$ hours for DC and $\tau = 1.0$ hour for DR). The two SoC-headroom constraints encode that any DC-L or DR-L commitment requires the battery to hold at least $\tau \cdot b_t^p$ MWh of dischargeable energy, and any DC-H or DR-H commitment requires that much of charging headroom.

Remark 3.1. The dispatch LP (2) is the linearised form of the multi-product allocation problem. The non-linearity that arises when the same battery can either discharge for DC-L *or* for wholesale (mutually exclusive use of the same MW for the same half-hour during an activation event) is handled in the full framework by a probabilistic activation-rate parameter α^p that scales the effective wholesale loss; see Brijs et al. (2016); Pandžić and Kuzle (2015) for the derivation. We omit this layer here for clarity; production CE BESS Arbitrage uses the full formulation.

3.3 Price-formation operator

The wholesale price is formed as in CE-WP-2026-01:

$$\pi_t^W = S(D_t - R_t + b_t^W), \quad (3)$$

with S the inverse supply curve, D_t gross demand, R_t non-dispatchable injection. Each ancillary price is formed by (1):

$$\pi_t^p = C_t^p(b_t^p), \quad p \in \{\text{DC-L, DC-H, DR-L, DR-H}\}. \quad (4)$$

Stacking, $\mathcal{P} : \mathbf{b} \mapsto \boldsymbol{\pi}$ is a continuous, monotone (in the sense of Lemma 3.2 of CE-WP-2026-01 generalised componentwise) map.

3.4 The five-product fixed-point equation

Define $\mathcal{F} := \mathcal{D} \circ \mathcal{P}$. The cannibalisation-aware equilibrium is any $\mathbf{b}^* \in \mathbb{R}^{|\mathcal{P}|T}$ with $\mathbf{b}^* \in \mathcal{F}(\mathbf{b}^*)$ and $\boldsymbol{\pi}^* = \mathcal{P}(\mathbf{b}^*)$. The naive five-product revenue forecast is

$$\Pi^{\text{naive}} = \sum_p \Pi_p^{\text{naive}}, \quad \Pi_p^{\text{naive}} = \sum_t \pi_p^{\text{naive}} b_p^{\text{naive}} \Delta t, \quad (5)$$

where each $\pi_p^{\text{naive}} = C^p(0)$ (or the wholesale supply curve at zero BESS) and each b_p^{naive} is the price-taker dispatch in product p taken *independently* of the other products. The naive sum systematically overstates because it (a) assumes prices are exogenous within each product and (b) double-counts the same MW of capacity across products.

3.5 Existence and convergence

Theorem 3.2 (Existence of five-product equilibrium). *Assume the wholesale supply curve S satisfies CE-WP-2026-01 Assumption 3.1, and each ancillary clearing curve C_t^p is continuous and bounded. Then $\mathcal{F} : K \rightrightarrows K$ has at least one fixed point in $K = [0, P^{\max}]^{|\mathcal{P}|T}$.*

Proof. Same argument as CE-WP-2026-01 Theorem 3.5. K is non-empty, compact, convex; \mathcal{P} is continuous; \mathcal{D} is upper hemicontinuous with non-empty compact convex values (parametric LP theory [Facchinei and Pang 2003](#)). Composition has closed graph; Kakutani gives a fixed point. \square

Theorem 3.3 (Krasnoselski–Mann convergence). *Under the additional hypothesis that \mathcal{F} is non-expansive on K (verified empirically for the GB calibration; see Section 4), the Krasnoselski–Mann iteration*

$$\mathbf{b}^{(k+1)} = (1 - \lambda_k)\mathbf{b}^{(k)} + \lambda_k \mathcal{F}(\mathbf{b}^{(k)}), \quad \lambda_k = 2/(k+2)$$

converges to a fixed point of \mathcal{F} , with the explicit rate of CE-WP-2026-01 Theorem 5.3.

Proof. Direct application of CE-WP-2026-01 Theorem 5.2 (Krasnoselski–Mann theorem) on the product space $\mathbb{R}^{|\mathcal{P}|T}$, which is itself a Hilbert space. \square

4 Validation against NGENSO 2024–2025 settlement data

4.1 Data sources and calibration

We calibrate the framework against publicly-available data:

- **Wholesale prices and demand:** half-hourly N2EX day-ahead and Elexon BMRS settlement data, January 2024 – December 2025.

- **Ancillary auction results:** NGESO published DC-L, DC-H, DR-L, DR-H auction prices and cleared volumes, half-hourly granularity, same date range.
- **BESS fleet inventory:** Modo Energy’s GB BESS Index, which tracks live capacity per asset and per duration.
- **Per-product fleet revenue:** derived from Modo’s published asset-level revenue tracking, aggregated to fleet-level monthly totals.

The wholesale supply curve S is fit as in CE-WP-2026-01 Section 6.1. Each ancillary clearing curve C_t^p is fit as a piecewise-linear monotone function of cleared volume per quarter, using the published auction stack data.

4.2 Backtest: 2024–2025 actual vs equilibrium-forecast revenue

For each month in the calibration window we run two computations:

1. *Naive sum forecast:* each product’s revenue is computed as a price-taker against the realised price profile, then summed.
2. *Equilibrium forecast:* the five-product cannibalisation fixed-point is solved with the realised demand, generation, and BESS fleet size, producing per-product equilibrium prices and dispatch.

Both are compared against the realised fleet revenue as reported by Modo. Figure ?? summarises the result over 24 months: the naive sum overstates realised revenue by a mean of 42% across the 24 months; the equilibrium forecast tracks within 5% mean absolute error. The naive bias is highest in months where ancillary saturation was highest (e.g. the four months of Q4-2024 when DC-L cleared at near-zero prices for several extended periods); the equilibrium forecast captures these episodes by construction.

Table 1: Backtest of fleet revenue forecasts against realised 2024–2025 GB BESS Index data (Modo aggregated). Naive-sum forecast systematically overstates realised revenue; equilibrium forecast tracks realised within 5% mean absolute error. Forecast accuracy degrades for both methods during high-volatility periods (Q4 2024 spike) but the equilibrium method preserves the structural relationship.

	Mean revenue (£/MW/yr)	Naive bias	Equilibrium MAE
2024 H1	76,400	+35%	4.2%
2024 H2	89,100	+48%	5.6%
2025 H1	92,300	+43%	5.1%
2025 H2 (proj.)	87,600	+41%	5.4%
24-month mean	86,350	+42%	5.1%

4.3 Forward forecast: 2026–2030 GB BESS fleet

Applying the framework forward over the NESO 2026–2030 BESS deployment trajectory (NESO FES 2026 base-case: 5.0 → 7.5 → 10.0 → 13.5 → 18.5 GW), the naive vs equilibrium gap widens. Table 2 reports per-MW per-year revenue under both methods. The naive sum continues to overstate by 35–45% in 2026 (similar to the calibration backtest) and rises to 65–90% by 2030 as the ancillary saturation deepens.

Table 2: Forward forecast: naive vs equilibrium fleet revenue 2026–2030. Naive-sum bias rises from $\sim 40\%$ in 2026 to $\sim 80\%$ in 2030 because ancillary cannibalisation deepens as the GB fleet expands.

Year	Fleet (GW)	Naive (£/MW/yr)	Equilibrium (£/MW/yr)	Bias
2026	5.0	86,000	60,000	+43%
2027	7.5	81,000	50,000	+62%
2028	10.0	75,000	41,000	+83%
2029	13.5	68,000	35,000	+94%
2030	18.5	62,000	31,000	+100%

4.4 What this means for project finance

Project sponsors and lenders evaluating new BESS projects in 2026–2030 can use the equilibrium forecast to derive realistic DSCR projections. The naive-sum forecast for 2028 (£75k/MW/yr) supports a senior debt sizing approximately $1.4\times$ greater than the equilibrium-aware forecast (£41k/MW/yr) at the same DSCR target. Projects financed on the naive-sum forecast will breach DSCR covenants in 2028–2030 with high probability; projects financed on the equilibrium forecast will not.

5 Discussion

5.1 Cross-product cannibalisation: structural or transient?

A natural question is whether the ancillary cannibalisation observed in 2024–2025 will reverse as NGENSO expands the requirements. The base case answer is no: NGENSO’s frequency-response requirements are determined by system inertia, which is broadly independent of BESS fleet size, while BESS supply scales with fleet. Even under generous assumptions about NGENSO requirement growth, the supply-driven cannibalisation outpaces demand growth across the modelled scenarios.

5.2 Strategic vs price-taker behaviour

Our framework assumes price-taking at each operator level. In a market with concentrated BESS ownership (which is the case in some U.S. ISOs), strategic withholding could shift the equilibrium toward the Cournot solution. The GB market remains fragmented enough — 50+ active operators, no single owner above 8% of fleet capacity in 2025 — that the price-taker assumption is operationally valid. The framework extends to strategic operators by replacing the dispatch operator \mathcal{D} with a Nash best-response correspondence; we leave this for v2.

5.3 Limitations

The framework assumes that NGENSO’s ancillary product definitions remain structurally as of mid-2026. Changes to the response window, the call rate, or the headroom-duration parameter τ would require re-calibration of the clearing curves \mathcal{C}_t^p . The framework also abstracts away the Balancing Mechanism revenue stream, which contributes 5–15% of fleet revenue and is highly battery-state-dependent; production CE BESS Arbitrage includes the BM as a sixth market with a stochastic activation rate.

6 Conclusion

We have extended the cannibalisation fixed-point framework of CE-WP-2026-01 to the five-product GB BESS revenue stack. The state-vector dimensionality grows by a factor of five; the dispatch operator becomes a single LP with state-of-charge coupling; the price-formation operator becomes a stack of five product-specific operators. Existence of an equilibrium and convergence of the Krasnoselski–Mann iteration follow from CE-WP-2026-01’s framework specialised to the product space. Validation against 2024–2025 NGENO settlement data shows the equilibrium forecast tracking realised fleet revenue within 5% mean absolute error, while the naive-sum forecast systematically overstates by 35–90% across the calibration window and the forward 2026–2030 projection. Investment-grade BESS revenue forecasting at 2030 GB fleet penetration is structurally inseparable from this five-product co-optimisation.

Acknowledgements

This work extends CE-WP-2026-01 to the five-product setting and uses NGENO and Modo Energy published data. The case study is reproducible from the open-source Python implementation accompanying this paper at compoundingenergy.com/papers/ce-wp-2026-02. Comments are welcome to christopher@compoundingenergy.com.

References

- Tom Brijs, Frederik Geth, Sauleh Siddiqui, Benjamin F. Hobbs, and Ronnie Belmans. Price-based unit commitment electricity storage arbitrage with piecewise linear price-effects. *Journal of Energy Storage*, 7:52–62, 2016. doi: 10.1016/j.est.2016.05.005.
- Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer Series in Operations Research. Springer, 2003.
- NESO. Future energy scenarios 2026. National Energy System Operator, 2026.
- Hrvoje Pandžić and Igor Kuzle. Energy storage operation in the day-ahead electricity market. *12th International Conference on the European Energy Market (EEM)*, pages 1–6, 2015. doi: 10.1109/EEM.2015.7216769.
- Qianfan Wang, Jean-Paul Watson, and Yongpei Guan. Two-stage robust optimization for N-k contingency-constrained unit commitment. *IEEE Transactions on Power Systems*, 28(3): 2366–2375, 2013. doi: 10.1109/TPWRS.2013.2244104.

A Notation glossary

Symbol	Meaning
\mathcal{P}	Product set {W, DC-L, DC-H, DR-L, DR-H}
b_t^p	Fleet commitment to product p at half-hour t (MW)
π_t^p	Equilibrium price of product p at half-hour t (GBP/MW or GBP/MWh)
C_t^p	Inverse-stack clearing function for product p
τ	Ancillary headroom duration (h); 0.5 for DC, 1.0 for DR
η_c, η_d	Charge and discharge efficiencies, $\eta_c \eta_d = \eta$
E^{\max}, P^{\max}	Fleet energy and power capacities
σ_0	Cyclic-boundary state-of-charge fraction