



COMPOUNDING ENERGY LTD

Working Paper Series • CE-WP-2026-01

Cannibalisation as a fixed point

revenue forecasting for battery energy storage in liberalised electricity markets

Dr Christopher T. M. Clack

Founder, Compounding Energy Ltd

christopher@compoundingenergy.com

September 2026

Abstract. Battery energy storage system (BESS) revenue forecasts that take wholesale prices as exogenous overstate revenue substantially in any market where the storage fleet is large enough to influence price formation. We frame the BESS revenue equilibrium as a fixed point of the composition of two operators — a price-formation operator $\mathcal{P}$ and a profit-maximising dispatch operator $\mathcal{D}$ — but observe that $\mathcal{D}$ is the solution map of a linear program and is therefore discontinuous, so the composition is not non-expansive, admits no continuous single-valued fixed-point map, and naive iteration on it does not converge. We instead characterise the competitive storage equilibrium as the unique minimiser of a strictly convex welfare program, with equilibrium price equal to the dual of power balance; existence and uniqueness follow from strict convexity rather than from any contraction property. The naive relaxed Picard iteration on the raw dispatch operator demonstrably fails to converge at the tolerance required for investment-grade revenue, at every fleet size, and its irreducible residual grows with fleet penetration; a μ -regularised Krasnoselski–Mann scheme, for which the regularised dispatch is a genuine contraction, converges at every fleet size in a small number of iterations. A stylised single-zone Great Britain case study

calibrated to a historical GB 2024–2025 day-ahead reference window (retained for comparability; as of August 2026), with headline revenues computed over a 26-week horizon of 8,736 half-hours scaled to annual, quantifies the bias of the naive (price-taker) revenue forecast: for two-hour batteries the naive forecast overstates revenue by 150% at 4 GW, 236% at 6 GW, 333% at 8 GW, and 440% at 10 GW; the public “100–300%+” range corresponds to the 4–8 GW representative band. The bias is monotone in fleet size, and monotone in storage duration for fleets of 2 GW and above. We discuss extensions to nodal markets, ancillary stacking, and risk-averse formulations, and locate the algorithm in the Compounding Energy stack as the analytical core of CECadence (wholesale and stacked BESS revenue, in private beta), with a nodal extension planned.

Keywords: battery energy storage, cannibalisation, fixed-point iteration, variational inequality, electricity markets, GB, Krasnoselski–Mann iteration

Contents

1	Introduction	2
1.1	Cannibalisation: the term and its history	3
1.2	Roadmap	3
2	The BESS revenue forecasting problem	4
2.1	Market structure and notation	4
2.2	The BESS dispatch problem	4
2.3	The fixed-point equation	5
2.4	The naive forecast as the zeroth Picard iterate	6
3	The fixed-point framework	6
3.1	Operator properties	6
3.2	Existence	7
3.3	Variational inequality formulation	8
3.4	Connection to Cournot equilibrium	8
4	Existence, uniqueness, and structure	8
4.1	The competitive equilibrium as a convex welfare program	9
5	Convergence: where naive Picard fails and KM averaging succeeds	11
5.1	Plain Picard iteration	11
5.2	Anderson acceleration	12
5.3	μ -regularised Krasnoselski–Mann averaging	12
5.4	Algorithmic summary	13
6	Stylised GB case study	14
6.1	Setup	14
6.2	Bias of the naive forecast	16
6.3	Equilibrium price duration curve	18
6.4	Convergence behaviour of iteration schemes	19
6.5	Translation to project finance	20
7	Discussion and extensions	21
7.1	What the stylised model omits	21
7.2	Strategic vs price-taker storage	22
7.3	Capacity expansion: the meta-fixed-point	22
8	Conclusion	23
A	Notation glossary	26
B	Reproducibility	26

1 Introduction

Battery energy storage systems (BESS) are being deployed at unprecedented rates in liberalised electricity markets. The Great Britain (GB) market alone has grown from a historical baseline of 0.4 GW of installed BESS in 2020 to approximately 5 GW as of August 2026 (Modo Energy, 2026), with NESO’s Clean Power 2030 analysis projecting on the order of 25 GW of installed GB battery capacity by 2030 (central range 23–27 GW) (NESO, 2024, 2026). Each project requires a revenue forecast on which capital allocation is decided: project finance terms, equity returns, and merchant exposure are all calibrated against the modelled revenue stream over the asset’s 15–20 year economic life.

The dominant revenue forecast methodologies in commercial use today — regression on historic price spreads, single-asset dispatch against an exogenous price profile, or combinations thereof — have a structural defect. They treat the wholesale price profile as an input to the BESS dispatch problem, when in fact the price profile is a function of aggregate BESS dispatch. As the storage fleet grows, the profile that an individual asset can capture is increasingly distorted by the activity of the rest of the fleet. The forecasts ignore this feedback. The error is small for early fleets, large for projected fleets, and dominant for the equilibrium that any new project’s revenue actually settles around over its lifetime.

This paper makes four contributions. First, we formalise the BESS revenue equilibrium as a fixed point of two operators acting on the fleet-level dispatch profile, and identify the structural obstruction to solving it pointwise: the dispatch operator $\mathcal{D}$ is the solution map of a linear program and is therefore discontinuous, so the composition $\mathcal{D} \circ \mathcal{P}$ is not non-expansive, admits no continuous single-valued fixed-point map, and defeats naive iteration (Section 3). Second, we resolve this by characterising the competitive storage equilibrium as the unique minimiser of a strictly convex welfare program: existence and uniqueness follow from strict convexity of the welfare potential, and the equilibrium price is recovered as the dual of power balance (Section 4). These conditions are physically plausible and satisfied in the GB market. Third, we exhibit the failure of naive relaxed Picard iteration on the raw dispatch operator — it does not converge at investment-grade tolerance at *any* fleet size, and its irreducible residual grows with fleet — and show that a μ -regularised Krasnoselski–Mann (KM) scheme, whose regularised dispatch is a genuine contraction, converges at every fleet size (Section 5). Fourth, we quantify the bias of the naive forecast against the equilibrium forecast in a stylised single-zone GB market calibrated to a historical 2024–2025 day-ahead reference window, retained for comparability and compared against the twelve months to August 2026 in Section 6. The headline numerical result is that for two-hour batteries the naive forecast overstates revenue by 150% at a 4 GW fleet, 236% at 6 GW, 333% at 8 GW, and 440% at 10 GW; the widely quoted “100–300%+” range — which our own analysis reproduces — corresponds to the 4–8 GW representative band (see also the subscription analyses of Modo Energy 2026; BloombergNEF 2026). The bias is monotone in fleet size, and monotone in storage duration for fleets of 2 GW and above (below 2 GW the ordering reverses and longer duration is mildly protective).

We deliberately strip away features that confound the cannibalisation effect: there is one zone, no transmission, no ancillary stack, perfect foresight, and homogenous fleet composition. These omissions are addressed in Section 7 but the cannibalisation effect itself is qualitatively robust to them. The mathematics — and the bias — survive any layering of additional structure. Our point is that the structure currently absent from commercial forecasts is what determines whether revenue projections are defensible at investment grade.

Practitioner sidebar

Why this matters for project finance. The following figures are *illustrative, not modelled*: they show the sensitivity of standard project-finance metrics to a revenue overstatement, not an output of the case study. A BESS project finance package typically

targets a debt-service-coverage ratio (DSCR) above $1.4\times$ over the bank case. If the bank case revenue is overstated by 30% — Table 2’s two-hour bias at roughly 0.9 GW of *effective* price-setting capacity, the order of magnitude implied for GB conditions as of August 2026 — then the realised DSCR drops to $\sim 1.05\times$, which is the breach point for most facilities. By 2030, in fleet conditions where the naive bias exceeds 100%, projects financed on naive forecasts would be priced for default in their first year. The asymmetry favours the cautious modeller: an under-forecast can be repriced after one year of operations; an over-forecast cannot. CECadence, our revenue product (in private beta), reports the *minimum* DSCR over the bank case together with its binding year (P50 target $1.40\times$; P90 lender covenant $1.10\times$), so that the year in which coverage is tightest is surfaced explicitly rather than averaged away.

1.1 Cannibalisation: the term and its history

The term *cannibalisation* has well-established usage in the variable renewables literature. Hirth (2013) introduced the framework formally for wind and solar, defining the *market value* of a technology as the time-weighted average price it captures, and showing that the ratio of wind/solar market value to base-load market value declines as the share of wind/solar in the system rises. López Prol et al. (2020) extended this analysis to California with a clean empirical decomposition. Frew et al. (2021) explored saturation of solar value at high penetrations. The mechanism is intuitive: when wind generates, all wind generates simultaneously; the price-setter is displaced by zero-marginal-cost supply; the price falls; the average price wind *captures* (rather than the price absent wind) declines.

Storage cannibalisation differs structurally from variable-renewables cannibalisation in three important ways. First, storage acts on the *spread* between high and low prices rather than on the absolute price level: storage flattens the price duration curve from both ends. Second, storage is intertemporally coupled — its dispatch in one period constrains its availability in another — so the cannibalisation effect on a storage asset is path-dependent in a way wind cannibalisation is not. Third, storage cannibalisation feeds back *into the storage dispatch decision itself*: a unit’s optimal charge/discharge profile depends on the price spread it expects to capture, which in turn depends on the dispatch of every other unit. There is no analogous feedback for an individual wind farm whose generation is exogenous.

The literature on storage in price-influencing settings is sparser. Sioshansi et al. (2009) and Walawalkar et al. (2007) consider storage arbitrage in real-time markets but treat prices as exogenous. Bradbury et al. (2014) estimate price-taker arbitrage revenues across U.S. ISO markets without endogenous price formation. Mauch et al. (2012) evaluate compressed-air energy storage in PJM and explicitly note the price-feedback problem but adopt a one-shot heuristic correction. Brijs et al. (2016) solve a price-based unit commitment with piecewise linear price effects but do not formalise the equilibrium. Pandžić and Kuzle (2015) simulate a single-iteration adjustment in a day-ahead context. Schill and Kemfert (2011) and Sioshansi et al. (2011) treat strategic storage in oligopoly settings rather than competitive equilibrium. Industry practice, as represented in the publicly documented methodology descriptions of commercial providers such as Modo Energy, LCP Delta, and Aurora, typically treats prices either as exogenous or applies a single-pass post-hoc adjustment (Modo Energy, 2026; LCP Delta, 2026; Aurora Energy Research, 2025). None of these approaches formalises or solves the equilibrium fixed-point problem we identify here.

1.2 Roadmap

Section 2 sets up the market and the dispatch problem and describes the naive forecast. Section 3 introduces the operators and the fixed-point equation, and discusses the equivalent variational

inequality formulation. Section 4 states existence and uniqueness results. Section 5 addresses the failure of Picard iteration and the convergence of KM averaging. Section 6 reports the GB case study. Section 7 discusses extensions and limitations. Section 8 concludes.

2 The BESS revenue forecasting problem

2.1 Market structure and notation

We work in a discretised time horizon $\mathcal{T} = \{1, 2, \dots, T\}$ of half-hours; the case study uses $T = 8736$ (twenty-six weeks) and the conclusions extrapolate by linearity to longer horizons.¹ Throughout the paper, vectors over $\mathcal{T}$ are denoted in bold lowercase, scalars in *italic*. We denote the half-hour timestep $\Delta t = 0.5$ (hours). We use the convention that e_t for $t \in \{1, \dots, T\}$ denotes state-of-charge at the *end* of half-hour t , and we adopt an implicit initial state-of-charge $e_0 := \sigma_0 E^{\max}$ where $\sigma_0 \in [0, 1]$ is fixed.

The aggregate gross demand profile is $\mathbf{D} \in \mathbb{R}_+^T$ with $D_t \geq 0$ in GW. The aggregate non-dispatchable injection from variable renewables and inflexible thermal is $\mathbf{R} \in \mathbb{R}_+^T$. The aggregate net injection of the BESS fleet is $\mathbf{b} \in \mathbb{R}^T$, with the convention $b_t > 0$ when the fleet is net charging (drawing from the grid) and $b_t < 0$ when net discharging (note the sign convention: b_t is the fleet's net contribution to *residual demand*, so $b_t > 0$ is charging — physically a net withdrawal from the grid). The *residual demand* faced by dispatchable generation is therefore

$$\rho_t(\mathbf{b}) = D_t - R_t + b_t, \quad t \in \mathcal{T}, \quad (1)$$

and the wholesale price at half-hour t is determined by an inverse supply function $S : \mathbb{R} \rightarrow [\underline{p}, \bar{p}]$ applied to the residual demand:

$$\pi_t(\mathbf{b}) = S(\rho_t(\mathbf{b})). \quad (2)$$

We assume throughout:

Assumption 2.1 (Supply curve). S is continuous, non-decreasing, and bounded: there exist $-\infty < \underline{p} < \bar{p} < \infty$ such that $\underline{p} \leq S(x) \leq \bar{p}$ for all $x \in \mathbb{R}$. Furthermore S is differentiable and Lipschitz on the interior of the no-clip region $S^{-1}((\underline{p}, \bar{p}))$, with finite Lipschitz constant L_S .

We distinguish three price levels, which the GB literature frequently conflates. The economic scarcity ceiling is the value of lost load, $\text{VOLL} = 10,000$ GBP/MWh: this is the price at which involuntary demand reduction clears, and it is the right ceiling for the welfare program of Section 4 because it is the marginal social cost of the last unit of unserved energy. Distinct from this is the *administrative* market price cap of 2,000 GBP/MWh applied on N2EX intraday; this is a regulatory clip on traded prices, not a statement about scarcity value, and we retain it as a separate cap ($\text{cap} \neq \text{VOLL}$). The lower bound $\underline{p}$ is the negative-pricing floor (set in GB by curtailment compensation and CCGT minimum-load economics, around -50 GBP/MWh), which is reached on an oversupply tail when non-dispatchable injection exceeds demand — in the case-study calibration the floor binds in 17 half-hours and the price is negative in 379 (Section 6). Differentiability and Lipschitz continuity are satisfied by all smooth merit-order surrogates and by piecewise-linear merit orders away from kink points.

2.2 The BESS dispatch problem

The aggregate BESS fleet is parameterised by total power capacity $P^{\max} \in \mathbb{R}_+$ (GW), total energy capacity $E^{\max} = hP^{\max}$ (GWh, where h is duration in hours), and round-trip efficiency

¹A full-year robustness run ($T = 17,520$) shifts equilibrium revenue per MW by -4.6% (6 GW, 2 h) and -6.6% (0.5 GW, 1 h) relative to the 26-week-scaled figures, and the headline bias by -6.0 and $+0.5$ percentage points respectively (`full_year_check.csv`); the cyclic-boundary seam — whose relative weight the full year halves — and the winter-half spread account for the difference, which changes no qualitative conclusion.

$\eta \in (0, 1]$. We split efficiency symmetrically across charge and discharge: $\eta_c = \eta_d = \sqrt{\eta}$. At each half-hour $t \in \mathcal{T}$ the fleet's primal decision variables are the charge power $c_t \in [0, P^{\max}]$, discharge power $d_t \in [0, P^{\max}]$, and end-of-period state-of-charge $e_t \in [0, E^{\max}]$.

Given a price profile $\pi \in \mathbb{R}^T$, the fleet's optimal dispatch is the solution of the linear program

$$\begin{aligned} & \max_{\mathbf{c}, \mathbf{d}, \mathbf{e}} \sum_{t=1}^T \pi_t (d_t - c_t) \Delta t \\ & \text{subject to } e_t - e_{t-1} - \eta_c c_t \Delta t + \frac{1}{\eta_d} d_t \Delta t = 0, \quad t \in \mathcal{T}, \\ & \quad 0 \leq c_t \leq P^{\max}, \quad 0 \leq d_t \leq P^{\max}, \quad t \in \mathcal{T}, \\ & \quad 0 \leq e_t \leq E^{\max}, \quad t \in \mathcal{T}, \\ & \quad e_0 = \sigma_0 E^{\max}, \quad e_T \geq \sigma_0 E^{\max}, \end{aligned} \tag{3}$$

where e_0 is the implicit initial state-of-charge and $e_T \geq \sigma_0 E^{\max}$ is the (one-sided) cyclic boundary. We choose the inequality form rather than $e_T = \sigma_0 E^{\max}$ because at any optimum either the cyclic constraint binds (so $e_T = \sigma_0 E^{\max}$) or strictly more energy is left at the horizon end than required, in which case the additional energy has no economic value within the horizon and the constraint is slack without affecting the optimal objective. The inequality form is a relaxation of the equality form; the two coincide in optimal revenue whenever terminal stored energy has non-negative marginal value (a negative price in the final half-hours could make it strictly profitable to be paid to charge and hold the energy, which the equality form forbids), which holds in our calibration, and the inequality form keeps the feasible set non-empty whenever the equality form is.

The dispatch LP (3) has feasible set $K_D \subset \mathbb{R}^{3T}$ (a non-empty bounded polyhedron, since $\mathbf{c} = \mathbf{d} = \mathbf{0}$, $e_t = \sigma_0 E^{\max}$ is feasible for all π). Write $\mathcal{B} : \mathbb{R}^{3T} \rightarrow \mathbb{R}^T$, $\mathcal{B}(\mathbf{c}, \mathbf{d}, \mathbf{e}) = \mathbf{c} - \mathbf{d}$, and let $K_b := \mathcal{B}(K_D)$ denote the projected polytope of feasible net-injection profiles; both sets are used throughout. Define the dispatch operator on net injection $b_t := c_t - d_t$:

$$\mathcal{D} : \mathbb{R}^T \rightarrow 2^{\mathbb{R}^T}, \quad \mathcal{D}(\pi) := \{\mathbf{c}^* - \mathbf{d}^* : (\mathbf{c}^*, \mathbf{d}^*, \mathbf{e}^*) \text{ solves (3)}\}. \tag{4}$$

$\mathcal{D}$ is generically single-valued but set-valued at LP-degenerate prices. Where notational economy demands a single-valued map we adopt the lexicographic selection rule: among the optima of (3), pick the one that minimises $\sum_t (c_t + d_t)$ (least cycling) and break further ties by lexicographic ordering on $(c_1, d_1, c_2, d_2, \dots)$. With this convention the symbol $\mathcal{D}(\pi)$ denotes a single point of $\mathbb{R}^T$.

Remark 2.2. Equation (3) describes a competitive price-taker fleet: the optimisation treats π as exogenous. This is the right formulation because in the GB market no individual BESS owner sees enough of the fleet to coordinate strategically; competitive price-taking is also the only formulation under which an equilibrium computed by a central planner is implementable as a market outcome without intervention. Cournot or strategic-storage formulations (Schill and Kemfert, 2011; Hobbs, 2001) are surveyed briefly in Section 7.

2.3 The fixed-point equation

The price-formation operator

$$\mathcal{P} : \mathbf{b} \in \mathbb{R}^T \mapsto \pi(\mathbf{b}) \in \mathbb{R}^T, \quad \pi_t(\mathbf{b}) = S(D_t - R_t + b_t), \tag{5}$$

is well-defined and continuous under Assumption 2.1. The dispatch operator $\mathcal{D}$ from (4) is also well-defined for any $\pi \in \mathbb{R}^T$ because the dispatch LP is bounded and feasible (the trivial $\mathbf{c} = \mathbf{d} = \mathbf{0}$ is always feasible). Compose them:

$$\mathcal{F} : \mathbb{R}^T \rightarrow \mathbb{R}^T, \quad \mathcal{F} := \mathcal{D} \circ \mathcal{P}. \tag{6}$$

The cannibalisation-aware equilibrium fleet dispatch is by definition any $\mathbf{b}^* \in \mathbb{R}^T$ such that $\mathbf{b}^* = \mathcal{F}(\mathbf{b}^*)$. The corresponding equilibrium price profile is $\boldsymbol{\pi}^* = \mathcal{P}(\mathbf{b}^*)$. The two together form the cannibalisation-aware equilibrium pair $(\mathbf{b}^*, \boldsymbol{\pi}^*)$.

The fixed-point equation $\mathbf{b}^* = \mathcal{F}(\mathbf{b}^*)$ is the right *conceptual* object — it states exactly that the fleet’s optimal dispatch against the prices it induces reproduces itself — but it is the wrong object to iterate on directly. The price-formation map $\mathcal{P}$ is smooth, but the dispatch map $\mathcal{D}$ is the solution map of the linear program (3): it is piecewise constant on the LP-cells of price space and jumps discontinuously across the degeneracy hyperplanes between cells. The composition $\mathcal{F} = \mathcal{D} \circ \mathcal{P}$ therefore inherits the discontinuity of $\mathcal{D}$, is not non-expansive, and need not possess a fixed point in the classical (single-valued) sense; Kakutani recovers a set-valued fixed point (Theorem 3.5) but says nothing about how to compute it, and naive iteration on $\mathcal{F}$ chatters between adjacent LP optima. Section 4 resolves both existence and uniqueness by a different route — a strictly convex welfare program whose unique minimiser *is* the equilibrium net injection — and Section 5 shows that a regularised, contractive surrogate of $\mathcal{D}$ is what makes the iteration converge in practice.

2.4 The naive forecast as the zeroth Picard iterate

The naive (price-taker) BESS revenue forecast in widespread industry use can be written as one half of the zeroth Picard iterate of $\mathcal{F}$. Specifically:

$$\boldsymbol{\pi}^{\text{naive}} = \mathcal{P}(\mathbf{0}), \quad \mathbf{b}^{\text{naive}} = \mathcal{D}(\boldsymbol{\pi}^{\text{naive}}), \quad (7)$$

and the naive revenue forecast is

$$\Pi^{\text{naive}} = \sum_t \pi_t^{\text{naive}} (-b_t^{\text{naive}}) \Delta t. \quad (8)$$

Equivalently, the naive forecast assumes that the fleet’s own dispatch has zero impact on the price profile. The cannibalisation-aware revenue is

$$\Pi^* = \sum_t \pi_t^* (-b_t^*) \Delta t. \quad (9)$$

The bias of the naive forecast is $B := (\Pi^{\text{naive}} - \Pi^*)/\Pi^*$. Quantifying B as a function of fleet size and duration is the case study of Section 6.

Practitioner sidebar

What the naive forecast actually computes. Most commercial BESS revenue tools take a price profile produced from a fundamental, regression, or hybrid model that does not include the proposed BESS fleet, then run an arbitrage optimisation against that profile. Even when the upstream price model is calibrated to current observed prices, the implied dispatch from running an optimiser against a frozen profile is not what the same optimiser would do once its dispatch is feeding back into prices. The error is invisible in current-fleet calibrations because today’s fleet is small enough that the feedback is weak; it becomes structural at the projected fleet sizes that determine debt sizing for new builds.

3 The fixed-point framework

3.1 Operator properties

We collect the four operator properties required for the existence and convergence theorems below.

Lemma 3.1 (Lipschitz continuity of $\mathcal{P}$). *Under Assumption 2.1, the price-formation operator $\mathcal{P} : \mathbb{R}^T \rightarrow \mathbb{R}^T$ defined by (5) is Lipschitz continuous with constant L_S :*

$$\|\mathcal{P}(\mathbf{b}_1) - \mathcal{P}(\mathbf{b}_2)\|_\infty \leq L_S \|\mathbf{b}_1 - \mathbf{b}_2\|_\infty.$$

Proof. Componentwise, $|\pi_t(\mathbf{b}_1) - \pi_t(\mathbf{b}_2)| = |S(\rho_t(\mathbf{b}_1)) - S(\rho_t(\mathbf{b}_2))|$. By Assumption 2.1, S is Lipschitz with constant L_S on the no-clip region; on the clipping plateaux S is constant. In either case the inequality $|S(x) - S(y)| \leq L_S|x - y|$ holds. Since $\rho_t(\mathbf{b}_1) - \rho_t(\mathbf{b}_2) = b_{1,t} - b_{2,t}$, we obtain $|\pi_t(\mathbf{b}_1) - \pi_t(\mathbf{b}_2)| \leq L_S|b_{1,t} - b_{2,t}|$. Taking the maximum over t gives the stated inequality. $\square$

Lemma 3.2 (Monotonicity of $\mathcal{P}$). *Under Assumption 2.1, the price-formation operator $\mathcal{P}$ is monotone in the sense that*

$$\langle \mathcal{P}(\mathbf{b}_1) - \mathcal{P}(\mathbf{b}_2), \mathbf{b}_1 - \mathbf{b}_2 \rangle_{\Delta t} \geq 0 \quad \text{for all } \mathbf{b}_1, \mathbf{b}_2 \in \mathbb{R}^T.$$

Proof. Componentwise, $(\pi_t(\mathbf{b}_1) - \pi_t(\mathbf{b}_2))(b_{1,t} - b_{2,t}) = (S(\rho_{1,t}) - S(\rho_{2,t}))(\rho_{1,t} - \rho_{2,t}) \geq 0$ because S is non-decreasing (Assumption 2.1). Sum over t with non-negative weights Δt . $\square$

Lemma 3.3 (Boundedness of $\mathcal{D}$). *For any $\boldsymbol{\pi} \in \mathbb{R}^T$, every selection $\mathbf{b} \in \mathcal{D}(\boldsymbol{\pi})$ lies in the box $K := [-P^{\max}, P^{\max}]^T$.*

Proof. The constraints $0 \leq c_t \leq P^{\max}$ and $0 \leq d_t \leq P^{\max}$ give $-P^{\max} \leq c_t - d_t \leq P^{\max}$ for every feasible $(\mathbf{c}, \mathbf{d}, \mathbf{e})$, hence for any optimum. $\square$

Lemma 3.4 (Upper hemicontinuity of $\mathcal{D}$). *The set-valued dispatch operator $\mathcal{D} : \mathbb{R}^T \rightrightarrows \mathbb{R}^T$ is upper hemicontinuous with non-empty, compact, convex values. Equivalently: for every $\boldsymbol{\pi} \in \mathbb{R}^T$, $\mathcal{D}(\boldsymbol{\pi})$ is a non-empty compact convex subset of K , and the graph $\{(\boldsymbol{\pi}, \mathbf{b}) : \mathbf{b} \in \mathcal{D}(\boldsymbol{\pi})\}$ is closed.*

Proof. The dispatch LP (3) has feasible polyhedron K_D independent of $\boldsymbol{\pi}$ and a linear objective in $\boldsymbol{\pi}$. The optimal value $v(\boldsymbol{\pi}) = \min_{\mathbf{b} \in K_b} \langle \boldsymbol{\pi}, \mathbf{b} \rangle_{\Delta t}$ is finite (since K_D is bounded), and the optimal-set correspondence $\boldsymbol{\pi} \mapsto \mathcal{D}(\boldsymbol{\pi}) = \{\mathbf{b} \in K_b : \langle \boldsymbol{\pi}, \mathbf{b} \rangle_{\Delta t} = v(\boldsymbol{\pi})\}$ is the intersection of K_b with a closed half-space whose boundary depends continuously on $\boldsymbol{\pi}$. Berge's maximum theorem and Robinson's parametric-LP stability theorem (both treated in Facchinei and Pang 2003) apply and give upper hemicontinuity. Convexity and compactness of $\mathcal{D}(\boldsymbol{\pi})$ are immediate from the LP structure. $\square$

When the lexicographic selection rule of Section 2 is applied, $\mathcal{D}$ becomes a single-valued map $\mathbb{R}^T \rightarrow \mathbb{R}^T$. Continuity of this single-valued selection holds at every $\boldsymbol{\pi}$ where the LP has a unique optimum — a property that fails on a measure-zero set of $\mathbb{R}^T$ (the LP-degenerate parameters). The fixed-point arguments below first prove existence using the set-valued formulation (Kakutani) and then specialise to a fixed point of the single-valued lexicographic selection where convenient.

3.2 Existence

Theorem 3.5 (Existence of equilibrium). *Under Assumption 2.1, the set-valued operator $\mathcal{F} : K \rightrightarrows K$ defined by $\mathcal{F}(\mathbf{b}) := \mathcal{D}(\mathcal{P}(\mathbf{b}))$ has at least one fixed point in $K = [-P^{\max}, P^{\max}]^T$. Equivalently, there exists $\mathbf{b}^* \in K$ and $\boldsymbol{\pi}^* = \mathcal{P}(\mathbf{b}^*) \in \mathbb{R}^T$ such that $\mathbf{b}^* \in \mathcal{D}(\boldsymbol{\pi}^*)$.*

Proof. Step 1 (domain). $K = [-P^{\max}, P^{\max}]^T \subset \mathbb{R}^T$ is non-empty, compact, and convex.

Step 2 (image inside K). For any $\mathbf{b} \in K$, $\mathcal{P}(\mathbf{b}) \in \mathbb{R}^T$ is well-defined (Lemma 3.1). For any $\boldsymbol{\pi} \in \mathbb{R}^T$, $\mathcal{D}(\boldsymbol{\pi}) \subset K$ (Lemma 3.3). Hence $\mathcal{F}(\mathbf{b}) \subset K$ and $\mathcal{F}$ maps K into the family of non-empty compact convex subsets of K .

Step 3 (closed graph). $\mathcal{P}$ is continuous (Lemma 3.1) and $\mathcal{D}$ has closed graph (Lemma 3.4). The composition of a continuous single-valued map with a closed-graph correspondence has closed graph: if $(\mathbf{b}_n, \mathbf{x}_n) \rightarrow (\mathbf{b}, \mathbf{x})$ with $\mathbf{x}_n \in \mathcal{F}(\mathbf{b}_n) = \mathcal{D}(\mathcal{P}(\mathbf{b}_n))$, then $\mathcal{P}(\mathbf{b}_n) \rightarrow \mathcal{P}(\mathbf{b})$ by continuity, and $\mathbf{x}_n \in \mathcal{D}(\mathcal{P}(\mathbf{b}_n))$ together with $\mathbf{x}_n \rightarrow \mathbf{x}$ implies $\mathbf{x} \in \mathcal{D}(\mathcal{P}(\mathbf{b})) = \mathcal{F}(\mathbf{b})$ by closed-graph of $\mathcal{D}$.

Step 4 (apply Kakutani). The conditions of Kakutani’s fixed-point theorem (Granás and Dugundji, 2003) are satisfied: K is non-empty, compact, and convex; $\mathcal{F} : K \rightrightarrows K$ has non-empty, compact, convex values; and $\mathcal{F}$ has closed graph (Step 3). Hence there exists $\mathbf{b}^* \in K$ with $\mathbf{b}^* \in \mathcal{F}(\mathbf{b}^*)$.

For the single-valued lexicographic selection of $\mathcal{D}$, the same conclusion follows from Brouwer’s fixed-point theorem (Granás and Dugundji 2003) at every fixed point in the (open) set of LP non-degeneracy. Existence on the closed set K requires Kakutani because the lexicographic selection is not continuous everywhere; the Kakutani argument above bypasses this. $\square$

Corollary 3.6 (Bounds on equilibrium price). *Any equilibrium price profile π^* satisfies $\underline{p} \leq \pi_t^* \leq \bar{p}$ for all $t \in \mathcal{T}$.*

Proof. Immediate from $\pi^* = \mathcal{P}(\mathbf{b}^*)$, the definition of $\mathcal{P}$, and the bounds in Assumption 2.1. $\square$

3.3 Variational inequality formulation

The fixed-point problem can be reformulated as a variational inequality (VI), which is the natural setting for further structural analysis. Recall the projection $\mathcal{B}$ and the projected feasible set $K_b = \mathcal{B}(K_D)$ from Section 2. By eliminating $\mathbf{e}$ via the SoC balance, K_b is the polyhedron in $\mathbb{R}^T$ defined by the bound and intertemporal constraints carried over from (3).

The objective of (3) written in $\mathbf{b} = \mathbf{c} - \mathbf{d}$ is $\sum_t \pi_t(d_t - c_t)\Delta t = -\langle \pi, \mathbf{b} \rangle_{\Delta t}$, where $\langle \mathbf{x}, \mathbf{y} \rangle_{\Delta t} := \sum_t x_t y_t \Delta t$. Since $\mathcal{D}(\pi)$ minimises $\langle \pi, \mathbf{b} \rangle_{\Delta t}$ over $\mathbf{b} \in K_b$, the standard linear-programming first-order condition gives

$$\langle \pi^*, \mathbf{b}' - \mathbf{b}^* \rangle_{\Delta t} \geq 0 \quad \text{for all } \mathbf{b}' \in K_b, \quad (10)$$

which is precisely the VI $\text{VI}(K_b, \pi^*)$. Substituting $\pi^* = \mathcal{P}(\mathbf{b}^*)$ closes the loop and gives the cannibalisation VI:

$$\langle \mathcal{P}(\mathbf{b}^*), \mathbf{b}' - \mathbf{b}^* \rangle_{\Delta t} \geq 0 \quad \text{for all } \mathbf{b}' \in K_b. \quad (11)$$

The pair $(\mathbf{b}^*, \pi^*) = (\mathbf{b}^*, \mathcal{P}(\mathbf{b}^*))$ solves (11) if and only if $\mathbf{b}^*$ is a fixed point of $\mathcal{F} = \mathcal{D} \circ \mathcal{P}$.

This VI formulation is convenient because the rich convergence theory of monotone variational inequalities (Facchinei and Pang, 2003, Ch. 12) applies whenever the operator $\mathbf{b} \mapsto \mathcal{P}(\mathbf{b})$ is monotone — which, by Lemma 3.2 above, follows directly from monotonicity of the supply curve S .

3.4 Connection to Cournot equilibrium

Treating $\mathcal{P}$ as the inverse demand of a single “commodity” (net injection of dispatchable plus storage power) and the dispatch problem as the supply side, the fixed point coincides with the competitive equilibrium of a Cournot-Walras-type game with a continuum of price-taking storage operators (Fudenberg and Tirole, 1991; Hobbs, 2001). The Nash equilibrium with finitely many strategic storage operators sits between the price-taker equilibrium considered here and the social-planner solution; we treat the price-taker case as the operationally relevant benchmark for GB-scale deployments where individual operators do not have market power.

4 Existence, uniqueness, and structure

Theorem 3.5 delivers existence of a set-valued fixed point via Kakutani, but the construction is non-constructive and silent on uniqueness: the dispatch LP has potentially many optima — generic when the price profile contains repeated values — and the composite map $\mathcal{F}$ is discontinuous, so naive iteration chatters between them. We now give a self-contained existence-and-uniqueness result through a strictly convex welfare program, which is both the cleaner statement and the object our solver actually targets.

Assumption 4.1 (Strict monotonicity of S in interior). S is strictly increasing on the interior of $S^{-1}((p, \bar{p}))$.

This assumption is satisfied by any merit-order curve with no flat sections in the relevant operating regime — equivalent to “no two distinct generators have identical short-run marginal cost up to the rounding precision of dispatch”, which holds in any real market.

4.1 The competitive equilibrium as a convex welfare program

Write $R_t^0 := D_t - R_t$ for the no-BESS residual demand, so that the residual faced by dispatchable generation at net injection $\mathbf{b}$ is $R_t^0 + b_t$. Define the per-period welfare potential as the antiderivative of the supply curve,

$$\Phi(x) := \int_0^x S(u) \, du, \quad (12)$$

and consider the convex program

$$\min_{\mathbf{b} \in K_b} W(\mathbf{b}) := \sum_{t=1}^T \Phi(R_t^0 + b_t) \Delta t, \quad (13)$$

over the projected net-injection polytope K_b of Section 3.3. Because $\Phi' = S$ is non-decreasing (Assumption 2.1), Φ is convex; because S is strictly increasing on the no-clip region (Assumption 4.1), Φ is *strictly* convex there. Note that W is the negative of social welfare net of a fixed generation-cost constant, so (13) is the standard market-clearing welfare problem with storage as a flexible participant.

Theorem 4.2 (Equilibrium as the unique welfare minimiser). *Under Assumptions 2.1 and 4.1, the convex program (13) attains its minimum on the non-empty compact polytope K_b . Let $\mathbf{b}^*$ be any minimiser and $R_t^* := R_t^0 + b_t^*$. Then:*

- (i) *the equilibrium residual $\mathbf{R}^*$ is unique on every half-hour at which R_t^* lies in the no-clip region, and consequently the price profile $\boldsymbol{\pi}^*$ with $\pi_t^* = S(R_t^*)$ is unique there;*
- (ii) *$(\mathbf{b}^*, \boldsymbol{\pi}^*)$ solves the cannibalisation VI (11), hence $\mathbf{b}^*$ is an equilibrium net injection in the sense of Section 3; and*
- (iii) *π_t^* is the Lagrange multiplier (dual variable) of the power-balance identity $R_t = R_t^0 + b_t$ at the optimum — i.e. the equilibrium price is the dual of power balance, $\boldsymbol{\pi}^* = \nabla W(\mathbf{b}^*)$ componentwise, where ∇W denotes the Riesz representative of the differential in the inner product $\langle \cdot, \cdot \rangle_{\Delta t}$ (so $[\nabla W(\mathbf{b}^*)]_t = S(R_t^*)$, with no stray Δt factor). The identity holds except at half-hours whose optimal residual is clipped (17 of 8,736 at the floor in the no-BESS calibration, and fewer at equilibrium), where the reported price is the clipped value while the welfare gradient is the unclipped one — the clipping remark following Proposition 4.4 discusses this.*

Proof. Existence of a minimiser is Weierstrass: W is continuous and K_b is a non-empty compact polytope (Section 3.3). The first-order optimality condition for the convex program (13) is the variational inequality

$$\langle \nabla W(\mathbf{b}^*), \mathbf{b}' - \mathbf{b}^* \rangle_{\Delta t} \geq 0 \quad \text{for all } \mathbf{b}' \in K_b, \quad [\nabla W(\mathbf{b}^*)]_t = \Phi'(R_t^0 + b_t^*) = S(R_t^*),$$

where the gradient is taken as the Riesz representative in $\langle \cdot, \cdot \rangle_{\Delta t}$ (the Euclidean partial derivative $\partial W / \partial b_t = S(R_t^*) \Delta t$; the Δt is absorbed by the weighted inner product, so no Δt^2 factor appears in the VI). Identifying $\pi_t^* = S(R_t^*)$ gives exactly (11), proving (ii) and the dual identification

(iii); the price is the gradient of the welfare potential, which is the multiplier on the residual-defining equality $R_t = R_t^0 + b_t$. For (i), suppose $\mathbf{b}_1^*, \mathbf{b}_2^*$ are two minimisers. By convexity of W the minimiser set is convex and W is constant on it; strict convexity of Φ on the no-clip region forces $R_t^0 + b_{1,t}^* = R_t^0 + b_{2,t}^*$, i.e. $b_{1,t}^* = b_{2,t}^*$, at every half-hour at which *both* optimal residuals — and hence the whole connecting segment $[R_{1,t}^*, R_{2,t}^*]$ — lie in the no-clip region (otherwise the midpoint would have strictly smaller W , contradicting optimality; the midpoint argument needs strict convexity along the whole segment, not only at its endpoints). Uniqueness of π^* on the no-clip portion follows from $\pi_t^* = S(R_t^*)$. $\square$

Remark 4.3 (Why the convex route is the operative one). Theorem 4.2 delivers what the fixed-point statement cannot: a unique equilibrium residual and price, obtained without any contraction or non-expansiveness hypothesis on $\mathcal{F}$, and an unambiguous economic reading of the price as the dual of power balance. The fixed-point and VI formulations of Sections 3–3.3 remain the right way to *state* the equilibrium and to connect it to the price-taker dispatch problem; the convex program is the right way to *prove* it exists and is unique, and — after the regularisation of Section 5 — to *compute* it.

The same uniqueness can be read off the variational-inequality formulation, which makes the role of monotonicity explicit and confirms that any fixed point of $\mathcal{F}$ coincides with the welfare minimiser of Theorem 4.2 on the no-clip portion.

Proposition 4.4 (Uniqueness of equilibrium net injection at the fleet level — VI route). *Under Assumptions 2.1 and 4.1, if $\mathbf{b}_1^*, \mathbf{b}_2^* \in K$ are both equilibrium fixed points of $\mathcal{F}$ then they coincide on every half-hour at which the residual demand lies in the strict interior of the no-clip region of S . In particular, the equilibrium price profile $\pi^* = \mathcal{P}(\mathbf{b}^*)$ is unique on the no-clip portion of the time horizon, and agrees with the welfare-program price π^* of Theorem 4.2.*

Proof. By Lemma 3.2, $\mathcal{P}$ is monotone, and by Assumption 4.1 it is strictly monotone on the no-clip portion of $\mathbb{R}^T$, i.e. $\langle \mathcal{P}(\mathbf{b}_1) - \mathcal{P}(\mathbf{b}_2), \mathbf{b}_1 - \mathbf{b}_2 \rangle_{\Delta t} > 0$ whenever $\mathbf{b}_1$ and $\mathbf{b}_2$ differ at a no-clip half-hour. From the VI (11) applied at both equilibria,

$$\langle \mathcal{P}(\mathbf{b}_1^*), \mathbf{b}_2^* - \mathbf{b}_1^* \rangle_{\Delta t} \geq 0, \quad \langle \mathcal{P}(\mathbf{b}_2^*), \mathbf{b}_1^* - \mathbf{b}_2^* \rangle_{\Delta t} \geq 0.$$

Adding the two inequalities gives $\langle \mathcal{P}(\mathbf{b}_1^*) - \mathcal{P}(\mathbf{b}_2^*), \mathbf{b}_2^* - \mathbf{b}_1^* \rangle_{\Delta t} \geq 0$, i.e. $\langle \mathcal{P}(\mathbf{b}_1^*) - \mathcal{P}(\mathbf{b}_2^*), \mathbf{b}_1^* - \mathbf{b}_2^* \rangle_{\Delta t} \leq 0$. Combined with monotonicity (≥ 0), the inner product equals zero. Strict monotonicity on the no-clip portion forces $b_{1,t}^* = b_{2,t}^*$ at every no-clip half-hour. Equality of π^* on the no-clip portion follows from $\pi^* = \mathcal{P}(\mathbf{b}^*)$ and the value of S at the same residual. $\square$

Remark 4.5 (Uniqueness of fleet revenue). The aggregate fleet revenue $\Pi^* = -\langle \pi^*, \mathbf{b}^* \rangle_{\Delta t}$ is uniquely determined at any equilibrium even when the fleet’s individual charge/discharge schedule is not (because at any LP optimum the objective value is the same). At equilibrium prices π^* , every selection $\mathbf{b}^* \in \mathcal{D}(\pi^*)$ yields the same value $\langle \pi^*, \mathbf{b}^* \rangle_{\Delta t}$ by the LP optimality of all elements of $\mathcal{D}(\pi^*)$. Practically, what the BESS revenue forecaster needs — the fleet revenue — is therefore a well-defined function of fleet specification, regardless of LP degeneracy.

Remark 4.6 (On clipping). Where the residual demand pushes S to its floor or cap, multiple net-injection profiles can yield the same clipped price. This is economically meaningful: when the price is at the cap, additional storage discharge that drops the price into the no-clip regime would change revenue, but storage discharge that merely shifts the residual within the clipped regime would not. The case study below stays in the no-clip regime for the great majority of half-hours, so the uniqueness of Proposition 4.4 obtains at almost every t .

Practitioner sidebar

Why uniqueness of revenue is enough. Investors do not need to predict the exact half-hour-by-half-hour dispatch a battery will execute fifteen years from now; they need

an unbiased estimate of average annual revenue. Proposition 4.4 and Remark 4.5 together show that the equilibrium prices and revenue are well-defined even when there are multiple optimal dispatch schedules. The methodological challenge is computing them; the economic concept itself is unambiguous.

5 Convergence: where naive Picard fails and KM averaging succeeds

5.1 Plain Picard iteration

The natural algorithm for finding a fixed point of $\mathcal{F}$ is Picard iteration:

$$\mathbf{b}^{(k+1)} \in \mathcal{F}(\mathbf{b}^{(k)}) = \mathcal{D}(\mathcal{P}(\mathbf{b}^{(k)})), \quad (14)$$

where, when $\mathcal{D}$ is set-valued, the iterate is taken under the lexicographic selection. Picard iteration converges if $\mathcal{F}$ is a strict contraction, in which case the Banach fixed-point theorem (Banach, 1922; Granas and Dugundji, 2003) gives the geometric rate $\|\mathbf{b}^{(k)} - \mathbf{b}^*\| \leq \kappa^k \|\mathbf{b}^{(0)} - \mathbf{b}^*\|$ with contraction modulus $\kappa < 1$.

In the BESS revenue problem $\mathcal{F}$ is not a contraction — indeed it is not even continuous. The Lipschitz constant of $\mathcal{P}$ is bounded by L_S (Lemma 3.1). The dispatch operator $\mathcal{D}$, however, is the solution map of a linear program: as a function of the price parameter $\boldsymbol{\pi}$, $\mathcal{D}(\boldsymbol{\pi})$ is piecewise constant on the strict interior of each LP-cell of the parameter space and jumps discontinuously across the LP-degeneracy hyperplanes that separate cells. Locally near a degeneracy the effective amplification factor of $\mathcal{F}$ is unbounded. In markets with steep scarcity tails (the GB calibration reaches L_S of about 120 GBP/MWh per GW at the peak realised residual demand, and exceeds 300 only beyond roughly 46 GW of residual demand, outside the simulated range — which is why a regularisation chosen above the realised slope contracts with margin, Section 5.3), the iterate (14) lands in adjacent LP-cells across iterations, so it oscillates rather than converges. Damping does not cure the discontinuity: a relaxed Picard iteration $\mathbf{b}^{(k+1)} = (1 - \alpha)\mathbf{b}^{(k)} + \alpha\mathcal{F}(\mathbf{b}^{(k)})$ with $\alpha \in (0, 1)$ smooths the oscillation but settles to an *irreducible residual floor* rather than to the equilibrium, and that floor grows monotonically with fleet penetration.

The case-study sweep (Section 6, Figure 6 and Table 3) makes this quantitative. At the tolerance required for investment-grade revenue — we use $\varepsilon = 0.05$ GBP/MWh, comfortably below the half-hour price resolution that materially moves arbitrage revenue — relaxed Picard does *not* converge at *any* fleet size, including the smallest tested (0.1 GW): across a full grid of dampings $\alpha \in \{0.1, \dots, 0.9\}$, fleet sizes 0.1–10 GW, durations 1/2/4 hours, and both a 4-week window and the full 26-week horizon (200 iterations per cell), no cell ever reaches the tolerance. Its irreducible residual floor (two-hour fleets, best damping per cell, 4-week window) rises from about 0.2 GBP/MWh at 0.1 GW, through 2.9 GBP/MWh at 1 GW and 7.1 GBP/MWh at 2 GW, to 13.0 GBP/MWh at 4 GW and 29.7 GBP/MWh at 10 GW — and grows further on the full horizon (17.2 at 4 GW, 38.3 at 10 GW). In words: Picard’s irreducible error exceeds the investment-grade tolerance everywhere, crosses 1 GBP/MWh by 0.5 GW, and exceeds 20 GBP/MWh beyond about 6 GW, so at the tolerance an investment-grade revenue forecast demands it fails throughout the deployment range. This is sharper than the heuristic “plain Picard fails beyond a few GW”: the failure is qualitative (a discontinuous map has no reason to admit a convergent Picard orbit) and is present at every penetration; it merely becomes economically unmissable above 1–2 GW. The regularised, contractive scheme below is therefore not an optimisation nicety but a correctness requirement.

5.2 Anderson acceleration

Anderson acceleration (Anderson, 1965; Walker and Ni, 2011) maintains a sliding window of the most recent $m + 1$ iterates and residuals and chooses the next iterate as the affine combination minimising the aggregate residual. Define the residual map $g(\mathbf{b}) := \mathcal{F}(\mathbf{b}) - \mathbf{b}$. Type-II Anderson with depth m replaces (14) with

$$\mathbf{b}^{(k+1)} = \sum_{i=0}^{m_k} \alpha_i^{(k)} \mathcal{F}(\mathbf{b}^{(k-i)}), \quad \sum_{i=0}^{m_k} \alpha_i^{(k)} = 1, \quad (15)$$

where $m_k = \min(m, k)$ and the weights solve the constrained least-squares problem

$$\alpha^{(k)} \in \arg \min_{\alpha \in \mathbb{R}^{m_k+1}, \sum_i \alpha_i = 1} \left\| \sum_{i=0}^{m_k} \alpha_i g(\mathbf{b}^{(k-i)}) \right\|_2^2. \quad (16)$$

The standard reduction to unconstrained least squares uses the substitution $\beta_i := \alpha_{i+1}$ for $i = 0, \dots, m_k - 1$ with $\alpha_0 = 1 - \sum \beta_i$, after which (16) becomes an ordinary least-squares problem in $\beta \in \mathbb{R}^{m_k}$ on the residual differences $g(\mathbf{b}^{(k)}) - g(\mathbf{b}^{(k-i-1)})$, solvable in closed form via the normal equations or QR decomposition. For contractive $\mathcal{F}$, Anderson acceleration converges locally r -linearly (Toth and Kelley, 2015); for non-contractive $\mathcal{F}$ it is empirically robust in many problems but lacks a general convergence theorem.

5.3 μ -regularised Krasnoselski–Mann averaging

The discontinuity of $\mathcal{D}$ that defeats Picard is removed by quadratically regularising the dispatch objective. Replace the linear dispatch objective $\langle \boldsymbol{\pi}, \mathbf{b} \rangle_{\Delta t}$ (which $\mathcal{D}$ minimises, Section 3.3) by the strongly convex objective $\langle \boldsymbol{\pi}, \mathbf{b} \rangle_{\Delta t} + \frac{\mu}{2} \sum_t (c_t^2 + d_t^2) \Delta t$ with $\mu > 0$ — which reduces to $\langle \boldsymbol{\pi}, \mathbf{b} \rangle_{\Delta t} + \frac{\mu}{2} \|\mathbf{b}\|_{\Delta t}^2$ at any optimum without simultaneous charge and discharge — and denote the resulting solution map $\mathcal{D}_\mu$. The regularised dispatch problem is a strictly convex quadratic program: its solution is *single-valued and Lipschitz* in $\boldsymbol{\pi}$ (with constant $1/\mu$), so the regularised composite $\mathcal{F}_\mu := \mathcal{D}_\mu \circ \mathcal{P}$ is continuous, and for $\mu > L_S$ it is a genuine contraction. This is the device that turns an intractable discontinuous fixed point into a convergent one; the QP is the Tikhonov regularisation of the welfare program (13) — the fixed point of $\mathcal{F}_\mu$ is exactly the minimiser of $W(\mathbf{b}) + \frac{\mu}{2} \|\mathbf{b}\|_{\Delta t}^2$ — and as $\mu \downarrow 0$ the regularised equilibrium converges to the unregularised one. In the case study the convergence certificate uses a fixed μ chosen comfortably above the realised supply-curve slope; the equilibrium revenue is reported from the unregularised welfare solution, with the QP used only to drive the iteration, and price-taker consistency is certified on the *reported* dispatch (the consistency gap is below 2×10^{-2} throughout the sweep and below 7.4×10^{-3} in the two-hour block; Section 6 and Appendix B).

The algorithm we recommend for production use, and which converges robustly in our case study, is Krasnoselski–Mann (KM) averaging (Krasnoselskiĭ, 1955; Mann, 1953) applied to the regularised map $\mathcal{F}_\mu$:

$$\mathbf{b}^{(k+1)} = (1 - \lambda_k) \mathbf{b}^{(k)} + \lambda_k \mathcal{F}_\mu(\mathbf{b}^{(k)}), \quad \lambda_k \in [0, 1], \quad (17)$$

with step-size schedule λ_k chosen to satisfy $\sum_k \lambda_k (1 - \lambda_k) = \infty$. We use $\lambda_k = 2/(k+2)$, which gives $\lambda_k \in (0, 1]$, $\lambda_k \rightarrow 0$, $\sum_k \lambda_k = \infty$, and $\sum_k \lambda_k (1 - \lambda_k) = \infty$. Because $\mathcal{F}_\mu$ is a contraction (hence *a fortiori* non-expansive) for the μ we use, the fundamental KM theorem applies directly in the real Hilbert space $(\mathbb{R}^T, \langle \cdot, \cdot \rangle_{\Delta t})$; no orbit-dependent hypothesis is needed.

Theorem 5.1 (Krasnoselski–Mann convergence; Krasnoselskiĭ 1955; Mann 1953). *Let H be a real Hilbert space, $C \subset H$ a non-empty closed convex subset, and let $\mathcal{F} : C \rightarrow C$ be non-expansive, i.e.*

$$\|\mathcal{F}(\mathbf{x}) - \mathcal{F}(\mathbf{y})\| \leq \|\mathbf{x} - \mathbf{y}\| \quad \text{for all } \mathbf{x}, \mathbf{y} \in C.$$

Suppose $\mathcal{F}$ has at least one fixed point in C . Let $(\lambda_k) \subset [0, 1]$ satisfy $\sum_k \lambda_k(1 - \lambda_k) = \infty$. Then for any $\mathbf{b}^{(0)} \in C$, the sequence $(\mathbf{b}^{(k)})$ generated by (17) converges weakly to a fixed point of $\mathcal{F}$. In finite dimension, weak convergence equals strong convergence in norm.

Proof. This is the classical Krasnoselski–Mann theorem; for a textbook proof in modern notation see Bauschke and Combettes (2017, Thm. 5.15). The original results are Krasnoselskiĭ (1955) (case $\lambda_k \equiv 1/2$) and Mann (1953) (general step sizes). $\square$

Theorem 5.2 (Convergence rate of KM averaging). *Under the hypotheses of Theorem 5.1, the residual of the KM iterates obeys*

$$\|\mathbf{b}^{(k)} - \mathcal{F}(\mathbf{b}^{(k)})\|^2 \leq \frac{\|\mathbf{b}^{(0)} - \mathbf{b}^*\|^2}{\sum_{j=0}^{k-1} \lambda_j(1 - \lambda_j)},$$

where $\mathbf{b}^* \in C$ is any fixed point of $\mathcal{F}$. With $\lambda_k = 2/(k+2)$ the denominator grows like $\Theta(\log k)$, so the residual decays at the rate $\|\mathbf{b}^{(k)} - \mathcal{F}(\mathbf{b}^{(k)})\| = O(1/\sqrt{\log k})$; with constant $\lambda_k \equiv 1/2$ the denominator grows like $k/4$ and the residual decays at the rate $O(1/\sqrt{k})$.

Proof. This is the standard residual bound for Krasnoselski–Mann iterates of a non-expansive map; see Bauschke and Combettes (2017). The rate is sharp without further structure on $\mathcal{F}$. $\square$

Remark 5.3 (Choice of step-size schedule). Theorem 5.2 suggests that constant $\lambda_k \equiv 1/2$ gives a faster worst-case rate than the harmonic schedule $\lambda_k = 2/(k+2)$. In practice the bound is loose: because $\mathcal{F}_\mu$ is a contraction, KM converges far faster than the non-expansive worst case — in the case-study sweep it reaches the 0.05 GBP/MWh tolerance in 10–19 iterations across all fleet sizes and durations (Section 6). A direct step-schedule comparison across fleet sizes 4–10 GW and all three durations confirms the theorem’s ordering on the contractive map: constant $\lambda = 2/3$ reaches tolerance in 6–7 iterations and constant $\lambda = 1/2$ in 8–10, against 10–14 for the harmonic schedule — so the harmonic schedule used in the reported runs (Figure 6) is conservative rather than optimal (`step_schedule_ab.csv`).

Remark 5.4 (Why we regularise rather than assume non-expansiveness). Theorem 5.1 requires the iterated map to be non-expansive. The *unregularised* composite $\mathcal{F} = \mathcal{D} \circ \mathcal{P}$ is *not* non-expansive — it is not even continuous, because $\mathcal{D}$ is the discontinuous solution map of a linear program (Section 3). We therefore do not iterate $\mathcal{F}$ and we do not claim non-expansiveness for it. Instead we iterate the regularised map $\mathcal{F}_\mu = \mathcal{D}_\mu \circ \mathcal{P}$, for which non-expansiveness is not an assumption but a consequence: $\mathcal{D}_\mu$ is the Lipschitz solution map of a strictly convex QP, and for $\mu > L_S$ the composition is a strict contraction, so Theorem 5.1 applies unconditionally.

Non-expansiveness of the *unregularised* orbit is then a quantity we *monitor*, not one we rely on: in production we monitor the empirical contraction ratio of the iteration and hold it below one. We deliberately avoid the tempting shortcut of declaring $\mathcal{P}$ firmly non-expansive after a coordinate rescaling: even where the Baillon–Haddad theorem (Baillon and Haddad, 1977; Bauschke and Combettes, 2017) would make $\mathcal{P}$ firmly non-expansive in a rescaled norm, it cannot repair the discontinuity of $\mathcal{D}$, so rescaling does not make $\mathcal{F}$ non-expansive. The regularisation does, and it does so with an explicit, controllable bias (the QP solution converges to the welfare minimiser as $\mu \downarrow 0$), which is the property an investment-grade pipeline needs.

5.4 Algorithmic summary

The algorithm we implement in the Compounding Energy stack, and reproduce below in a standalone form for the case study, is:

Algorithm 1 Cannibalisation-aware BESS revenue computation (μ -regularised KM iteration)

Require: Demand $\mathbf{D}$, renewables $\mathbf{R}$, supply curve S , fleet $(\mathbf{P}^{\max}, h, \eta, \sigma_0)$, regularisation $\mu > 0$, tolerance $\varepsilon > 0$, max-iter K

Ensure: Equilibrium prices $\boldsymbol{\pi}^*$ and revenue Π^*

```

1:  $\mathbf{b}^{(0)} \leftarrow \mathbf{0}$ 
2: for  $k = 0, 1, \dots, K - 1$  do
3:    $\boldsymbol{\pi}^{(k)} \leftarrow S(\mathbf{D} - \mathbf{R} + \mathbf{b}^{(k)})$  ▷ price-formation step
4:    $\mathbf{b}^{(\text{tmp})} \leftarrow \mathcal{D}_\mu(\boldsymbol{\pi}^{(k)})$  ▷ regularised dispatch QP (Section 5.3)
5:    $\lambda_k \leftarrow 2/(k + 2)$ 
6:    $\mathbf{b}^{(k+1)} \leftarrow (1 - \lambda_k)\mathbf{b}^{(k)} + \lambda_k\mathbf{b}^{(\text{tmp})}$ 
7:   if  $\|\boldsymbol{\pi}^{(k+1)} - \boldsymbol{\pi}^{(k)}\|_\infty < \varepsilon$  then
8:     break
9:   end if
10: end for
11:  $\mathbf{b}^* \leftarrow$  minimiser of the welfare program (13), reported as its SoC-preserving non-simultaneous dispatch
12:  $\boldsymbol{\pi}^* \leftarrow S(\mathbf{D} - \mathbf{R} + \mathbf{b}^*)$  ▷  $\boldsymbol{\pi}^*$  is the dual of power balance
13: certify  $|\Pi^{\text{taker}}(\boldsymbol{\pi}^*) - \Pi^*|/|\Pi^*| \leq \varepsilon_{\text{cert}}$  ▷ price-taker / KKT consistency on the reported dispatch (Appendix B)
14:  $\Pi^* \leftarrow -\langle \boldsymbol{\pi}^*, \mathbf{b}^* \rangle_{\Delta t}$ 

```

The complexity per outer iteration is dominated by the regularised dispatch solve; the banded state-of-charge structure keeps its cost close to linear in T (measured scaling exponent ≈ 1.2 over $T = 672\text{--}8,736$; one solve takes 0.86 s at the full horizon, and the full-horizon convergence certificate completes in under 6 s within 280 MB). Convergence in our case study is reached in 10–19 outer iterations across the relevant fleet range (Section 6); the final equilibrium price and revenue are taken from the unregularised welfare solution, reported as its SoC-preserving non-simultaneous dispatch, and certified against the price-taker dispatch problem, with a consistency gap below 2×10^{-2} throughout and below 7.4×10^{-3} in the two-hour block (Appendix B).

6 Stylised GB case study

We now report numerical results from a stylised Great Britain market calibrated to a historical 2024–2025 day-ahead reference window, retained for comparability (Table 1 compares the calibration against realised prices for the twelve months to August 2026). All figures and tables in this section are reproducible from the Python implementation accompanying this paper (available on request).

6.1 Setup

The market has one zone, $T = 8736$ half-hours (twenty-six weeks; revenue scaled to annual by the factor $365/182 \approx 2.005$), and uses the smooth supply curve

$$S(R) = \text{clip}_{[\underline{p}, \bar{p}_{\text{cap}}]} \left(\min \left(a(R/R_{\text{ref}})_+^k + b e^{(R-R_h)/\sigma} - \frac{p_{\text{neg}}}{1 + e^{(R-R_{\text{neg}})/\sigma_{\text{neg}}}}, \text{VOLL} \right) \right), \quad (18)$$

with base term $a = 75$ GBP/MWh, $k = 1.5$, $R_{\text{ref}} = 25$ GW; scarcity term $b = 350$ GBP/MWh, $R_h = 40$ GW, $\sigma = 4.5$ GW; and a logistic oversupply term of depth $p_{\text{neg}} = 65$ GBP/MWh centred at $R_{\text{neg}} = 6$ GW with width $\sigma_{\text{neg}} = 3$ GW that drives the price negative when non-dispatchable injection swamps demand. The inner min caps the economic curve at the value of lost load $\text{VOLL} = 10,000$ GBP/MWh; the outer clip applies the negative-price floor $\underline{p} = -50$ GBP/MWh and the $\bar{p}_{\text{cap}} = 2,000$ GBP/MWh administrative cap under which cleared prices are reported

(cap $\neq$ VOLL, Section 2); the welfare program of Section 4 integrates the unclipped economic curve; the cap and the VOLL ceiling are never reached on the realised operating range, while the -50 GBP/MWh floor binds only at the 17 oversupply half-hours reported below. Equation (18) is transcribed term-by-term from the implementation, and the shipped pipeline asserts on a fine grid that the printed form and the code agree and that both are non-decreasing everywhere and strictly increasing on the no-clip region (the `strict_monotone_supply_curve` invariant in `summary.json`), so Assumptions 2.1 and 4.1 are checked certificates rather than hypotheses. This calibration produces a no-BESS price duration curve with mean 84.0 GBP/MWh, P50 67.8 GBP/MWh, P95 213.6 GBP/MWh, and maximum 672.8 GBP/MWh; the oversupply tail yields 379 negative half-hours, of which 17 sit at the -50 GBP/MWh floor. Table 1 compares these statistics against realised GB half-hourly market prices for the twelve months to 28 August 2026 (Elexon Market Index Data, APX and N2EX volume-weighted (Elexon, 2026)): the price level is broadly consistent (synthetic mean 84.0 against realised 89.5 GBP/MWh), while the calibration deliberately retains the wider spread of its 2024–25 reference window (standard deviation 72.2 against 37.9; P95 213.6 against 145.0) — cannibalisation acts on the spread, so the calibration sits in the regime in which the effect is material, and the bias ratios rather than the revenue levels are the object of interest.

Table 1: Synthetic no-BESS calibration statistics against realised GB half-hourly market prices, twelve months to 28 August 2026 (Elexon Market Index Data, APX + N2EX volume-weighted, retrieved 29 August 2026). Prices in GBP/MWh; counts in half-hours (the realised window spans 17,515 half-hours against the synthetic 8,736). Source: `real_data_calibration.csv`.

Statistic	Synthetic (no-BESS)	Realised MID, 12m to Aug 2026
Mean	84.0	89.5
P50	67.8	90.3
P95	213.6	145.0
Maximum	672.8	560.8
Minimum	-50.0	-99.0
Standard deviation	72.2	37.9
Negative half-hours	379	428
At/below -50 GBP/MWh	17	13

Figure 1 plots the supply curve. Figure 2 shows the synthetic gross demand and non-dispatchable injection over the first two weeks; demand has the characteristic GB diurnal-plus-seasonal pattern with weekday-weekend differentiation, and renewables are anti-correlated with demand at synoptic and diurnal scales.

The fleet is parameterised by aggregate power capacity $P^{\max} \in \{0.5, 1, 2, 4, 6, 8, 10\}$ GW and duration $h \in \{1, 2, 4\}$ hours, with round-trip efficiency $\eta = 0.86$ and cyclic boundary $\sigma_0 = 0.5$. As of August 2026 the GB BESS fleet is approximately 5 GW with average duration just under 1.5 hours (Modo Energy, 2026), and NESO’s Clean Power 2030 analysis projects on the order of 25 GW of installed GB battery capacity by 2030 (central range 23–27 GW) (NESO, 2024, 2026). A clarification on scope is important here: the stylised single-zone sweep treats $P^{\max}$ as an *aggregate* price-influencing fleet acting on one national supply curve, and the cannibalisation bias becomes severe over a 4–8 GW band of *effective* price-setting capacity. This is *not* the same as the ~ 25 GW system-wide GB figure: in a real network, geographic dispersion, duration heterogeneity, ancillary commitments, and binding transmission mean that only a fraction of installed capacity competes on any one price node at any one time. The single-zone result should therefore be read as the cannibalisation of a representative price-setting cohort, not as a GB-wide revenue forecast for the full installed base; the GB-wide treatment requires the nodal extension (Section 7) and a full scenario set rather than this stylised sweep.




Figure 1: Stylised GB inverse supply curve $S(R)$. The base term $a(R/R_{\text{ref}})^k$ captures the slow rise from baseload through mid-merit gas; the scarcity term $b \exp((R - R_h)/\sigma)$ captures the sharp rise as the merit order is exhausted. Calibrated to a representative historical GB 2024–2025 price duration curve (reference window retained for comparability; see Table 1).




Figure 2: Synthetic gross demand, non-dispatchable injection (wind plus solar), and residual demand for the first two weeks of the simulation horizon. The residual demand is the input to the supply curve in equation (18). Anti-correlation between renewable injection and demand peak is qualitatively correct for the GB market.

6.2 Bias of the naive forecast

Table 2 reports the headline result: revenue per MW (annualised) under the naive (price-taker) forecast and under the cannibalisation-aware equilibrium forecast computed via Algorithm 1, together with the bias of the naive forecast. Each row is a fleet/duration pair; rows within a duration block share the same naive revenue (because the naive prices do not depend on fleet size — they are computed assuming zero BESS injection).

The headline two-hour figures are: at a 4 GW effective price-setting fleet the naive forecast

Table 2: Annualised BESS revenue per MW: naive (price-taker) vs cannibalisation-aware equilibrium forecast across fleet sizes and durations. *Bias* is the percentage by which the naive forecast overstates revenue; positive throughout, increasing in fleet size, and increasing in duration for fleets of 2 GW and above (below 2 GW the ordering in duration reverses). Source: case study reproducible from accompanying code (`bias_vs_fleet.csv`).

Fleet (GW)	Duration (h)	Naive rev. (£/MW/yr)	Equilibrium rev. (£/MW/yr)	Bias
0.5	1	69,680	58,219	+19.7%
1.0	1	69,680	50,729	+37.4%
2.0	1	69,680	41,142	+69.4%
4.0	1	69,680	31,163	+123.6%
6.0	1	69,680	25,711	+171.0%
8.0	1	69,680	21,949	+217.5%
10.0	1	69,680	18,987	+267.0%
0.5	2	112,176	95,601	+17.3%
1.0	2	112,176	83,169	+34.9%
2.0	2	112,176	65,450	+71.4%
4.0	2	112,176	44,957	+149.5%
6.0	2	112,176	33,341	+236.4%
8.0	2	112,176	25,926	+332.7%
10.0	2	112,176	20,788	+439.6%
0.5	4	169,080	145,418	+16.3%
1.0	4	169,080	125,914	+34.3%
2.0	4	169,080	95,764	+76.6%
4.0	4	169,080	56,901	+197.1%
6.0	4	169,080	35,879	+371.3%
8.0	4	169,080	23,977	+605.2%
10.0	4	169,080	17,134	+886.8%

overstates revenue by 150% (£44,957/MW/yr equilibrium against £112,176 naive), rising to 236% at 6 GW (£33,341), 333% at 8 GW (£25,926), and 440% at 10 GW (£20,788). Duration sharpens the effect at scale: at a fixed 6 GW fleet the bias is 171% for one-hour, 236% for two-hour, and 371% for four-hour batteries. The bias is monotone in fleet size, monotone in duration for fleets of 2 GW and above (below 2 GW the ordering reverses and longer duration is mildly protective), and unambiguously large across the representative band. Equilibrium revenue per MW collapses with fleet size — at 8 GW of two-hour BESS, equilibrium revenue is roughly one quarter of naive — because the cannibalisation effect compresses the price duration curve enough that the fleet competes for a much smaller spread. A caution on levels: the absolute per-MW figures are artefacts of the stylised calibration and of perfect foresight — on the realised GB price series for the twelve months to 28 August 2026, perfect-foresight wholesale-only arbitrage earns approximately £24,300 / £43,800 / £72,400 per MW-year for 1/2/4-hour batteries (computed from the Table 1 source series), well below the synthetic naive levels — so only the ratio B is intended to be interpreted.

Figure 3 plots the bias against fleet size for each duration. Figure 4 compares the per-MW revenue surfaces of the naive and equilibrium forecasts. The contrast is most stark for four-hour batteries: the naive forecast says they earn the most per MW (because they capture the full daily spread); the equilibrium says they cannibalise themselves the fastest (because their longer duration extracts more arbitrage per cycle, flattening prices more aggressively).




Figure 3: Bias of the naive (price-taker) revenue forecast as a function of fleet size, by storage duration. Positive bias means the naive forecast overstates revenue. Bias is monotone in fleet size, and monotone in duration for fleets of 2 GW and above (below 2 GW the ordering reverses). For two-hour batteries the naive overstatement is 150% at a 4 GW effective price-setting fleet, 236% at 6 GW, and 333% at 8 GW; the 4–8 GW band is the one in which the widely quoted “100–300%+” figures arise. As discussed in the text, this aggregate single-zone capacity is not the ~ 25 GW system-wide GB installed base projected for 2030.




Figure 4: Annual revenue per MW (k£/MW/yr) under the naive (price-taker) forecast (left) and the cannibalisation-aware equilibrium forecast (right), for fleet sizes 0.5–10 GW and durations 1, 2, 4 hours. The naive forecast is independent of fleet size by construction; the equilibrium forecast falls steeply as fleet size grows. In equilibrium, 4-hour BESS becomes worse than 2-hour between 6 and 8 GW, and worse than 1-hour between 8 and 10 GW.

6.3 Equilibrium price duration curve

Figure 5 compares the price duration curves implied by the naive and equilibrium forecasts at a fleet of 6 GW, 2 hour. The equilibrium price duration curve is materially compressed: peaks are flattened (storage discharges into them) and troughs are lifted (storage charges from them). The shaded regions highlight the bias regimes — where naive overstates (peak hours) and where

naive understates (trough hours). The arbitrage opportunity that the naive forecast assumes is significantly smaller in equilibrium than in the no-fleet price profile.




Figure 5: Price duration curves: naive (no-BESS feedback) versus cannibalisation-aware equilibrium with a 6 GW, 2-hour BESS fleet. The equilibrium peaks are flattened and the troughs lifted, compressing the daily spread. The naive arbitrage opportunity (the area between the price duration curve and a flat reference) is materially overstated.

6.4 Convergence behaviour of iteration schemes

Figure 6 plots the L^∞ convergence residual $\|\pi^{(k+1)} - \pi^{(k)}\|_\infty$ over outer iterations for five schemes applied to the same 4 GW, 2-hour problem: relaxed Picard on the *unregularised* dispatch operator with damping $\alpha = 0.3$ and $\alpha = 0.6$, and — on the regularised map — plain Picard, Anderson acceleration with depth $m = 4$, and μ -regularised Krasnoselski–Mann averaging. To keep the 21-row sweep runtime bounded, the per-row convergence certificate is computed on a representative 4-week window of 1,344 half-hours; the headline revenues in Table 2 use the full 8,736 half-hours, and the full-horizon certificate behaves the same way (KM reaches tolerance in 14 iterations against 11 on the window, completing in under 6 s within 280 MB). The behaviour is unambiguous. Raw Picard does not converge at the 0.05 GBP/MWh tolerance under either damping: with $\alpha = 0.6$ it oscillates between roughly 240 and 300 GBP/MWh, and with $\alpha = 0.3$ it stagnates on a floor in the tens of GBP/MWh (minimum ≈ 46 , never reaching tolerance in 40 iterations). On the regularised map *every* scheme converges: Anderson reaches the tolerance in 6 iterations, relaxed Picard ($\alpha = 0.6$) in 7, plain Picard in 9, constant-step KM ($\lambda = 1/2$) in 9, and harmonic-step KM in 11. The comparison isolates the mechanism: *regularisation, not averaging, is the enabling device* — once $\mathcal{F}_\mu$ is a contraction, Banach applies and plain Picard already converges geometrically. Anderson is fastest but lacks a general convergence theorem on this class; plain Picard is the fastest scheme *with* a guarantee while the contraction margin is comfortable. We deploy KM in production because its guarantee (Theorem 5.1) is the one that survives at the non-expansive boundary $\mu \approx L_S$, where Banach no longer applies — and a μ -sweep confirms the regime is real: as μ is lowered toward the realised slope, plain Picard degrades from 9 to 19 iterations while the averaged schemes hold up (16 for harmonic KM, 10 for constant-step KM; `mu_sweep.csv`) — with the empirical contraction ratio monitored in production to keep the iteration inside the certified regime (Remark 5.4). Across the full sweep, μ -regularised KM converges at every fleet size and duration in 10–19 iterations (certificate diagnostics in `bias_vs_fleet.csv`; Table 3).




Figure 6: Convergence of five fixed-point iteration schemes on the same 4 GW, 2-hour problem (4-week certificate window; regularised schemes use the certificate μ , chosen above the realised supply-curve slope). Relaxed Picard on the raw dispatch operator fails at the 0.05 GBP/MWh tolerance under both dampings: $\alpha = 0.6$ oscillates, $\alpha = 0.3$ stagnates in the tens of GBP/MWh. On the regularised map every scheme converges — Anderson (depth $m = 4$) in 6 iterations, plain Picard in 9, KM in 11: regularisation, not averaging, is what makes the iteration converge. KM is the production scheme for the guarantee that survives at the non-expansive boundary; its advantage over raw Picard is consistent across all fleet sizes considered.

Table 3 reports the irreducible residual floor of relaxed Picard (best damping per cell over $\alpha \in \{0.1, \dots, 0.9\}$, 200 iterations, 4-week window) against fleet size, alongside the iteration count of μ -regularised KM at the same investment-grade tolerance. Picard’s floor grows monotonically with penetration, crossing 1 GBP/MWh by 0.5 GW and exceeding 20 GBP/MWh beyond about 6 GW (the floors grow further on the full 26-week horizon), while KM reaches 0.05 GBP/MWh in 10–19 iterations throughout. At the tolerance an investment-grade revenue number requires, Picard converges at no fleet size; the contraction safeguard is what makes the computation well-posed.

6.5 Translation to project finance

The bias estimates above translate directly into project-finance metrics. The following worked figures are *illustrative, not modelled*: they sketch the sensitivity of a standard financing structure to a revenue overstatement, and are not an output of the case study. A typical UK BESS project closing today targets a senior-debt DSCR of $1.40\times$ against a P50 revenue case and $1.10\times$ against a P90 case, with a 50/50 debt/equity split, 7.5% gross debt cost, and a 15-year term. Under these standard terms, a revenue overstatement of 30% would reduce realised DSCR to approximately $1.05\times$ (a covenant breach in most facilities) and reduce equity IRR from a target 11% to approximately -3% — equity wipeout in the bank case. At the 100%+ overstatements implied by the representative 2030 fleet band, naive-forecast-financed projects would default within months of commissioning. In the CECadence product (in private beta), we do not report a single point estimate of this kind: it reports the *minimum* DSCR over the financing horizon together with the year in which it binds, against a P50 target of $1.40\times$ and a P90 lender covenant of $1.10\times$, so that the tightest-coverage year is surfaced rather than averaged into a headline.

These calculations are robust to a wide range of assumptions because the bias is a structural feature of the forecast methodology, not of any particular project’s cost structure. The cleanest

Table 3: Irreducible residual floor of relaxed Picard (raw dispatch operator, two-hour fleets; best damping per cell over $\alpha \in \{0.1, \dots, 0.9\}$, 200 iterations, 4-week window — floors grow further on the full 26-week horizon) versus convergence of μ -regularised KM at the 0.05 GBP/MWh tolerance. Picard never reaches the investment-grade tolerance at any fleet size, duration, damping or horizon tested; its floor (the smallest residual it attains) grows with fleet. KM reaches tolerance at every fleet size; the iteration count is the two-hour full-sweep certificate at tolerance 0.05 GBP/MWh. Sources: `picard_floor.csv` (floor column), `bias_vs_fleet.csv` (KM column), `picard_threshold.csv` (raw-Picard control).

Fleet (GW)	Picard floor (£/MWh, best α)	Picard reaches 0.05 £/MWh?	KM iterations to tol.
0.1	0.17	no	—
0.25	0.39	no	—
0.5	1.05	no	16
1.0	2.89	no	13
2.0	7.10	no	14
4.0	13.00	no	11
6.0	18.95	no	12
8.0	24.21	no	12
10.0	29.71	no	12

mitigation is to forecast on the equilibrium prices computed via Algorithm 1, not the naive prices.

Practitioner sidebar

The asymmetry favours the cautious modeller. If you under-forecast revenue by 30% and operations beat expectations, you can refinance into a more aggressive package after one year of operating data. If you over-forecast by 30% and operations miss expectations, you breach DSCR covenants in the first year and lose equity. The optionality is asymmetric. This is why we view cannibalisation-aware forecasting as a defensiveness measure for the project sponsor and the lender, not just an academic refinement.

7 Discussion and extensions

7.1 What the stylised model omits

The case study makes several deliberate simplifications that we now address.

Single zone. Real markets have transmission constraints; when binding, they create locational marginal prices that differ across nodes. The fixed-point structure carries over directly to nodal price formation (a planned extension of the Compounding Energy stack), and the bias result strengthens in nodal markets because constraint-bound nodes have steeper effective supply curves and therefore stronger cannibalisation.

Ancillary stack. Real BESS revenue stacks energy arbitrage with frequency response (Dynamic Containment, Dynamic Moderation, Dynamic Regulation in GB, each with Low and High variants), Capacity Market, and Balancing Mechanism. Ancillary cannibalisation has its own saturation structure — frequency response markets saturate when BESS deployment exceeds the system requirement — and composes with the wholesale equilibrium; the multi-market formulation is developed in the companion working paper CE-WP-2026-02. The stylised case study here is a wholesale-energy-only special case, and a single illustrative calibration rather than a

scenario set spanning a range of 2030 BESS build-outs and gas-price paths, which is the basis for any committed revenue number.

Perfect foresight. The dispatch LP (3) assumes perfect foresight over the horizon. Real BESS operators dispatch under price uncertainty, typically via day-ahead optimisation followed by within-day rolling re-optimisation against forecast-updates. Stochastic extension is straightforward via scenario decomposition (Birge and Louveaux, 2011): replace the deterministic LP by a multistage stochastic LP solved via SDDP or Lagrangian decomposition, and recover the equilibrium fixed point by averaging across scenarios. Quantitatively, the bias of the naive forecast is robust to the foresight assumption; the absolute revenue numbers shift slightly because of the value of imperfect optionality.

Homogeneous fleet. We treat the fleet as a single representative asset. In reality the fleet is heterogeneous in duration, efficiency, location, and contracting structure. The fixed-point framework extends directly to heterogeneous fleets: $\mathbf{b}$ becomes a fleet of distinct dispatch profiles indexed by asset class, and the dispatch operator is solved per asset class. The aggregate cannibalisation effect is approximately the duration-weighted average of the per-class effects; the asymmetry is that longer-duration assets cannibalise themselves *and* shorter-duration assets, which is captured naturally in the multi-class formulation.

Degradation and round-trip efficiency dynamics. Our model treats η as fixed and the degradation hook at its default of zero. In practice round-trip throughput carries a marginal degradation cost (Wankmüller et al., 2017), which we model as a per-MWh throughput charge entering the dispatch objective. As a robustness check, a £2/MWh throughput cost *raises* equilibrium revenue by 2.5% at the 4 GW two-hour fleet, 3.8% at 6 GW, and 5.5% at 8 GW relative to the zero-cost case (a full sweep over costs £0.5–£5/MWh and fleets 0.5–10 GW ships as `degradation_sensitivity.csv`; the effect grows with penetration, reaching +14.7% at 10 GW under a £5/MWh cost, and turns slightly negative only for very small fleets under aggressive costs): the charge discourages low-value cycling, the fleet competes less aggressively for thin spreads, and the resulting reduction in self-cannibalisation more than offsets the direct cost. The cannibalisation effect and the qualitative bias result are therefore robust to plausible degradation costs; if anything, ignoring degradation modestly overstates the bias.

7.2 Strategic vs price-taker storage

We have assumed price-taking behaviour by all storage operators. In a market dominated by a few large operators, strategic withholding can shift the equilibrium toward the Cournot solution (Hobbs, 2001; Schill and Kemfert, 2011), which sits between the price-taker equilibrium considered here and the social-planner solution that an integrated operator would compute. The price-taker equilibrium is the appropriate benchmark for the GB market where deployment is fragmented across many independent operators; markets with concentrated ownership (e.g. early stages of grid-scale BESS in some U.S. ISOs) require the strategic extension.

7.3 Capacity expansion: the meta-fixed-point

The case study takes the BESS fleet as given. In a capacity expansion problem the fleet itself is endogenous, determined by entry and exit decisions that respond to forecast revenue. A consistent capacity-expansion-cum-dispatch equilibrium is a meta-fixed-point: the entry-and-exit operator selects fleets that earn at least their cost of capital under the dispatch equilibrium, and the dispatch equilibrium is computed for that fleet. The bias of naive capacity expansion forecasting — using naive revenue forecasts to choose fleet size — compounds the dispatch-level bias and is materially larger over a 2027–2040 expansion horizon.

8 Conclusion

We have framed the BESS revenue forecasting problem as a fixed-point problem in price-quantity space, shown that the composite operator is discontinuous so the fixed point cannot be solved pointwise, and established existence and uniqueness through an equivalent strictly convex welfare program whose dual is the equilibrium price. We identified μ -regularisation of the dispatch operator as the device that makes the fixed-point computation convergent — demonstrating that raw Picard fails at investment-grade tolerance at every fleet size, while on the regularised map plain Picard, Krasnoselski–Mann averaging (the production scheme, whose guarantee survives at the non-expansive boundary), and Anderson acceleration all converge — and quantified the bias of the prevailing naive forecasting methodology against a stylised GB calibration. The bias is structural, monotone in fleet size (and in duration for fleets of 2 GW and above), and large enough across the representative 2030 fleet band to invalidate any project finance package built on naive forecasts.

Three implications follow. First, project sponsors and lenders evaluating new BESS projects should require revenue forecasts explicitly stated as cannibalisation-aware equilibrium forecasts, with the underlying fleet and price-formation assumptions disclosed. Second, the methodological contribution of this paper — the convex welfare characterisation of the equilibrium with μ -regularised KM convergence — generalises straightforwardly to nodal markets, multi-market revenue stacks, and stochastic operational settings; the implementation forms the analytical core of the Compounding Energy CECadence product (wholesale and stacked BESS revenue, in private beta), with a nodal extension planned. Third, the same framework underpins capacity-expansion-cum-dispatch equilibrium computations relevant for system planning at the regulator level (NESO, OFGEM) and the investor level (infrastructure funds, utility planners).

The bigger point is methodological. As deployment of dispatchable resources — storage, demand response, controllable renewables, hydrogen — crosses thresholds at which they materially influence prices, the analytical machinery used to forecast their revenue must transition from price-taker single-asset optimisation to equilibrium fixed-point computation. This is not an incremental improvement; it is a regime change. The defensibility of capital allocation decisions in the energy transition depends on whether the modelling layer keeps pace.

Acknowledgements

This work develops material first prototyped in internal working notes. The case study is fully reproducible from the Python implementation accompanying this paper, available on request from the corresponding author. Comments are welcome to christopher@compoundingenergy.com.

References

- Donald G. Anderson. Iterative procedures for nonlinear integral equations. *Journal of the ACM*, 12(4):547–560, 1965. doi: 10.1145/321296.321305.
- Aurora Energy Research. GB battery storage strategic outlook. Aurora Energy Research market report, 2025.
- Jean-Bernard Baillon and Georges Haddad. Quelques propriétés des opérateurs angle-bornés et n -cycliquement monotones. *Israel Journal of Mathematics*, 26:137–150, 1977. doi: 10.1007/BF03007664.
- Stefan Banach. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3:133–181, 1922.

- Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer, 2nd edition, 2017.
- John R. Birge and François Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer, 2nd edition, 2011.
- BloombergNEF. Energy storage market outlook 2026. Bloomberg New Energy Finance, 2026.
- Kyle Bradbury, Lincoln Pratson, and Dalia Patiño-Echeverri. Economic viability of energy storage systems based on price arbitrage potential in real-time U.S. electricity markets. *Applied Energy*, 114:512–519, 2014. doi: 10.1016/j.apenergy.2013.10.010.
- 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.
- Elexon. Market index data (MID), BMRS insights API. <https://data.elexon.co.uk>, 2026. Half-hourly APX and N2EX market index prices, volume-weighted; window 29 August 2025 – 28 August 2026, retrieved 29 August 2026.
- Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer Series in Operations Research. Springer, 2003.
- Bethany Frew, Wesley Cole, Paul Denholm, A. Will Frazier, Nina Vincent, and Robert Margolis. Sunny with a chance of curtailment: Operating the U.S. grid with very high levels of solar photovoltaics. *iScience*, 24(6):102609, 2021. doi: 10.1016/j.isci.2021.102609.
- Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.
- Andrzej Granas and James Dugundji. *Fixed Point Theory*. Springer Monographs in Mathematics. Springer, 2003.
- Lion Hirth. The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy Economics*, 38:218–236, 2013. doi: 10.1016/j.eneco.2013.02.004.
- Benjamin F. Hobbs. Linear complementarity models of Nash-Cournot competition in bilateral and POOLCO power markets. *IEEE Transactions on Power Systems*, 16(2):194–202, 2001. doi: 10.1109/59.918286.
- Mark A. Krasnoselskiĭ. Two remarks on the method of successive approximations. *Uspekhi Matematicheskikh Nauk*, 10(1):123–127, 1955. (In Russian).
- LCP Delta. GB storage market outlook. LCP Delta market report, 2026.
- Javier López Prol, Karl W. Steininger, and David Zilberman. The cannibalization effect of wind and solar in the Californian wholesale electricity market. *Energy Economics*, 85:104552, 2020. doi: 10.1016/j.eneco.2019.104552.
- W. Robert Mann. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 4:506–510, 1953. doi: 10.1090/S0002-9939-1953-0054846-3.
- Brandon Mauch, Pedro M. S. Carvalho, and Jay Apt. Can a wind farm with CAES survive in the day-ahead market? *Energy Policy*, 48:584–593, 2012. doi: 10.1016/j.enpol.2012.05.061.
- Modo Energy. GB battery energy storage index: 2025 retrospective. Modo Energy industry report, 2026. Retrieved August 2026.

- NESO. Clean Power 2030: Advice on achieving clean power for Great Britain by 2030. National Energy System Operator, 2024.
- 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.
- Wolf-Peter Schill and Claudia Kemfert. Modeling strategic electricity storage: The case of pumped hydro storage in Germany. *The Energy Journal*, 32(3):59–87, 2011. doi: 10.5547/ISSN0195-6574-EJ-Vol32-No3-3.
- R. Sioshansi, P. Denholm, T. Jenkin, and J. Weiss. Estimating the value of electricity storage in PJM: arbitrage and some welfare effects. *Energy Economics*, 31(2):269–277, 2009.
- Ramteen Sioshansi, Paul Denholm, and Thomas Jenkin. A comparative analysis of the value of pure and hybrid electricity storage. *Energy Economics*, 33(1):56–66, 2011. doi: 10.1016/j.eneco.2010.06.004.
- Alex Toth and C. T. Kelley. Convergence analysis for Anderson acceleration. *SIAM Journal on Numerical Analysis*, 53(2):805–819, 2015. doi: 10.1137/130919398.
- Rahul Walawalkar, Jay Apt, and Rick Mancini. Economics of electric energy storage for energy arbitrage and regulation in New York. *Energy Policy*, 35(4):2558–2568, 2007. doi: 10.1016/j.enpol.2006.09.005.
- Homer F. Walker and Peng Ni. Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis*, 49(4):1715–1735, 2011. doi: 10.1137/10078356X.
- Florian Wankmüller, Prakash R. Thimmapuram, Brian P. Gallachóir, and Audun Botterud. Impact of battery degradation on energy arbitrage revenue of grid-level energy storage. *Journal of Energy Storage*, 10:56–66, 2017. doi: 10.1016/j.est.2016.12.004.

A Notation glossary

Symbol	Meaning
T	Number of half-hours in horizon (case study: $T = 8736$)
Δt	Half-hour step length, 0.5 hours
D	Gross demand profile, GW per half-hour
R	Non-dispatchable injection (wind plus solar), GW
b	Aggregate BESS net injection (+ charging, – discharging), GW
$\rho_t(b)$	Residual demand at time t , equation (1)
$\pi_t(b)$	Wholesale price at time t , equation (2)
$S(\cdot)$	Inverse supply curve, equation (18)
R_t^0	No-BESS residual demand $D_t - R_t$, GW
VOLL	Value of lost load, 10,000 GBP/MWh (economic scarcity ceiling)
$P^{\max}$	Aggregate fleet power capacity, GW
h	Storage duration, hours; $E^{\max} = hP^{\max}$
η	Round-trip efficiency; $\eta_c = \eta_d = \sqrt{\eta}$
σ_0	Cyclic boundary state-of-charge fraction
$\mathcal{P}$	Price-formation operator, equation (5)
$\mathcal{D}$	Dispatch operator (LP), equation (4)
$\mathcal{D}_\mu$	Regularised dispatch operator (strictly convex QP), Section 5.3
$\mathcal{F}, \mathcal{F}_\mu$	Composite operators $\mathcal{D} \circ \mathcal{P}$ and $\mathcal{D}_\mu \circ \mathcal{P}$
Φ	Welfare potential, $\Phi' = S$, equation (12)
μ	Quadratic dispatch regularisation parameter (chosen above the realised supply-curve slope in the certificate)
b^*, π^*	Equilibrium net injection and price (welfare minimiser, Theorem 4.2)
$\Pi^*, \Pi^{\text{naive}}$	Equilibrium and naive revenue, equations (9) and (8)
B	Bias of naive forecast: $(\Pi^{\text{naive}} - \Pi^*)/\Pi^*$
λ_k	KM step size, $\lambda_k = 2/(k + 2)$

B Reproducibility

The case study reported in Section 6 is reproducible end-to-end from the Python implementation accompanying this paper, available on request from the corresponding author. The implementation requires Python 3.11+, NumPy, SciPy (HiGHS via `linprog` for the dispatch LP), Pandas, Matplotlib, OSQP (regularised dispatch QP), and CVXPY with Clarabel and SCS (equilibrium welfare program). The random seed is fixed at 20,260,901 throughout, so all numerical results are reproducible to the precision reported here. The headline revenue sweep runs over the full $T = 8,736$ half-hours; the per-row convergence certificate (Figure 6) runs over a representative 4-week window of 1,344 half-hours to keep the 21-row sweep runtime bounded — the quadratic term is sparse, and the full-horizon certificate completes in under 6 s within 280 MB with the same qualitative behaviour (14 KM iterations against 11 on the window). The script is organised into stages, each runnable independently: `setup`, `sweep`, `convergence`, `picard`, `pfloor`, `price`, `degradation`, `schedules`, `mu_sweep`, `fullyear`, and `summary`. Output figures appear as PDFs in `figures/`, numerical results as CSVs and a machine-readable `summary.json` in `results/`.

Reproducibility key. We tag every committed result with the triple of `scenario_id`, `code_version`, and `timestamp`. For the run reported here:

- `scenario_id`: CE-WP-2026-01-stylised-GB-2026-09
- `code_version`: fcaf71bbdb8c5223...69b201cc14 (SHA-256 of the source tree; full digest)

in `summary.json`)

- `timestamp: 2026-08-29T14:50:14Z`

The same `summary.json` records the power-systems invariant checks (SoC bounded, no simultaneous charge/discharge in the reported dispatch, price within $[p, \bar{p}_{\text{cap}}]$ with VOLL ceiling, negative-price floor reachable, and the strict-monotone-supply-curve check that the printed equation (18) and the implementation agree and satisfy Assumptions 2.1 and 4.1 on a fine grid), the welfare-relaxation dual identity (price-taker profit gap below 6.5×10^{-7} across the sweep on the relaxation dispatch), and the price-taker/KKT consistency gap on the *reported* dispatch: between 2.7×10^{-4} and 1.9×10^{-2} across the sweep, and below 7.4×10^{-3} in the two-hour block ($\varepsilon_{\text{cert}} = 2 \times 10^{-2}$ in Algorithm 1).

Simultaneity reconstruction. The welfare relaxation admits simultaneous charge and discharge: it is objective-neutral there (simultaneity nets out in the residual while destroying round-trip energy, which has value at negative prices), and the convex program returns it at up to 589 of 8,736 half-hours, at up to the full fleet power, depending on the sweep cell. We therefore reconstruct the unique SoC-preserving non-simultaneous dispatch: the SoC trajectory e_t is retained exactly, and each half-hour is assigned charge-only or discharge-only according to the sign of its SoC change. The reconstruction shifts total revenue by at most 0.79% across the sweep (at most 0.24% in the two-hour block; -0.03% at the 6 GW two-hour reference), and every figure in Table 2 is reported *from the reconstructed dispatch*, as is the price-taker consistency gap quoted above — the certificate is computed on the dispatch the paper reports, not on the relaxation iterate.