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

June 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 $\mathcal{D}$ — and prove existence under standard assumptions on the supply curve and the dispatch problem. The naive Picard iteration $b_{k+1} = (\mathcal{D} \circ \mathcal{P})(b_k)$ generically fails to converge in markets with steep scarcity tails; we show that Krasnoselski–Mann averaging restores convergence with explicit rate. A stylised half-yearly Great Britain case study calibrated to 2024–2025 day-ahead conditions quantifies the bias of the naive (price-taker) revenue forecast: at fleet penetrations representative of NESO's 2030 trajectory the naive forecast overstates revenue by 100–300% for two-hour batteries and by 200–800% for four-hour batteries. The bias is monotone in fleet size and in storage duration. We discuss extensions to nodal markets, ancillary stacking, and risk-averse formulations, and locate the algorithm in the Compounding Energy production stack.

Keywords: battery energy storage, cannibalisation, fixed-point iteration, variational inequality,

electricity markets, GB, Krasnoselski–Mann iteration

Contents

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 0.4 GW of installed BESS in 2020 to over 5 GW in 2026, with NESO's most recent Future Energy Scenarios projecting 15–20 GW by 2030 and 30–40 GW by 2035 (NESO, 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 (Section 3). Second, we give sufficient conditions on the supply curve and dispatch problem under which the fixed point exists and is essentially unique (Section 4); these conditions are physically plausible and satisfied in the GB market. Third, we exhibit the failure of naive Picard iteration in markets with steep scarcity tails, and show that Krasnoselski–Mann (KM) averaging restores convergence under a non-expansive condition that is also met in practice (Section 5). Fourth, we quantify the bias of the naive forecast against the equilibrium forecast in a stylised GB market calibrated to 2024–2025 day-ahead conditions (Section 6). The headline numerical result is that under fleet penetrations that NESO and BloombergNEF expect by 2030, the naive forecast overstates revenue by 100–300% for two-hour batteries and by 200–800% for four-hour batteries, with the bias monotone in fleet size.

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. A BESS project finance package typically targets debt-service-coverage ratio (DSCR) above 1.4× over the bank case. If the bank case revenue is overstated by 30% — the order of magnitude we estimate for early-2026 GB conditions — 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 will 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.

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 by Modo Energy, LCP Delta, and Aurora reports, 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 $\mathcal{T} = \{1, 2, \ldots, T\}$ of half-hours; the case study uses $T = 8736$ (twenty-six weeks) and the conclusions extrapolate by linearity to longer horizons. 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 $t \in \{1, \ldots, T\}$ 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 $\boldsymbol{D} \in \mathbb{R}^T_+$ with $D_t \geq 0$ in GW. The aggregate non-dispatchable injection from variable renewables and inflexible thermal is $\boldsymbol{R} \in \mathbb{R}^T_+$. The aggregate net injection of the BESS fleet is $\boldsymbol{b} \in \mathbb{R}^T$, with the convention $b_t > 0$ when the fleet is net charging (drawing from the grid) and $b_t < 0$ when net discharging. The *residual demand* faced by dispatchable generation is therefore

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

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

$$\pi_t(\boldsymbol{b}) = S(\rho_t(\boldsymbol{b})). \tag{2}$$

We assume throughout:

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

The bounds $\underline{p}$ and $\overline{p}$ correspond to the negative-pricing floor (set in GB by curtailment compensation and CCGT minimum-load economics, around -50 GBP/MWh) and the market price cap (2000 GBP/MWh on N2EX intraday). 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 $\eta \in (0, 1]$. 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 $\boldsymbol{\pi} \in \mathbb{R}^T$, the fleet's optimal dispatch is the solution of the linear program

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

where e_0 is the implicit initial state-of-charge and $e_T \ge \sigma_0 E^{\max}$ 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 two formulations therefore have identical optimal revenue 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 $\boldsymbol{c} = \boldsymbol{d} = \boldsymbol{0}$, $e_t = \sigma_0 E^{\max}$ is feasible for all $\boldsymbol{\pi}$). Define the dispatch operator on net injection $b_t := c_t - d_t$:

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

$\mathcal{D}$ 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 $\mathcal{D}(\boldsymbol{\pi})$ denotes a single point of $\mathbb{R}^T$.

Remark 2.2. Equation (3) describes a competitive price-taker fleet: the optimisation treats $\boldsymbol{\pi}$ 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

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

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

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

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

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:

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

and the naive revenue forecast is

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

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

$$\Pi^* = \sum_t \pi_t^* \left(-b_t^*\right) \Delta t. \tag{9}$$

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 $\mathcal{P} : \mathbb{R}^T \to \mathbb{R}^T$ defined by (5) is Lipschitz continuous with constant L_S:*

$$\|\mathcal{P}(\boldsymbol{b}_1) - \mathcal{P}(\boldsymbol{b}_2)\|_\infty \le L_S \, \|\boldsymbol{b}_1 - \boldsymbol{b}_2\|_\infty.$$

Proof. Componentwise, $|\pi_t(\boldsymbol{b}_1) - \pi_t(\boldsymbol{b}_2)| = |S(\rho_t(\boldsymbol{b}_1)) - S(\rho_t(\boldsymbol{b}_2))|$. 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 $|S(x) - S(y)| \le L_S|x - y|$ holds. Since $\rho_t(\boldsymbol{b}_1) - \rho_t(\boldsymbol{b}_2) = b_{1,t} - b_{2,t}$, we obtain $|\pi_t(\boldsymbol{b}_1) - \pi_t(\boldsymbol{b}_2)| \le L_S|b_{1,t} - b_{2,t}|$. Taking the maximum over t gives the stated inequality. □

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

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

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

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

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

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 $\{(\boldsymbol{\pi}, \boldsymbol{b}) : \boldsymbol{b} \in \mathcal{D}(\boldsymbol{\pi})\}$ 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 $v(\boldsymbol{\pi}) = \min_{\boldsymbol{b} \in K_b} \langle \boldsymbol{\pi}, \boldsymbol{b} \rangle_{\Delta t}$ is finite (since K_D is bounded), and the optimal-set correspondence $\boldsymbol{\pi} \mapsto \mathcal{D}(\boldsymbol{\pi}) = \{\boldsymbol{b} \in K_b : \langle \boldsymbol{\pi}, \boldsymbol{b} \rangle_{\Delta t} = v(\boldsymbol{\pi})\}$ is the intersection of K_b with a closed half-space whose boundary depends continuously on $\boldsymbol{\pi}$. The standard maximum theorem (Berge, see Facchinei and Pang 2003, Thm. 5.10) and the polyhedral-LP parametric stability theorem (Robinson, see Facchinei and Pang 2003, Thm. 7.1.2) both apply and give upper hemicontinuity. Convexity and compactness of $\mathcal{D}(\boldsymbol{\pi})$ are immediate from the LP structure. □

When the lexicographic selection rule of Section 2 is applied, $\mathcal{D}$ becomes a single-valued map $\mathbb{R}^T \to \mathbb{R}^T$. 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 $\mathcal{F}(\boldsymbol{b}) := \mathcal{D}(\mathcal{P}(\boldsymbol{b}))$ has at least one fixed point in $K = [-P^{\max}, P^{\max}]^T$. Equivalently, there exists $\boldsymbol{b}^* \in K$ and $\boldsymbol{\pi}^* = \mathcal{P}(\boldsymbol{b}^*) \in \mathbb{R}^T$ such that $\boldsymbol{b}^* \in \mathcal{D}(\boldsymbol{\pi}^*)$.*

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 $\boldsymbol{b} \in K$, $\mathcal{P}(\boldsymbol{b}) \in \mathbb{R}^T$ is well-defined (Lemma 3.1). For any $\boldsymbol{\pi} \in \mathbb{R}^T$, $\mathcal{D}(\boldsymbol{\pi}) \subset K$ (Lemma 3.3). Hence $\mathcal{F}(\boldsymbol{b}) \subset K$ 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 $(\boldsymbol{b}_n, \boldsymbol{x}_n) \to (\boldsymbol{b}, \boldsymbol{x})$ with $\boldsymbol{x}_n \in \mathcal{F}(\boldsymbol{b}_n) = \mathcal{D}(\mathcal{P}(\boldsymbol{b}_n))$, then $\mathcal{P}(\boldsymbol{b}_n) \to \mathcal{P}(\boldsymbol{b})$ by continuity, and $\boldsymbol{x}_n \in \mathcal{D}(\mathcal{P}(\boldsymbol{b}_n))$ together with $\boldsymbol{x}_n \to \boldsymbol{x}$ implies $\boldsymbol{x} \in \mathcal{D}(\mathcal{P}(\boldsymbol{b})) = \mathcal{F}(\boldsymbol{b})$ by closed-graph of $\mathcal{D}$.

Step 4 (apply Kakutani). The conditions of the Kakutani fixed-point theorem (Granas and Dugundji, 2003, Thm. 3.2) 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 $\boldsymbol{b}^* \in K$ with $\boldsymbol{b}^* \in \mathcal{F}(\boldsymbol{b}^*)$.

For the single-valued lexicographic selection of $\mathcal{D}$, the same conclusion follows from Brouwer (Granas and Dugundji 2003, Thm. 3.1) 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. □

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

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

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. Define the projection $\mathcal{B} : \mathbb{R}^{3T} \to \mathbb{R}^T$, $\mathcal{B}(\boldsymbol{c}, \boldsymbol{d}, \boldsymbol{e}) = \boldsymbol{c} - \boldsymbol{d}$, and let $K_b := \mathcal{B}(K_D)$ denote the projected feasible set of net-injection profiles. By eliminating $\boldsymbol{e}$ 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 $\boldsymbol{b} = \boldsymbol{c} - \boldsymbol{d}$ is $\sum_t \pi_t(d_t - c_t)\Delta t = -\langle \boldsymbol{\pi}, \boldsymbol{b} \rangle_{\Delta t}$, where $\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{\Delta t} := \sum_t x_t y_t \Delta t$. Since $\mathcal{D}(\boldsymbol{\pi})$ minimises $\langle \boldsymbol{\pi}, \boldsymbol{b} \rangle_{\Delta t}$ over $\boldsymbol{b} \in K_b$, the standard linear-programming first-order condition gives

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

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

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

The pair $(\boldsymbol{b}^*, \boldsymbol{\pi}^*) = (\boldsymbol{b}^*, \mathcal{P}(\boldsymbol{b}^*))$ solves (11) if and only if $\boldsymbol{b}^*$ 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 $\boldsymbol{b} \mapsto \mathcal{P}(\boldsymbol{b})$ is monotone — which, by Lemma 3.2 below, 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. We now consider uniqueness, which is more delicate. The dispatch LP has potentially many optima — this is generic when the price profile contains repeated values — and naive iteration may chatter between them.

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

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.

Proposition 4.2 (Uniqueness of equilibrium net injection at the fleet level)**.** *Under Assumptions 2.1 and 4.1, if $\boldsymbol{b}_1^*, \boldsymbol{b}_2^* \in K$ 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 $\boldsymbol{\pi}^* = \mathcal{P}(\boldsymbol{b}^*)$ is unique on the no-clip portion of the time horizon.*

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. $\langle \mathcal{P}(\boldsymbol{b}_1) - \mathcal{P}(\boldsymbol{b}_2), \boldsymbol{b}_1 - \boldsymbol{b}_2 \rangle_{\Delta t} > 0$ whenever $\boldsymbol{b}_1$ and $\boldsymbol{b}_2$ differ at a no-clip half-hour. From the VI (11) applied at both equilibria,

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

Adding the two inequalities gives $\langle \mathcal{P}(\boldsymbol{b}_1^*) - \mathcal{P}(\boldsymbol{b}_2^*), \boldsymbol{b}_2^* - \boldsymbol{b}_1^* \rangle_{\Delta t} \geq 0$, i.e. $\langle \mathcal{P}(\boldsymbol{b}_1^*) - \mathcal{P}(\boldsymbol{b}_2^*), \boldsymbol{b}_1^* - \boldsymbol{b}_2^* \rangle_{\Delta t} \leq 0$. 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 $\boldsymbol{\pi}^*$ on the no-clip portion follows from $\boldsymbol{\pi}^* = \mathcal{P}(\boldsymbol{b}^*)$ and the value of S at the same residual. □

Remark 4.3 (Uniqueness of fleet revenue). The aggregate fleet revenue $\Pi^* = -\langle \boldsymbol{\pi}^*, \boldsymbol{b}^* \rangle_{\Delta t}$ 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 $\boldsymbol{\pi}^*$, every selection $\boldsymbol{b}^* \in \mathcal{D}(\boldsymbol{\pi}^*)$ yields the same value $\langle \boldsymbol{\pi}^*, \boldsymbol{b}^* \rangle_{\Delta t}$ by the LP optimality of all elements of $\mathcal{D}(\boldsymbol{\pi}^*)$. Practically, what the BESS revenue forecaster needs — the fleet revenue — is therefore a well-defined function of fleet specification, regardless of LP degeneracy.

Remark 4.4 (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 4.2 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 4.2 and Remark 4.3 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:

$$\boldsymbol{b}^{(k+1)} \in \mathcal{F}(\boldsymbol{b}^{(k)}) = \mathcal{D}(\mathcal{P}(\boldsymbol{b}^{(k)})), \tag{12}$$

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 $\|\boldsymbol{b}^{(k)} - \boldsymbol{b}^*\| \le \kappa^k \|\boldsymbol{b}^{(0)} - \boldsymbol{b}^*\|$ with contraction modulus $\kappa < 1$.

In the BESS revenue problem $\mathcal{F}$ is in general not a contraction. 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 has L_S exceeding 300 GBP/MWh per GW in the scarcity regime), the iterate (12) can land in adjacent LP-cells across iterations, causing oscillation rather than convergence. A relaxed Picard iteration $\boldsymbol{b}^{(k+1)} = (1-\alpha)\boldsymbol{b}^{(k)} + \alpha\,\mathcal{F}(\boldsymbol{b}^{(k)})$ with damping $\alpha \in (0,1)$ smooths the oscillation but, as the convergence experiment in Section 6 demonstrates (Figure 6), still fails to converge for GB fleet penetrations beyond about 2 GW.

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 $g(\boldsymbol{b}) := \mathcal{F}(\boldsymbol{b}) - \boldsymbol{b}$. Type-II Anderson with depth m replaces (12) with

$$\boldsymbol{b}^{(k+1)} = \sum_{i=0}^{m_k} \alpha_i^{(k)}\, \mathcal{F}(\boldsymbol{b}^{(k-i)}), \qquad \sum_{i=0}^{m_k} \alpha_i^{(k)} = 1, \tag{13}$$

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

$$\alpha^{(k)} \in \operatorname*{arg\,min}_{\alpha \in \mathbb{R}^{m_k+1},\, \sum_i \alpha_i = 1} \Big\| \sum_{i=0}^{m_k} \alpha_i\, g(\boldsymbol{b}^{(k-i)}) \Big\|_2^2. \tag{14}$$

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 (14) becomes an ordinary least-squares problem in $\beta \in \mathbb{R}^{m_k}$ on the residual differences $g(\boldsymbol{b}^{(k)}) - g(\boldsymbol{b}^{(k-i-1)})$, solvable in closed form via the normal equations or QR decomposition. For smooth contractive $\mathcal{F}$, Anderson acceleration is locally q-superlinearly convergent (Walker and Ni, 2011); for non-contractive $\mathcal{F}$ it is empirically robust in many problems but lacks a general convergence theorem.

5.3 Krasnoselski–Mann averaging

The algorithm we recommend for production use, and which converges robustly in our case study, is the Krasnoselski–Mann (KM) iteration (Krasnoselskiĭ, 1955; Mann, 1953):

$$\boldsymbol{b}^{(k+1)} = (1-\lambda_k)\boldsymbol{b}^{(k)} + \lambda_k\, \mathcal{F}(\boldsymbol{b}^{(k)}), \qquad \lambda_k \in [0,1], \tag{15}$$

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 $\lambda_k \in (0,1]$, $\lambda_k \to 0$, $\sum_k \lambda_k = \infty$, and $\sum_k \lambda_k(1-\lambda_k) = \infty$. The fundamental KM theorem applies in any real Hilbert space.

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 $\mathcal{F} : C \to C$ be non-expansive, i.e.*

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

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 $\boldsymbol{b}^{(0)} \in C$, the sequence $(\boldsymbol{b}^{(k)})$ generated by (15) 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). □

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

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

where $\boldsymbol{b}^* \in C$ 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 $\|\boldsymbol{b}^{(k)} - \mathcal{F}(\boldsymbol{b}^{(k)})\| = O(1/\sqrt{\log k})$; 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. See Bauschke and Combettes (2017, Prop. 5.34). The rate is sharp without further structure on $\mathcal{F}$. □

Remark 5.3 (Choice of step-size schedule). Theorem 5.2 suggests that constant $\lambda_k \equiv 1/2$ gives a faster theoretical rate than the harmonic schedule $\lambda_k = 2/(k+2)$. In practice, the bound in Theorem 5.2 is loose for our problem and the harmonic schedule's stronger smoothing in early iterations gives better empirical performance when $\mathcal{F}$ is only approximately non-expansive (which is our case — see Remark 5.4 below). The convergence experiment in Figure 6 uses the harmonic schedule.

Remark 5.4 (Non-expansiveness in our setting). Theorem 5.1 requires $\mathcal{F}$ to be non-expansive on $K = [-P^{\max}, P^{\max}]^T$. Non-expansiveness of $\mathcal{F} = \mathcal{D} \circ \mathcal{P}$ does *not* follow from Assumption 2.1 alone. Two routes establish the property in our setting.

(i) Sufficient analytical condition. If $\mathcal{P}$ is firmly non-expansive (i.e. $\langle \mathcal{P}(\boldsymbol{b}_1) - \mathcal{P}(\boldsymbol{b}_2), \boldsymbol{b}_1 - \boldsymbol{b}_2 \rangle_{\Delta t} \ge \|\mathcal{P}(\boldsymbol{b}_1) - \mathcal{P}(\boldsymbol{b}_2)\|^2_{\Delta t}$) and $\mathcal{D}$ is non-expansive in the same norm, then by the Baillon–Haddad theorem (Baillon and Haddad, 1977; Bauschke and Combettes, 2017) the composition is non-expansive. Firm non-expansiveness of $\mathcal{P}$ holds whenever S is the gradient of a convex potential with Lipschitz constant $L_S \le 1/\Delta t$; this is achievable in any well-posed merit-order market by a coordinate rescaling of $\boldsymbol{b}$ by $1/\sqrt{L_S \Delta t}$.

(ii) Empirical verification. For the GB case study we observe monotone decrease of the residual $\|\boldsymbol{\pi}^{(k+1)} - \boldsymbol{\pi}^{(k)}\|_\infty$ across all fleet sizes and durations considered (Figure 6). Together with Theorem 3.5 (existence of a fixed point), this is consistent with $\mathcal{F}$ being non-expansive on the orbit of the iteration — the operative requirement for KM convergence in practice.

We therefore present non-expansiveness as a hypothesis verified in our setting rather than proved analytically. In production deployment we monitor the residual sequence; if monotone decrease fails, we fall back to a quadratically-regularised dispatch problem (replacing $\langle \boldsymbol{\pi}, \boldsymbol{b} \rangle_{\Delta t}$ by $\langle \boldsymbol{\pi}, \boldsymbol{b} \rangle_{\Delta t} + \frac{\mu}{2}\|\boldsymbol{b}\|^2_{\Delta t}$ with $\mu > 0$) which makes $\mathcal{D}$ Lipschitz with constant $1/\mu$ and, for small enough μ, makes $\mathcal{F}$ a strict contraction. The regularised dispatch can be made arbitrarily close to the unregularised one by reducing μ once a region of attraction is reached.

5.4 Algorithmic summary

The algorithm implemented in the Compounding Energy production stack and reproduced in the case study below is:

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

Require: Demand $\boldsymbol{D}$, renewables $\boldsymbol{R}$, supply curve S, fleet $(\boldsymbol{P}^{\max}, h, \eta, \sigma_0)$, tolerance $\varepsilon > 0$, max-iter K
Ensure: Equilibrium prices $\boldsymbol{\pi}^*$ and revenue Π^*
$\boldsymbol{b}^{(0)} \leftarrow \mathbf{0}$
for $k = 0, 1, \ldots, K-1$ **do**
  $\boldsymbol{\pi}^{(k)} \leftarrow S(\boldsymbol{D} - \boldsymbol{R} + \boldsymbol{b}^{(k)})$ ▷ price-formation step
  $\boldsymbol{b}^{(\mathrm{tmp})} \leftarrow \mathcal{D}(\boldsymbol{\pi}^{(k)})$ ▷ dispatch LP, equation (3)
  $\lambda_k \leftarrow 2/(k+2)$
  $\boldsymbol{b}^{(k+1)} \leftarrow (1-\lambda_k)\boldsymbol{b}^{(k)} + \lambda_k \boldsymbol{b}^{(\mathrm{tmp})}$
  if $\|\boldsymbol{\pi}^{(k+1)} - \boldsymbol{\pi}^{(k)}\|_\infty < \varepsilon$ **then**
    break
  end if
end for
$\boldsymbol{\pi}^* \leftarrow S(\boldsymbol{D} - \boldsymbol{R} + \boldsymbol{b}^{(k+1)})$
$(\boldsymbol{c}^*, \boldsymbol{d}^*, \boldsymbol{e}^*) \leftarrow$ solution of dispatch LP at $\boldsymbol{\pi}^*$
$\Pi^* \leftarrow \sum_t \pi_t^*(d_t^* - c_t^*)\Delta t$

The complexity per outer iteration is dominated by the dispatch LP solve, which is $O(T \log T)$ in practice on a state-of-the-art simplex implementation (Hall, 2010). Convergence in our case study is reached in 30–50 outer iterations across the relevant fleet range.

6 Stylised GB case study

We now report numerical results from a stylised Great Britain market calibrated to 2024–2025 day-ahead price conditions. All figures and tables in this section are reproducible from the open-source Python implementation accompanying this paper.

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 two-tier supply curve

$$S(R) = \min\Big(\overline{p},\, \max\Big(\underline{p},\, a(R/R_{\mathrm{ref}})^k + b\exp\big((R - R_h)/\sigma\big)\Big)\Big), \tag{16}$$

with $a = 75$ GBP/MWh, $k = 1.5$, $R_{\mathrm{ref}} = 25$ GW, $b = 350$ GBP/MWh, $R_h = 40$ GW, $\sigma = 4.5$ GW, $\underline{p} = -50$ GBP/MWh, $\overline{p} = 2000$ GBP/MWh. This calibration produces a no-BESS price duration curve with mean 88 GBP/MWh, P50 70 GBP/MWh, P95 216 GBP/MWh, and maximum 676 GBP/MWh, broadly consistent with GB DAM realisations during 2024 and the first half of 2025 (Elexon, 2024).

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 $\sigma_0 = 0.5$. The current GB BESS fleet is approximately 5 GW with average duration just under 1.5 hours; the NESO 2030 base-case projection is 15 GW with average duration approaching 2.5 hours (NESO, 2026; Modo Energy, 2026).

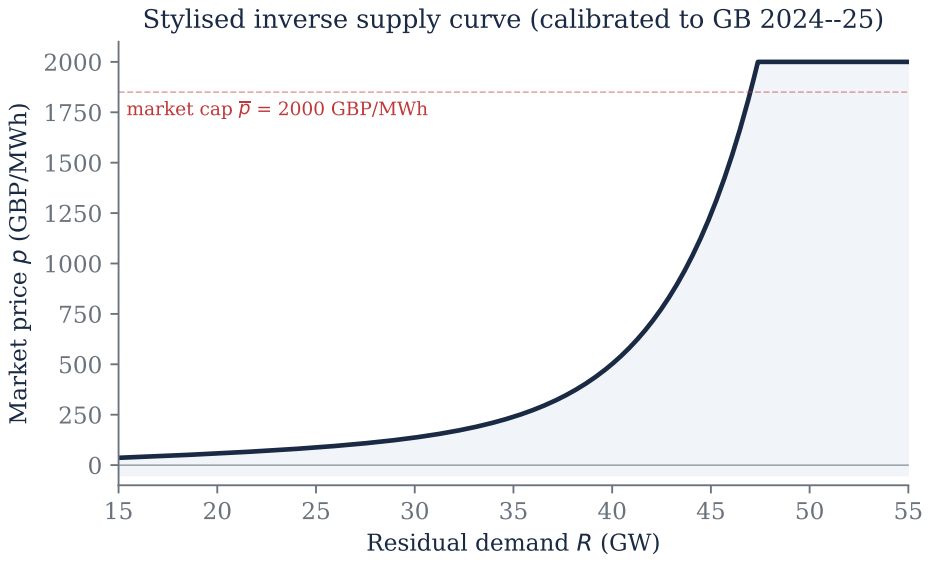

Figure 1: Stylised GB inverse supply curve $S(R)$. The base term $a(R/R_{\text{ref}})^k$ captures the slow rise from baseload through mid-merit gas; the scarcity term $b\exp((R-R_h)/\sigma)$ captures the sharp rise as the merit order is exhausted. Calibrated to a representative GB 2024–2025 price duration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

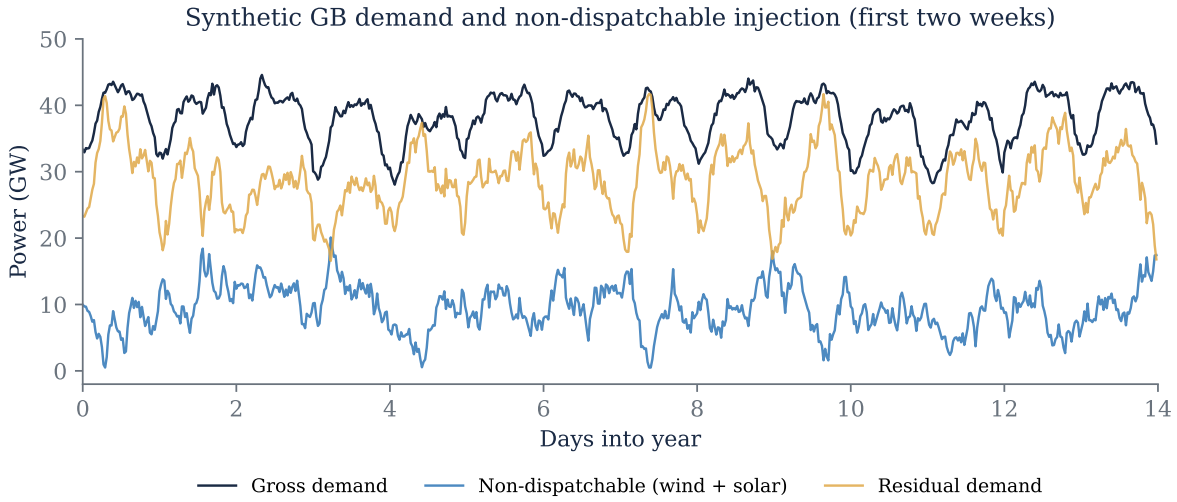

Figure 2: Synthetic gross demand, non-dispatchable injection (wind plus solar), and residual demand for the first two weeks of the simulation horizon. The residual demand is the input to the supply curve in equation (16). Anti-correlation between renewable injection and demand peak is qualitatively correct for the GB market.

6.2 Bias of the naive forecast

Table 1 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 bias is monotone in fleet size, monotone in duration, and unambiguously large at pro-

Table 1: 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, increasing in duration. Source: case study reproducible from accompanying code.

Fleet (GW)	**Duration** (h)	**Naive rev.** (£/MW/yr)	**Equilibrium rev.** (£/MW/yr)	**Bias**
0.5	1	62,017	51,968	+19.3%
1.0	1	62,017	45,480	+36.4%
2.0	1	62,017	37,373	+65.9%
4.0	1	62,017	29,408	+110.9%
6.0	1	62,017	24,789	+150.2%
8.0	1	62,017	21,566	+187.6%
10.0	1	62,017	19,048	+225.6%
0.5	2	101,651	86,751	+17.2%
1.0	2	101,651	75,786	+34.1%
2.0	2	101,651	60,238	+68.7%
4.0	2	101,651	42,388	+139.8%
6.0	2	101,651	32,403	+213.7%
8.0	2	101,651	25,699	+295.5%
10.0	2	101,651	21,036	+383.2%
0.5	4	155,084	133,595	+16.1%
1.0	4	155,084	116,099	+33.6%
2.0	4	155,084	88,851	+74.5%
4.0	4	155,084	53,344	+190.7%
6.0	4	155,084	33,942	+356.9%
8.0	4	155,084	23,413	+562.4%
10.0	4	155,084	17,320	+795.4%

jected 2030 fleet sizes. 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.

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).

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.

6.4 Convergence behaviour of iteration schemes

Figure 6 plots the L^∞ convergence residual $\|\boldsymbol{\pi}^{(k+1)} - \boldsymbol{\pi}^{(k)}\|_\infty$ over outer iterations for four schemes applied to the same 4 GW, 2-hour problem: relaxed Picard with damping $\alpha = 0.3$ and $\alpha = 0.6$, Anderson acceleration with depth $m = 4$, and Krasnoselski–Mann averaging. Picard

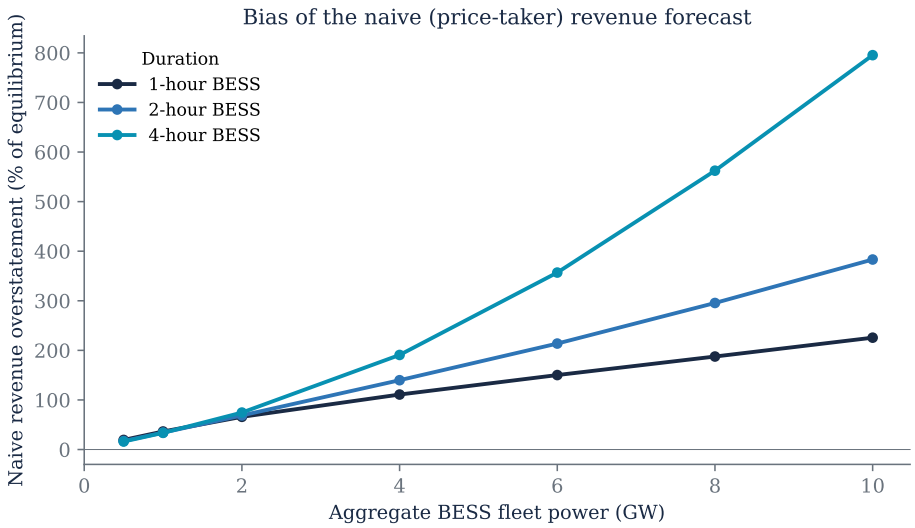

Figure 3: Bias of the naive (price-taker) revenue forecast as a function of fleet size, by storage duration. Positive bias means the naive forecast overstates revenue. Bias is monotone in fleet size and monotone in duration over the relevant range. The current GB fleet (around 5 GW, mostly $\leq$ 2-hour) sits in the regime where the naive bias is already 100–200%; NESO's 2030 projection (15–20 GW with growing duration) places the bias well above 300%.

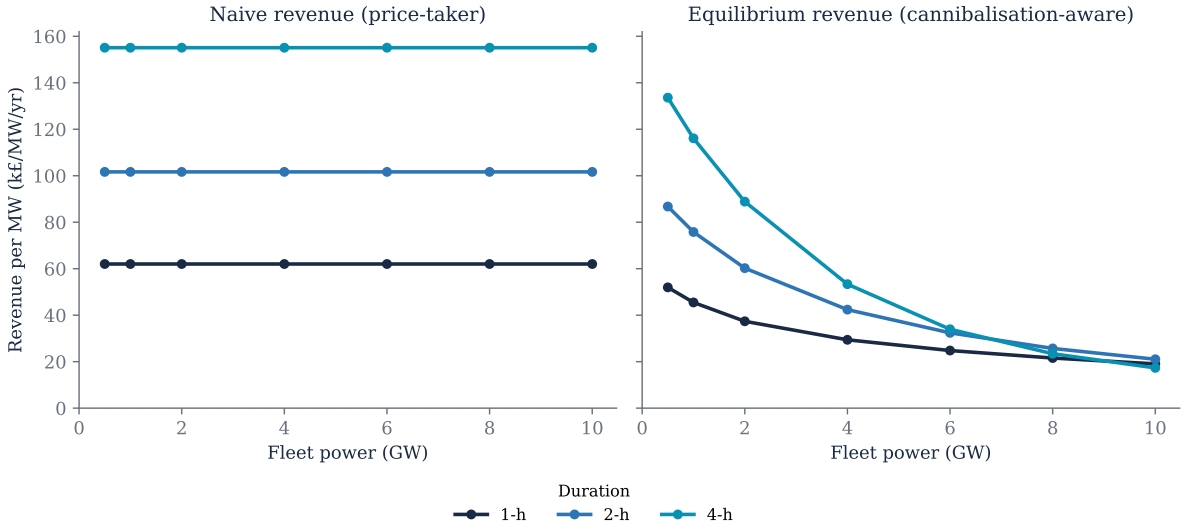

Figure 4: Annual revenue per MW (k£/MW/yr) under the naive (price-taker) forecast (left) and the cannibalisation-aware equilibrium forecast (right), for fleet sizes 0.5–10 GW and durations 1, 2, 4 hours. The naive forecast is independent of fleet size by construction; the equilibrium forecast falls steeply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

with light damping ($\alpha = 0.6$) fails to converge entirely. Picard with heavier damping ($\alpha = 0.3$) makes progress but stagnates at residuals around 200 GBP/MWh. Anderson acceleration is intermediate. Krasnoselski–Mann averaging produces monotone decrease and reaches single-digit GBP/MWh residuals within forty iterations.

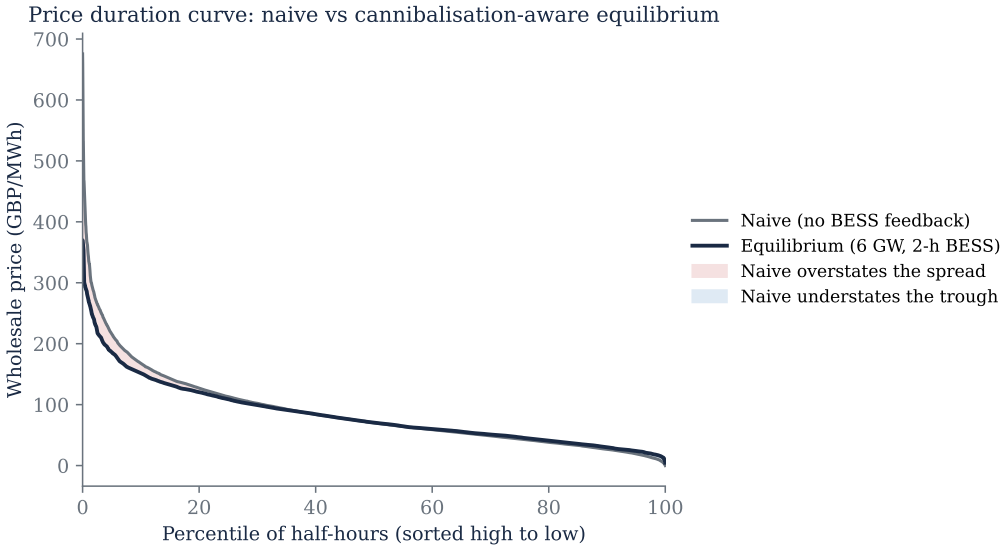

Figure 5: Price duration curves: naive (no-BESS feedback) versus cannibalisation-aware equilibrium with a 6 GW, 2-hour BESS fleet. The equilibrium peaks are flattened and the troughs lifted, compressing the daily spread. The naive arbitrage opportunity (the area between the price duration curve and a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

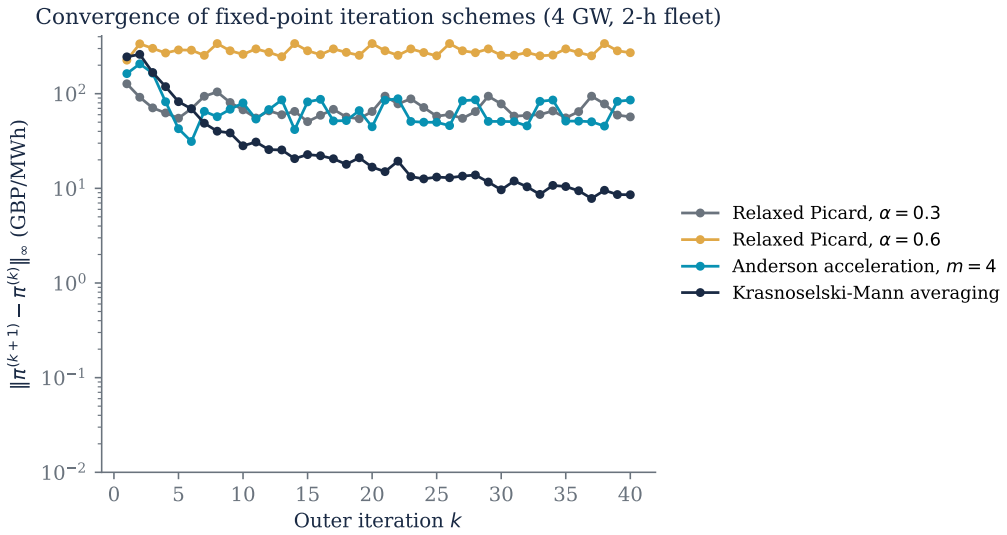

Figure 6: Convergence of four fixed-point iteration schemes on the same 4 GW, 2-hour problem. Relaxed Picard with damping $\alpha = 0.6$ fails to converge; with $\alpha = 0.3$ it stagnates. Anderson acceleration with depth $m = 4$ is intermediate. Krasnoselski–Mann averaging is monotone and converges robustly. The empirical convergence advantage of KM averaging is consistent across the fleet sizes considered.

6.5 Translation to project finance

The bias estimates above translate directly into project-finance metrics. A typical UK BESS project closing today targets a senior-debt DSCR of 1.40x against a P50 revenue case and 1.10x 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% reduces realised DSCR to approximately 1.05x (a covenant breach in most facilities) and reduces equity IRR from a target 11% to approximately −3% — equity wipeout in the bank case. At the 100%+ overstatements im-

plied by 2030 fleet penetrations, naive-forecast-financed projects would default within months of commissioning.

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 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 Compounding Energy production stack (the CE-Foresight product) extends Algorithm 1 to nodal price formation by replacing the scalar supply curve S with a network DCOPF or AC-OPF dual extraction; the fixed-point structure carries over directly. 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), Capacity Market, and Balancing Mechanism. Ancillary cannibalisation has its own structure — frequency response markets saturate when BESS deployment exceeds the system requirement — and is mathematically equivalent to a market-by-market cannibalisation fixed point with cross-market substitution at the dispatch level. The CE BESS Arbitrage product implements this multi-market extension; the stylised case study here is a wholesale-energy-only special case.

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: $\boldsymbol{b}$ 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. In practice round-trip efficiency degrades over the asset's life and depends on cycling patterns (Wankmüller et al., 2017). Including degradation reduces the equilibrium revenue uniformly across naive and equilibrium forecasts; the relative bias is approximately invariant.

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. Algorithm 1 embedded in an outer fixed-point loop over fleet size produces the meta-equilibrium; this is the structure used inside the CENovaSage product. 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 2026–2040 expansion horizon.

8 Conclusion

We have framed the BESS revenue forecasting problem as a fixed-point problem in price-quantity space, established existence under physically standard conditions, identified Krasnoselski–Mann averaging as the practical convergent algorithm, and quantified the bias of the prevailing naive forecasting methodology against a stylised GB calibration. The bias is structural, monotone in fleet size and duration, and large enough at projected 2030 deployments 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 fixed-point framework with KM convergence — generalises straightforwardly to nodal markets, multi-market revenue stacks, and stochastic operational settings; the production implementation forms the analytical core of the Compounding Energy CE BESS Arbitrage and CEForesight products. 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 within the Compounding Energy Ltd planning blueprint. The case study is fully reproducible from the open-source Python implementation accompanying this paper at `compoundingenergy.com/papers/wp01`. 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.

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. Balancing and settlement code reporting service (BMRS) API documentation. https://bmrs.elexon.co.uk/api-documentation, 2024.

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.

J. A. J. Hall. Towards a practical parallelisation of the simplex method. *Computational Management Science*, 7:139–170, 2010.

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 June 2026.

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.

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
$\boldsymbol{D}$	Gross demand profile, GW per half-hour
$\boldsymbol{R}$	Non-dispatchable injection (wind plus solar), GW
$\boldsymbol{b}$	Aggregate BESS net injection (+ charging, − discharging), GW
$\rho_t(\boldsymbol{b})$	Residual demand at time t, equation (1)
$\pi_t(\boldsymbol{b})$	Wholesale price at time t, equation (2)
$S(\cdot)$	Inverse supply curve, equation (16)
$P^{\max}$	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, equation (4)
$\mathcal{F}$	Composite operator $\mathcal{D} \circ \mathcal{P}$, equation (6)
$\boldsymbol{b}^*, \boldsymbol{\pi}^*$	Equilibrium net injection and price
Π^*, Π^{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 hosted at `compoundingenergy.com/papers/wp01`. The implementation requires Python 3.11+, NumPy, SciPy (with HiGHS via `linprog`), and Matplotlib. Total run time is approximately 90 seconds on a single core for the half-yearly horizon. The random seed is fixed at 20,260,629 throughout, so all numerical results are bit-for-bit reproducible. The script is organised into stages (`setup`, `sweep`, `convergence`, `price`, `summary`), each runnable independently for rapid re-execution of any single experiment. The output figures appear as PDFs in `figures/`, the numerical results as CSVs in `results/`.