Supplementary Information

Continuation-specific reliability assessment of receding-horizon wind–battery–reserve dispatch

Xin He[1]†, Yingshuo Qi[1]†, Xinrong Xie[1]† and Yongpeng Tong[1]*

[1] Department of Mathematics, School of Science, Shihezi University, Shihezi 832003, Xinjiang, China (Xin He: Mathematics and Applied Mathematics; Yingshuo Qi and Xinrong Xie: Information and Computing Science)

† These authors contributed equally to this work.

* Correspondence: tongyp@shzu.edu.cn

This file contains the prespecified and post-freeze robustness details supporting the main Article. Headline numerical claims remain the frozen 20-min baseline. Sensitivity analyses were not used to tune the model, select a controller, or replace the primary results. Percentile ranges are exploratory; Tables S14–S17 report origin-selection and boundary-weighting limitations.

Supplementary Note: Frozen evidence hierarchy and denominators

The UK stream is historically exposed development evidence. The Australian COOPGWF1 cohort is a frozen external temporal stress test. The UK model, scaler, residual bank, physical resource settings and controller definitions were transferred without outcome-based refitting. Complete-horizon paired statistics exclude terminal origins whose future label would leave the same contiguous episode; their executed rows remain in the operational ledger.

Supplementary Table S1. Frozen evidence hierarchy and denominators.

Cohort	Executed intervals	Complete hours	Complete paired origins	Role
UK development	9,028	1,503	9,010	Historically exposed development
Australian check	2,592	432	2,586	Frozen external stress test

Supplementary Note: Paired continuation labels and raw discordance

For each origin, the pre-action resource state, exact first command, future wind path and delivery obligation are fixed. Only the continuation after the common first command changes. The table below reports raw DMPC discordance counts; “deployed-only” means the rolling continuation fails while the named reference does not, and “reference-only” is the reverse.

Supplementary Table S2. Raw discordance counts for paired DMPC continuations.

Cohort	Reference	N	Deployed-only	Reference-only	Mismatch (%)	Gap (pp)
UK development	Frozen-origin FULL	9010	462	1025	16.50	-6.25
UK development	Origin-fixed DMPC plan	9010	552	49	6.67	+5.58
Australian check	Frozen-origin FULL	2586	84	403	18.83	-12.34
Australian check	Origin-fixed DMPC plan	2586	114	6	4.64	+4.18

Supplementary Note: Target-matched probability scores

Scenario probabilities are scored only against realized labels generated by the same continuation. Mean predicted probability and observed event rate are descriptive summaries and are not, by themselves, tests of perfect calibration.

Supplementary Table S3. Target-matched scenario-probability scores.

Cohort	Target continuation	N	Mean score (%)	Observed event (%)	Brier	Mean score - event (pp)
UK development	Frozen-origin FULL	9010	77.03	75.85	0.1740	+1.18
UK development	Origin-fixed DMPC plan	9010	66.52	64.02	0.1866	+2.51
Australian check	Frozen-origin FULL	2586	78.08	69.61	0.2052	+8.47
Australian check	Origin-fixed DMPC plan	2586	65.67	53.09	0.1844	+12.57

Supplementary Note: Dependence-aware block-length sensitivity

The main Article uses seven-day moving blocks with 2,000 paired resamples within contiguous episodes. One- and fourteen-day blocks are sensitivity checks. The 14-day UK analysis contains fixed contributions. All ranges are exploratory; see Supplementary Tables S14–S17 for thinning and boundary-weighting qualifications.

Supplementary Table S4. Moving-block-length sensitivity for paired event gaps.

Cohort	Reference	Block (days)	Estimate (pp)	95% percentile range (pp)	Fixed fragment (%)
Australian check	Frozen-origin FULL	1	-12.34	[-20.30, -5.72]	0.00
Australian check	Frozen-origin FULL	7	-12.34	[-25.48, -8.00]	0.00
Australian check	Frozen-origin FULL	14	-12.34	[-17.40, -10.98]	0.00
Australian check	Origin-fixed DMPC plan	1	+4.18	[+2.47, +5.96]	0.00
Australian check	Origin-fixed DMPC plan	7	+4.18	[+2.55, +5.03]	0.00
Australian check	Origin-fixed DMPC plan	14	+4.18	[+3.09, +5.38]	0.00
UK development	Frozen-origin FULL	1	-6.25	[-11.18, -2.43]	0.00
UK development	Frozen-origin FULL	7	-6.25	[-12.82, -1.80]	0.00
UK development	Frozen-origin FULL	14	-6.25	[-11.00, -4.57]	16.84
UK development	Origin-fixed DMPC plan	1	+5.58	[+4.34, +6.89]	0.00
UK development	Origin-fixed DMPC plan	7	+5.58	[+4.15, +6.73]	0.00
UK development	Origin-fixed DMPC plan	14	+5.58	[+4.42, +6.16]	16.84

Supplementary Note: Observation-availability sensitivity

The primary availability assumption is 20 min. The entire causal pipeline was rebuilt after freeze for 10, 30 and 40 min. The table reports DMPC only because it carries both continuation targets. Eligibility changes slightly with the declared availability gate, so the sweep is a robustness analysis rather than a monotonic latency experiment.

Supplementary Table S5. Observation-availability sensitivity for DMPC.

Cohort	Delay (min)	Reference	N	Gap (pp)	7-day 95% range (pp)
Australian check	10	Frozen-origin FULL	2586	-12.72	[-26.34, -7.89]
Australian check	10	Origin-fixed plan	2586	+3.94	[+2.20, +4.45]
Australian check	20	Frozen-origin FULL	2586	-12.34	[-25.48, -8.00]
Australian check	20	Origin-fixed plan	2586	+4.18	[+2.55, +5.03]
Australian check	30	Frozen-origin FULL	2586	-11.21	[-26.14, -5.80]
Australian check	30	Origin-fixed plan	2586	+5.53	[+2.94, +7.27]
Australian check	40	Frozen-origin FULL	2586	-11.87	[-26.06, -5.95]
Australian check	40	Origin-fixed plan	2586	+4.80	[+2.71, +6.65]
UK development	10	Frozen-origin FULL	9012	-4.88	[-13.38, -0.10]
UK development	10	Origin-fixed plan	9012	+7.49	[+4.42, +8.52]
UK development	20	Frozen-origin FULL	9010	-6.25	[-12.82, -1.80]
UK development	20	Origin-fixed plan	9010	+5.58	[+4.15, +6.73]
UK development	30	Frozen-origin FULL	9008	-7.57	[-14.85, -3.33]
UK development	30	Origin-fixed plan	9008	+4.48	[+3.27, +5.26]
UK development	40	Frozen-origin FULL	9006	-8.46	[-15.12, -3.81]
UK development	40	Origin-fixed plan	9006	+3.80	[+2.84, +5.10]

Supplementary Note: One-factor resource sensitivity

Seven fixed resource configurations perturb battery power, usable-energy bounds, or directional reserve power/energy one factor at a time. The baseline remains the only primary configuration. Values are normalized by the chosen farm-capacity reference.

Supplementary Table S6. One-factor resource sensitivity.

Config	Cohort	Batt. P	E range	Reserve P	Reserve E	FULL gap (pp)	Own-plan gap (pp)
baseline	UK development	0.40	0.11-0.99	0.20	0.0500	-6.25	+5.58
baseline	Australian check	0.40	0.11-0.99	0.20	0.0500	-12.34	+4.18
pb_low	UK development	0.30	0.11-0.99	0.20	0.0500	-6.23	+5.64
pb_low	Australian check	0.30	0.11-0.99	0.20	0.0500	-12.18	+4.29
pb_high	UK development	0.50	0.11-0.99	0.20	0.0500	-6.25	+5.59
pb_high	Australian check	0.50	0.11-0.99	0.20	0.0500	-12.34	+4.14
energy_low	UK development	0.40	0.22-0.88	0.20	0.0500	-6.27	+5.58
energy_low	Australian check	0.40	0.22-0.88	0.20	0.0500	-12.34	+4.18
energy_high	UK development	0.40	0.00-1.10	0.20	0.0500	-6.39	+5.53
energy_high	Australian check	0.40	0.00-1.10	0.20	0.0500	-12.18	+4.25
reserve_low	UK development	0.40	0.11-0.99	0.15	0.0375	-5.99	+5.32
reserve_low	Australian check	0.40	0.11-0.99	0.15	0.0375	-12.53	+3.63
reserve_high	UK development	0.40	0.11-0.99	0.25	0.0625	-6.29	+5.84
reserve_high	Australian check	0.40	0.11-0.99	0.25	0.0625	-12.34	+3.79

Supplementary Note: Leave-one-day-out diagnostic

This post-freeze diagnostic omits each eligible calendar day in turn and recomputes the two DMPC continuation gaps. It is not a resampling confidence interval; it tests whether one calendar day alone is necessary for the headline sign.

Supplementary Table S7. Leave-one-day-out diagnostic.

Cohort	Days	LODO FULL gap range (pp)	LODO own-plan gap range (pp)	Daily FULL <0	Daily FULL >0	Daily own >0	Daily own =0
UK development	66	-6.58 to -5.21	+5.33 to +5.68	33	26	53	11
Australian check	18	-13.80 to -10.97	+3.73 to +4.42	12	6	15	3

Supplementary Note: AEMO payload timing gate

A file-level audit re-parsed the archived AEMO payload-creation metadata. This supports only the non-use of not-yet-created public payloads. It does not identify operator-internal telemetry receipt or the time an external client could retrieve a file.

Supplementary Table S8. AEMO payload-creation timing margins.

Availability gate (min)	N proxies	All created before gate	Min margin (s)	P01 margin (s)	Median margin (s)
10	2592	Yes	262	278	288
20	2592	Yes	862	878	888
30	2592	Yes	1462	1478	1488
40	2592	Yes	2062	2078	2088

Supplementary Note: Operational endpoint details

Energy totals are integrated over each cohort and should not be compared directly between cohorts of different duration. The objective coefficients are engineering proxies, not tariffs.

Supplementary Table S9. Operational endpoints and solver fallback.

Cohort	Controller	Failed hours (%)	Excess shortfall (pu h)	Mean excess P (pu)	Battery (pu h)	Reserve (pu h)	Curtailment (pu h)	Fallback
UK development	F_FULL	78.84	35.20	0.0234	90.84	10.15	40.19	0
UK development	F_BR0	58.02	41.28	0.0274	10.64	25.74	66.60	0
UK development	DMPC	58.68	39.04	0.0259	19.74	28.33	66.57	0
Australian check	F_FULL	60.65	9.70	0.0224	27.93	2.59	9.48	0
Australian check	F_BR0	47.45	14.78	0.0342	2.36	7.03	21.39	0
Australian check	DMPC	47.22	13.76	0.0319	4.01	8.04	21.25	0

Supplementary Note: Audit and claim boundary

Independent checks reconstructed the training/residual/development indices, re-solved ridge coefficients, perturbed not-yet-available inputs, checked resource-prefix constraints and mutually exclusive actions, replayed saved execution steps, rebuilt paired labels, and recomputed target-matched probability scores. Exact replay of an identical fixed command sequence produced zero directional label discrepancies. All full baseline analyses completed with zero solver fallback.

The evidence does not establish field telemetry latency, settlement-metered energy delivery, commercial profit, battery degradation economics, universal superiority of any controller, distribution-free conditional safety, or broad cross-site generalization. The Australian 18-day cohort is an external frozen-cohort stress test rather than a prospective multi-site validation.

Supplementary Note: Supplementary Data contents

Supplementary Data 1 contains the core reproducibility package, compact timing/provenance audit outputs, frozen robustness tables and scripts, the leave-one-day-out/discordance diagnostic, and the verification records added on 6 October 2026. The four auxiliary construction scripts are retained for provenance only (Supplementary Note “Verification and portability update of 6 October 2026”). Raw AEMO daily ZIP archives are excluded from this file. Supplementary Data 2 contains the 30 official AEMO daily DISPATCH_SCADA source archives together with source/download and timing-provenance records. Baseline analysis-only regeneration and controller recomputation require Supplementary Data 1 alone. The new post-review analyses additionally require Supplementary Data 3 (Supplementary_Data_3_Revision_Analyses.zip). Supplementary Data 2 is needed only to repeat raw AEMO archive parsing and provenance reconstruction. For that offline raw-source check, copy Supplementary Data 2's raw/ directory to core_reproducibility/confirmation/raw/ inside Supplementary Data 1 and follow the package instructions. Each archive contains its own SHA-256 manifest.

Supplementary Data 1 file

Supplementary_Data_1_Reproducibility_20261006.zip

Supplementary Data 2 file

Supplementary_Data_2_AEMO_Raw_Sources.zip

Supplementary Note: Post-freeze scoring-definition sensitivity

The primary Article defines a directional shortfall event using a 0.01 pu scoring tolerance over the full active 7-12 step (70-120 min) window. To test whether the reference-dependent sign reversal was an artifact of that event definition, we performed a post-freeze relabeling analysis on the frozen DMPC replays. Controller actions, the original 0.01 pu local-curtailment executor, physical state trajectories, realized wind paths and synthetic delivery obligations were held fixed. Only the event-scoring rule was changed. Accordingly, these checks assess scoring-definition sensitivity; they are not alternative controller designs or physical-service-tolerance experiments.

As a reconstruction check, relabeling at the archived 0.01 pu threshold reproduced every stored shortfall label: 0 mismatches among 9,010 UK paired origins and 0 among 2,586 Australian paired origins. Each condition uses the same paired all-origin estimator and 2,000 seven-day moving-block resamples as the primary analysis.

Threshold sensitivity over the full active window

Supplementary Table S10. Scoring-threshold sensitivity over the full active window.

Cohort	Scoring threshold (pu)	Reference	N	Reference event (%)	Deployed event (%)	Gap (pp)	7-day 95% range (pp)
UK development	0.005	Frozen-origin FULL	9,010	78.81	72.45	-6.36	[-13.02, -1.53]
UK development	0.005	Origin-fixed DMPC plan	9,010	66.91	72.45	+5.54	[3.92, 6.98]
UK development	0.010	Frozen-origin FULL	9,010	75.85	69.60	-6.25	[-12.82, -1.80]
UK development	0.010	Origin-fixed DMPC plan	9,010	64.02	69.60	+5.58	[4.15, 6.73]
UK development	0.020	Frozen-origin FULL	9,010	69.71	63.91	-5.80	[-12.89, -1.21]
UK development	0.020	Origin-fixed DMPC plan	9,010	58.11	63.91	+5.79	[3.81, 7.26]
UK development	0.030	Frozen-origin FULL	9,010	62.97	58.31	-4.66	[-10.70, -0.61]
UK development	0.030	Origin-fixed DMPC plan	9,010	53.10	58.31	+5.22	[2.91, 7.06]
Australian stress test	0.005	Frozen-origin FULL	2,586	71.77	59.67	-12.10	[-24.94, -7.69]
Australian stress test	0.005	Origin-fixed DMPC plan	2,586	55.49	59.67	+4.18	[2.09, 4.68]
Australian stress test	0.010	Frozen-origin FULL	2,586	69.61	57.27	-12.34	[-25.48, -8.00]
Australian stress test	0.010	Origin-fixed DMPC plan	2,586	53.09	57.27	+4.18	[2.55, 5.03]
Australian stress test	0.020	Frozen-origin FULL	2,586	61.48	53.94	-7.54	[-17.52, -3.75]
Australian stress test	0.020	Origin-fixed DMPC plan	2,586	48.96	53.94	+4.99	[3.02, 6.26]
Australian stress test	0.030	Frozen-origin FULL	2,586	52.51	47.76	-4.76	[-11.37, -2.47]
Australian stress test	0.030	Origin-fixed DMPC plan	2,586	44.97	47.76	+2.78	[0.19, 4.25]

Across all four scoring thresholds, the deployed-minus-frozen-origin-FULL gap remained negative and the deployed-minus-own-plan gap remained positive in both cohorts. The corresponding seven-day percentile ranges excluded zero at every tested threshold. Thus the headline sign reversal is not confined to the single 0.01 pu scoring threshold.

Prefix-window sensitivity at the primary 0.01 pu scoring threshold

Supplementary Table S11. Prefix-window sensitivity at the primary 0.01 pu scoring threshold.

Cohort	Scoring window	Reference	N	Reference event (%)	Deployed event (%)	Gap (pp)	7-day 95% range (pp)
UK development	30 min (3 steps)	Frozen-origin FULL	9,010	57.55	49.18	-8.37	[-13.86, -4.68]
UK development	30 min (3 steps)	Origin-fixed DMPC plan	9,010	46.80	49.18	+2.38	[1.69, 2.94]
UK development	60 min (6 steps)	Frozen-origin FULL	9,010	68.36	62.10	-6.26	[-13.44, -1.79]
UK development	60 min (6 steps)	Origin-fixed DMPC plan	9,010	57.48	62.10	+4.62	[3.24, 5.88]
UK development	Full 70-120 min	Frozen-origin FULL	9,010	75.85	69.60	-6.25	[-12.82, -1.80]
UK development	Full 70-120 min	Origin-fixed DMPC plan	9,010	64.02	69.60	+5.58	[4.15, 6.73]
Australian stress test	30 min (3 steps)	Frozen-origin FULL	2,586	55.26	44.43	-10.83	[-20.77, -7.97]
Australian stress test	30 min (3 steps)	Origin-fixed DMPC plan	2,586	43.12	44.43	+1.31	[0.93, 1.43]
Australian stress test	60 min (6 steps)	Frozen-origin FULL	2,586	64.04	52.32	-11.72	[-24.44, -7.65]
Australian stress test	60 min (6 steps)	Origin-fixed DMPC plan	2,586	49.30	52.32	+3.02	[1.78, 3.44]
Australian stress test	Full 70-120 min	Frozen-origin FULL	2,586	69.61	57.27	-12.34	[-25.48, -8.00]
Australian stress test	Full 70-120 min	Origin-fixed DMPC plan	2,586	53.09	57.27	+4.18	[2.55, 5.03]

At the primary 0.01 pu threshold, the same reference-dependent sign reversal was present when scoring only the first 30 min, the first 60 min, or the full active window. In the UK cohort, frozen-origin-FULL gaps ranged from -8.37 to -6.25 pp and own-plan gaps from +2.38 to +5.58 pp. In the Australian stress test, the corresponding ranges were -12.34 to -10.83 pp and +1.31 to +4.18 pp. Every displayed seven-day percentile range excluded zero.

Interpretation boundary. These analyses strengthen the assessment claim that the observed sign reversal is not an artifact of one scoring threshold or one window length. They do not show that the result is invariant to a different commercial obligation, a redesigned executor, alternative battery dynamics, or a different physical service tolerance. Those remain outside the evidence of the present study.

Machine-readable values, the relabeling script and a validation record are included in revised Supplementary Data 1 under robustness/.

Supplementary Note: Post-freeze synthetic-obligation-level sensitivity

The primary Article uses a synthetic hourly delivery obligation constructed from the power observation labelled 80 min before delivery. To test whether the headline reference-dependent sign reversal depended narrowly on the level of that synthetic service obligation, we performed a full post-freeze DMPC rerun at obligation scales 0.90 and 1.10, with the frozen primary analysis retained as scale 1.00.

Each obligation was replaced by $clip(\kappa s, 0, 1)$. Information timing, forecasts, realized wind, resource limits, controller coefficients and continuation definitions were unchanged. Unlike the scoring-definition analysis in Supplementary Note “Post-freeze scoring-definition sensitivity”, the controller and both paired continuations were recomputed from scratch at $\kappa = 0.90$ and 1.10.

Each row below uses the same all-origin paired estimator as the primary Article and 2,000 seven-day ordinary moving-block resamples within contiguous episodes. A negative gap means fewer deployed shortfall events than under the named reference continuation.

UK development (N = 9,010 paired origins)

Supplementary Table S12. Synthetic-obligation-level sensitivity in UK development.

Scale	Reference	Reference event (%)	Deployed event (%)	Gap (pp)	7-day 95% range (pp)
0.90	Frozen-origin FULL	71.41	59.67	-11.74	[-17.96, -8.33]
0.90	Origin-fixed DMPC plan	56.89	59.67	+2.78	[2.03, 3.25]
1.00	Frozen-origin FULL	75.85	69.60	-6.25	[-12.82, -1.80]
1.00	Origin-fixed DMPC plan	64.02	69.60	+5.58	[4.15, 6.73]
1.10	Frozen-origin FULL	82.13	77.77	-4.36	[-10.98, -0.87]
1.10	Origin-fixed DMPC plan	71.93	77.77	+5.84	[4.42, 6.13]

Australian external stress test (N = 2,586 paired origins)

Supplementary Table S13. Synthetic-obligation-level sensitivity in the Australian external stress test.

Scale	Reference	Reference event (%)	Deployed event (%)	Gap (pp)	7-day 95% range (pp)
0.90	Frozen-origin FULL	65.82	52.09	-13.73	[-25.91, -9.74]
0.90	Origin-fixed DMPC plan	49.19	52.09	+2.90	[2.09, 3.60]
1.00	Frozen-origin FULL	69.61	57.27	-12.34	[-25.48, -8.00]
1.00	Origin-fixed DMPC plan	53.09	57.27	+4.18	[2.55, 5.03]
1.10	Frozen-origin FULL	76.84	65.58	-11.25	[-24.87, -5.84]
1.10	Origin-fixed DMPC plan	60.60	65.58	+4.99	[2.09, 5.96]

Result. Across all three obligation levels, the deployed-minus-frozen-origin-FULL gap remained negative and the deployed-minus-origin-fixed-DMPC-plan gap remained positive in both cohorts; every displayed seven-day percentile range excluded zero. Thus the headline sign reversal is not confined to the exact baseline synthetic-obligation level.

Interpretation boundary. These reruns change the level of the already-defined synthetic obligation while preserving its construction and issue-time rule. They do not validate a different commercial commitment product, a different issuance rule, field settlement delivery, or alternative battery dynamics. Machine-readable values, the reproducible driver and a validation record are included in revised Supplementary Data 1 under robustness/.

Supplementary Note: Verification and portability update of 6 October 2026

A fresh verification reproduced the scientific endpoints in all three archived analysis groups and independently replayed every controlled baseline step and paired label in the UK development and Australian cohorts. The check also reconstructed the non-anticipative input pipeline under its declared simulation clock, re-parsed all 30 AEMO daily archives, checked payload-creation timing, and recomputed target-matched score summaries. Independent scenario-target checks covered 54 UK and 18 Australian origins per controller, 32 scenarios and both branches; these sampled checks do not establish universal calibration.

The complete scoring-definition analysis was rerun, with 34,788 baseline label comparisons and zero mismatches. Four nonbaseline DMPC obligation-level reruns (two cohorts at scales 0.90 and 1.10) completed without solver fallback; all 12 obligation-table rows matched the archived CSV at its three-decimal precision. The headline 20-min baseline, all existing supplementary table values and the frozen analysis choices were retained.

The latency and resource drivers were repaired only to replace historical filesystem roots with package-relative paths. No scientific parameter or algorithm changed. The repaired 20-min latency builder matched the archived input arrays to a maximum absolute difference of 3.50×10^{-15}. The resource driver reproduced all saved Australian baseline controller arrays except measured solve time. The complete altered latency/resource grids were not independently rerun on 6 October 2026; their displayed sweep results remain archived evidence with freshly verified baseline equivalence.

Verification records, the original manifest and script hash changes are included in Supplementary Data 1 under verification_20261006/. The four files in auxiliary_scripts/ (build_round7_manuscript.py, build_round8_revision.py, build_sr_timing_audit.py and regenerate_figure3.py) are retained for provenance only because they depend on historical paths or upstream files; they are not turnkey entry points for rebuilding the final manuscript or figures.

The earlier checks support reproducibility of the model-conditional assessment. They do not establish live telemetry receipt, field delivery validity or controller superiority. The obligations remain synthetic. The UK remains historically exposed development evidence, and Australia remains one frozen external 18-day cohort after a prior-exposure review rather than pristine prospective blinding. Moving-block percentile ranges are exploratory and do not establish population coverage; center shifts and zero-crossing origin-thinning results are documented in Supplementary Tables S14–S17.

Supplementary Note: Additional diagnostics and evidence hierarchy

The following analyses were added after the primary results were known. They do not create a new blinded cohort or replace the frozen 20-min baseline. Tables S1–S13 retain the original numerical evidence; Tables S14–S20 add previously omitted sampling designs, finite-sample resampling diagnostics, score-target comparisons and new targeted simulations. Both supporting and limiting findings are reported. All resampling ranges in this supplement are exploratory conditional percentiles, not validated population confidence intervals.

Origin selection, phase composition and non-overlap

All-origin comparisons weight every complete 7–12-step origin equally. Hourly phase-zero selection retains 12-step windows at exact clock hours. Greedy non-overlap selection starts with the first complete origin in each run and excludes subsequent origins before that window ends. These rules change the origin population and window/phase composition as well as the sample size. No controller trajectory is rerun for this diagnostic. All 12 rows were recalculated from frozen labels and matched the archived bootstrap.json outputs, including the seeds for each sampling design.

Supplementary Table S14. Archived origin-selection sensitivity, freshly reproduced.

Cohort	Ref.	Origin selection	N	Gap (pp)	7-day range (pp)
UK	B	All origins	9,010	-6.25	[-12.82, -1.80]
UK	B	Hourly phase 0	1,500	-5.20	[-11.53, 0.07]
UK	B	Greedy non-overlap	752	-3.72	[-10.11, 2.39]
UK	P	All origins	9,010	+5.58	[4.15, 6.73]
UK	P	Hourly phase 0	1,500	+5.27	[4.40, 6.67]
UK	P	Greedy non-overlap	752	+5.98	[4.12, 7.85]
Australia	B	All origins	2,586	-12.34	[-25.48, -8.00]
Australia	B	Hourly phase 0	431	-13.23	[-26.68, -9.28]
Australia	B	Greedy non-overlap	216	-12.50	[-26.39, -8.80]
Australia	P	All origins	2,586	+4.18	[2.55, 5.03]
Australia	P	Hourly phase 0	431	+5.57	[3.25, 6.73]
Australia	P	Greedy non-overlap	216	+5.09	[2.31, 6.48]

B denotes frozen-origin FULL; P denotes the origin-fixed DMPC plan. Both UK B thinned ranges cross zero despite negative point estimates. Other listed ranges retain their original signs, but this does not remove the change of estimand or establish population coverage. Supplementary Fig. S2 shows all rows.

Supplementary Note: Exact finite-sample resampling centers

The original non-circular algorithm samples full blocks uniformly by allowed starts, preserves each run's sample count and truncates the final drawn block. Uniform block starts do not imply uniform expected weights on observations near run boundaries. The table reports the exact expectation over starts, rather than only a Monte Carlo estimate. No original point estimate or range has been recentered or replaced.

Supplementary Table S15. Ordinary moving-block center diagnostics.

Cohort	Ref.	Days	Observed (pp)	Exact mean (pp)	Shift (pp)	Original range (pp)
UK	B	1	-6.25	-6.49	-0.24	[-11.18, -2.43]
UK	B	7	-6.25	-6.91	-0.66	[-12.82, -1.80]
UK	B	14	-6.25	-7.73	-1.48	[-11.00, -4.57]
UK	P	1	+5.58	+5.58	-0.00	[4.34, 6.89]
UK	P	7	+5.58	+5.42	-0.16	[4.15, 6.73]
UK	P	14	+5.58	+5.30	-0.28	[4.42, 6.16]
Australia	B	1	-12.34	-13.00	-0.67	[-20.30, -5.72]
Australia	B	7	-12.34	-16.89	-4.55	[-25.48, -8.00]
Australia	B	14	-12.34	-14.38	-2.04	[-17.40, -10.98]
Australia	P	1	+4.18	+4.14	-0.04	[2.47, 5.96]
Australia	P	7	+4.18	+3.75	-0.42	[2.55, 5.03]
Australia	P	14	+4.18	+4.27	+0.09	[3.09, 5.38]

For Australia/B at seven days, the observed difference is −12.33565 pp while the exact resampling mean is −16.88744 pp (shift −4.55179 pp). A longer block is not automatically more conservative. At fourteen days, 1,517/9,010 UK origins (16.84%) belong to a fixed short run. Increasing repetitions would reduce simulation noise but would not change these finite-sample weights or supply more independent data.

Exact expectation for one run

For a complete-origin run of length n with differences $a_i = Y_i(D) - Y_i(c)$, take block length $b < n$, $q = \text{floor}(n/b)$, remainder $r = n - qb$, and $A(s,\ell)$ is the sum of a_i for $i = s,\ldots,s+\ell-1$. Starts s range from 0 to $n - b$. The expected unnormalized resampled sum S_{boot} is:

$$E_{boot}(S_{boot}) = \frac{q}{n-b+1} \sum_{s=0,\ldots,n-b} A(s,b) + \frac{1}{n-b+1} \sum_{s=0,\ldots,n-b} A(s,r) \tag{S1}$$

Define $A(s,0) = 0$. For $n \le b$ the run sum is fixed. Summing expectations over runs and multiplying by 100/N gives the table's exact mean. This identity diagnoses the implemented finite-sample resampling distribution; it is not a theorem about coverage of a population confidence interval.

Supplementary Note: Circular boundary sensitivity and temporal dependence

For the circular sensitivity, starts are uniform among all n positions of each run; blocks wrap within that same run and the final draw is truncated to preserve n. Short runs n ≤ b remain fixed as in the primary implementation. Every original observation has equal expected weight, so the exact center equals the observed mean. Wrapping nevertheless creates artificial endpoint adjacency. The sensitivity is not a coverage correction and does not represent physical reconnection of disconnected episodes.

Supplementary Table S16. Within-run circular sensitivity (2,000 replicates).

Cohort	Ref.	Days	Exact mean (pp)	Circular range (pp)	Fixed (%)
UK	B	1	-6.25	[-10.75, -2.13]	0.00
UK	B	7	-6.25	[-12.51, -0.42]	0.00
UK	B	14	-6.25	[-11.82, -0.81]	16.84
UK	P	1	+5.58	[4.28, 6.91]	0.00
UK	P	7	+5.58	[4.18, 7.09]	0.00
UK	P	14	+5.58	[4.48, 6.73]	16.84
Australia	B	1	-12.34	[-19.33, -5.57]	0.00
Australia	B	7	-12.34	[-23.43, -3.98]	0.00
Australia	B	14	-12.34	[-18.60, -6.15]	0.00
Australia	P	1	+4.18	[2.55, 5.92]	0.00
Australia	P	7	+4.18	[2.78, 5.57]	0.00
Australia	P	14	+4.18	[3.02, 5.38]	0.00

Seeds are 20261703 + block-days. The UK B seven-day circular range is [−12.51, −0.42] pp; the Australian B range is [−23.43, −3.98] pp. These values quantify sensitivity to boundary handling, not evidence that circular sampling should replace the frozen ordinary analysis. The data do not establish a valid seven-day independence or stationarity scale.

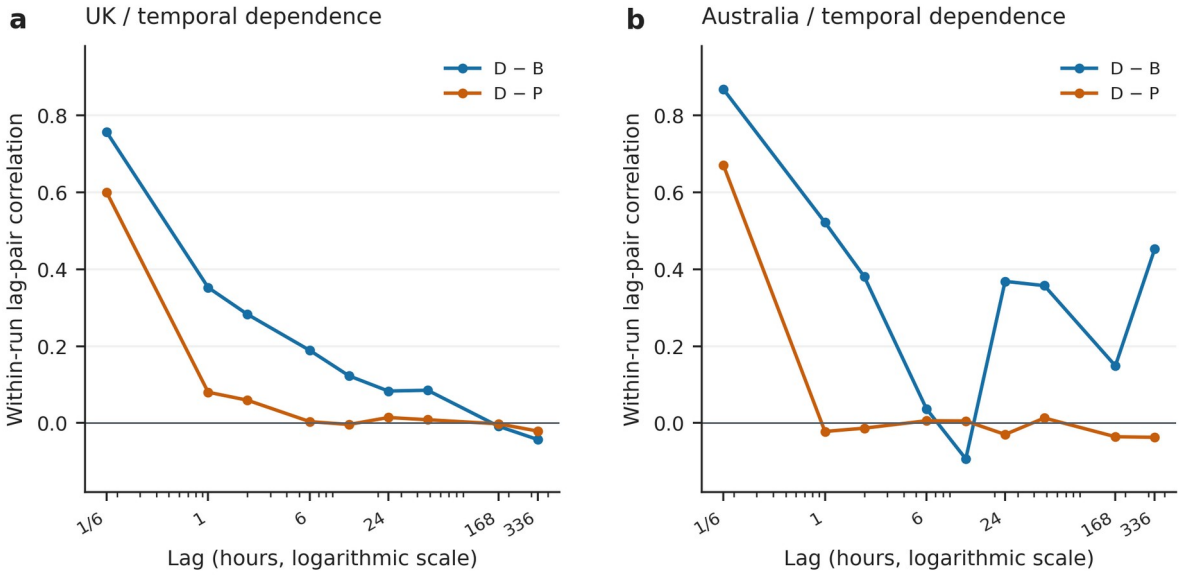
Supplementary Note: Lag-pair correlations

For each lag, the following Pearson correlation uses only pairs whose two origins lie inside the same contiguous complete-origin run. The differences are not detrended. Counts vary with lag; these descriptive correlations are not tests of stationarity and were not used to select a favorable block length. Daily paired summaries and all underlying values are provided in Supplementary Data 3.

Supplementary Table S17. Within-run correlation of paired event differences.

Lag (h)	UK pairs	UK D−B	UK D−P	AU pairs	AU D−B	AU D−P
1/6	9007	0.757	0.600	2585	0.868	0.671
1	8992	0.353	0.080	2580	0.522	-0.022
2	8974	0.283	0.059	2574	0.381	-0.013
6	8902	0.189	0.004	2550	0.037	0.006
12	8794	0.123	-0.004	2514	-0.093	0.005
24	8578	0.083	0.014	2442	0.369	-0.030
48	8146	0.085	0.009	2298	0.357	0.013
168	5986	-0.008	-0.002	1578	0.149	-0.036
336	3461	-0.043	-0.021	570	0.452	-0.037

The Australian D−B difference remains temporally structured: correlations are 0.369 at 24 h, 0.357 at 48 h and 0.149 at seven days. The fourteen-day estimate uses only 570 pairs and is not evidence of a stable long-lag population correlation. The own-plan difference has a different dependence pattern. A single block duration does not have the same empirical rationale for both targets.



Supplementary Figure S3. Lag-pair correlations for the UK (a) and Australian (b) cohorts. Only within-run pairs are included. Horizontal axes are logarithmic; no confidence bands or inferred decorrelation length are asserted.

Supplementary Note: Target-matched and cross-target Brier scores

For a fixed probability vector p_t,c, target-matched scoring uses Y_t(c), whereas the cross-target diagnostic uses Y_t(D). Probabilities, origins and all labels are unchanged from the primary replay. The diagnostic changes only which existing event is treated as the scoring outcome. This is not training or recalibration, and the cross-target value is not substituted for the matched score.

Supplementary Table S18. Identical probabilities scored against different continuations.

Cohort	Probability target	N	Matched Brier	Cross-target D	Cross − matched
UK	B	9,010	0.1740	0.2075	+0.03349
UK	P	9,010	0.1866	0.1669	-0.01964
Australia	B	2,586	0.2052	0.2652	+0.06001
Australia	P	2,586	0.1844	0.1729	-0.01153

Binary score identity

For binary y,d, $(p-d)^2 - (p-y)^2 = (d-y)(1-2p)$. Averaging at the paired origins gives main-text equation (4). Hence a net event-rate difference is insufficient to determine the Brier change: probabilities at the discordant origins supply the weighting. The four numerical residuals between the two sides have maximum absolute value 1.39×10^{-17}. This is an elementary diagnostic identity, not a new risk guarantee.

Own-plan probabilities yield lower cross-target scores in both cohorts even though deployment is not their designated event. FULL-target probabilities yield higher cross-target scores. Neither comparison can establish improved or worsened calibration on an unchanged target. Both targets' failure frequencies and the two directional discordance counts remain necessary to interpret the change.

Supplementary Note: Fixed-original-endpoint replanning

Continuation S shares the archived DMPC pre-action resource state, exact first command, realized wind and original obligation vector. It then repeatedly solves the unchanged DMPC using the currently available point forecast restricted to the remaining original window, while retaining the original endpoint and terminal target. Resource states are updated from S's own commands. All 9,010 UK and 2,586 Australian complete paired origins were processed. S is an origin-specific diagnostic, not one globally implemented replacement policy.

Supplementary Table S19. Sequential continuation contrasts.

Cohort	Contrast	Reference (%)	Target (%)	Gap (pp)	Target only	Reference only
UK	S-P	64.02	67.57	+3.55	388	68
UK	D-S	67.57	69.60	+2.03	183	0
UK	D-P	64.02	69.60	+5.58	552	49
UK	S-B	75.85	67.57	-8.28	343	1089
Australia	S-P	53.09	56.11	+3.02	84	6
Australia	D-S	56.11	57.27	+1.16	30	0
Australia	D-P	53.09	57.27	+4.18	114	6
Australia	S-B	69.61	56.11	-13.50	63	412

The rows S−P and D−S sum algebraically to D−P for every origin and in the average. They are not independent causal components and their interpretation depends on the sequence of interventions. Updating forecasts/reoptimizing can change S relative to P even before endpoint extension is allowed; D relative to S additionally permits the endpoint and associated future obligation set to move. A claim about longer-run controller superiority does not follow from scoring only the original window.

There were 76,580 additional UK LP solves and 21,981 Australian solves, with zero failures or fallbacks. Independently written scalar physics reproduced all 85,590 and 24,567 S window steps, respectively. Maximum absolute discrepancy was 2.22×10^{-16}. Commands before the first clock-hour rollover agreed exactly with D in 31,530 UK and 9,051 Australian step comparisons. At least one D/S command difference occurred in 4,259 UK and 1,224 Australian origins, never before rollover. These are code-level checks under the model, not human review or field validation.

Main Figure 4 illustrates the chronologically earliest origin with YD = 1 and YS = 0 in each cohort, a deterministic outcome-defined rule. The UK origin is 2021-10-14 20:00 UTC (input row 1032); the Australian origin is 2026-09-11 05:00 UTC (row 90). Both have H = 12 and P = S = B = 0, D = 1. At minute 100, D charges 0.09129 pu in the UK and 0.06685 pu in Australia, while S does not charge; their delivery errors are −0.01375 and −0.02837 pu. These cases illustrate one resource-reallocation path and were not selected to estimate prevalence. Full CSV trajectories and constraint-checked ledgers are archived.

Supplementary Note: Alternative synthetic-obligation construction

The original obligation is a clipped single power observation at h−80 min, issued at h−60 min. The new sensitivity uses the mean of six ten-minute power labels from h−130 to h−80 min, then clips to [0,1]. At the unchanged 20-min availability gate, each source is available no later than issue time. Each source index must lie in the same source episode; finite values, timestamp alignment and future obligation suffix consistency are explicitly checked. No date or origin is removed on the basis of an outcome.

Supplementary Table S20. Full DMPC reruns with six-label trailing-mean obligations.

Cohort	Ref.	N	Reference (%)	Deployed (%)	Gap (pp)	Exploratory range (pp)
UK	B	9,010	76.12	66.85	-9.27	[-15.79, -5.16]
UK	P	9,010	62.54	66.85	+4.31	[3.32, 4.93]
Australia	B	2,586	69.41	57.70	-11.72	[-24.63, -7.04]
Australia	P	2,586	53.91	57.70	+3.79	[2.36, 5.34]

Forecast model, scenario bank, controller formulation, resource limits, original cohort dates and issue-time rules remain fixed. The entire chronological controller trajectory and paired P/B/D labels are recomputed; the new results are not relabelings of old actions. The original single-observation baseline remains primary. Both runs had zero fallback. Independent scalar replay checked 9,028 UK and 2,592 Australian executed steps and 27,030 plus 7,758 paired window labels. Maximum executed-step discrepancy was 2.22×10^{-16}.

The result reduces dependence on one synthetic persistence construction, but both rules are synthetic. It does not verify a market product, historical operator commitment, field delivery measurement or a different issuance schedule. Seven-day ordinary ranges use the original procedure and remain subject to the center and dependence limitations in Tables S15–S17.

Supplementary Note: Exploratory 2022 Kelmarsh temporal extension

This evaluation used the first 14 UTC calendar days of 2022 from a continuous quality-screened Kelmarsh aggregate, with the frozen 2021 predictor, phase-specific residual bank, 32 scenarios, physical resource settings and controller definitions. The evaluation consists of 1,991 chronological ten-minute origins, 1,985 complete paired origins and 331 complete hours. The source underwent prior quality assurance, so this is not an annual or prospectively blinded validation. The original 20-min observation gate and synthetic commitment construction are unchanged. The model adaptation altered only date/path validation, not the physical solver; the state was initialised once and propagated over the uninterrupted test episode. Supplementary Data 4 records source checksums, file hashes, scripts, logs and complete origin-wise outcomes.

Supplementary Table S21. Controller-specific continuation comparisons in the 14-day UK 2022 extension.

Controller	Reference	Reference event (%)	Deployment event (%)	Gap (pp)	D only	Ref. only	Discordance (%)
F_FULL	FULL	74.91	91.89	+16.98	359	22	19.19
F_BR0	FULL	78.49	77.03	−1.46	106	135	12.14
DMPC	FULL	76.88	76.52	−0.35	124	131	12.85
DMPC	Own plan	69.72	76.52	+6.80	139	4	7.20

All contrasts are deployment minus the specified reference. The risk tolerance is 0.01 pu. The complete-hour failure rates and resource endpoints appear in main-text Table 4; the 2022 time interval is shorter than either of the original evaluation periods.

Supplementary Table S22. The same 2022 DMPC scenario-probability vectors evaluated against their original targets or the deployment target.

Probability vector	N	Mean probability (%)	Original event (%)	Matched Brier	Brier vs D
Own fixed plan	1,985	70.28	69.72	0.1701	0.1491
Frozen FULL	1,985	78.71	76.88	0.1658	0.1807

The two Brier values in each row use an identical probability vector. A change in the score is a target-switching diagnostic, not a calibration improvement or deterioration on a fixed prediction task.

Supplementary Note: Three-cohort origin-design sensitivity

All-origin, phase-zero and greedy non-overlapping estimates were recomputed from the original paired-event arrays for UK 2021 and Australia 2026, and from the newly generated UK 2022 arrays. The selection rules change the horizon and ten-minute phase composition as well as the number of observations, so the results should not be interpreted as estimates of one invariant target population.

Supplementary Table S23. DMPC deployment–reference contrasts under three prespecified origin-selection schemes.

Cohort	Reference	Selection	N	Gap (pp)	Discordance (%)
UK2021	FULL	All	9010	-6.25	16.50
UK2021	FULL	Hourly	1500	-5.20	16.00
UK2021	FULL	Non-overlap	752	-3.72	16.76
UK2021	Own plan	All	9010	+5.58	6.67
UK2021	Own plan	Hourly	1500	+5.27	6.87
UK2021	Own plan	Non-overlap	752	+5.98	7.58
Australia2026	FULL	All	2586	-12.34	18.83
Australia2026	FULL	Hourly	431	-13.23	20.65
Australia2026	FULL	Non-overlap	216	-12.50	19.91
Australia2026	Own plan	All	2586	+4.18	4.64
Australia2026	Own plan	Hourly	431	+5.57	5.57
Australia2026	Own plan	Non-overlap	216	+5.09	5.09
UK2022_Jan14	FULL	All	1985	-0.35	12.85
UK2022_Jan14	FULL	Hourly	330	+1.52	13.03
UK2022_Jan14	FULL	Non-overlap	166	+1.81	11.45
UK2022_Jan14	Own plan	All	1985	+6.80	7.20
UK2022_Jan14	Own plan	Hourly	330	+8.18	8.18
UK2022_Jan14	Own plan	Non-overlap	166	+7.23	7.23

The UK 2022 frozen-FULL point estimate changes from −0.35 pp over all origins to +1.52 pp over hourly origins and +1.81 pp over non-overlapping origins. A sign switch between designs is retained and not treated as a computational error.

Supplementary Table S24. Within-run circular moving-block sensitivity with 2,000 replicates per cell.

Cohort	Reference	Block (d)	Observed (pp)	Exploratory percentile range (pp)
UK2021	FULL	1	-6.25	[-10.88, -2.16]
UK2021	FULL	3	-6.25	[-12.00, -1.04]
UK2021	FULL	7	-6.25	[-12.73, -0.52]
UK2021	Own plan	1	+5.58	[+4.31, +6.89]
UK2021	Own plan	3	+5.58	[+4.28, +6.99]
UK2021	Own plan	7	+5.58	[+4.22, +6.96]
Australia2026	FULL	1	-12.34	[-19.41, -5.72]
Australia2026	FULL	3	-12.34	[-22.00, -2.94]
Australia2026	FULL	7	-12.34	[-22.35, -3.98]
Australia2026	Own plan	1	+4.18	[+2.63, +5.88]
Australia2026	Own plan	3	+4.18	[+2.47, +5.99]
Australia2026	Own plan	7	+4.18	[+2.75, +5.53]
UK2022_Jan14	FULL	1	-0.35	[-7.46, +4.84]
UK2022_Jan14	FULL	3	-0.35	[-6.80, +4.69]
UK2022_Jan14	FULL	7	-0.35	[-5.19, +4.53]
UK2022_Jan14	Own plan	1	+6.80	[+3.98, +9.77]
UK2022_Jan14	Own plan	3	+6.80	[+4.33, +9.32]
UK2022_Jan14	Own plan	7	+6.80	[+4.84, +8.87]

Ranges are simulation-based sensitivity summaries conditional on the fitted controller, available periods and circular resampling scheme. Wrapping creates artificial endpoint adjacency; short scored runs can remain fixed. None of these rows establishes independent-site population coverage. Original ordinary-block resampling and exact finite-sample centre diagnostics remain in Tables S4 and S15–S16.

Supplementary Methods: 2022 source construction and scope

The 2022 power source was the archived six-turbine Kelmarsh quality-screened aggregate. The predeclared calendar selection comprised 1–14 January 2022 UTC; an additional day of source values was available for feature/horizon construction. Input preparation retained source time stamps and normalised by 12.3 MW, using the frozen 2021 model, information gate and residual bank. Each eligible origin used only lagged, arrived information for prediction and scenario selection. Physical states and commands were evolved sequentially across the uninterrupted episode; no reset occurred between origins. The terminal incomplete horizons were excluded from paired event statistics but remained in chronological execution. Prior quality review precludes claims of fully blinded validation. All inputs, controller outputs, selection masks and reproducibility scripts are archived in Supplementary Data 4.

Supplementary Methods: Physical executor and complete DMPC program

The following specification describes the unchanged implementation in Supplementary Data 1, core_reproducibility/baselines/held_mpc.py. Variables refer to one planning origin and steps j = 0,…,H−1. The normalized time step is Δt = 1/6 h; η = √0.85; battery bounds are E_min = 0.11 and E_max = 0.99 pu h; battery power limit is 0.4 pu; directional reserve power is 0.2 pu; each direction receives 0.05 pu h per clock hour. Initial episode state is (0.55,0.05,0.05). Episode resets are simulation conventions, not charging events.

Held commands and local curtailment

For held charge c, discharge d, upward reserve u and downward reserve v, define the wind-side export ceiling L. Only curtailment q observes realized wind w in the same interval. The implementation is:

$$L_j = s_j + \varepsilon - d_j + c_j - u_j + v_j \tag{S2}$$

$$q_j = min\{1, max(w_j, 0), max(w_j - L_j, 0)\} \tag{S3}$$

Delivery is w − q + d − c + u − v. An interval shortfall is recorded when delivery − s < −(0.01 + 10^{-7}) pu. Excess energy integrates max(s − delivery − 0.01,0)Δt. Opposite commands and all before-state power/energy constraints are checked before execution. No grid-purchase price or degradation cost is included.

Linear-program objective and variables

For each j, decision variables are nonnegative c_j, d_j, u_j, v_j, predicted curtailment $\hat{q}_j$, shortfall slack z_j^s, overdelivery slack z_j^o and energy E_{j+1}. A nonnegative r represents absolute terminal-energy deviation. Point forecast $\hat{w}_j$ and obligation s_j are fixed inputs. The objective is:

$$min\, \Delta t \sum_{j=0,\ldots,H-1} [100(z_j^s + z_j^o) + c_j + d_j + 0.25(u_j + v_j)] + r \tag{S4}$$

There is no additional terminal coefficient: r has weight exactly 1. The constraints below are the complete LP specification, with hourly reserve indices defined on the next page.

$$E_{j+1} = E_j + \Delta t\left(\eta c_j - \frac{d_j}{\eta}\right), 0.11 \le E_{j+1} \le 0.99 \tag{S5}$$

$$0 \le c_j, d_j \le 0.4; 0 \le u_j, v_j \le 0.2 \tag{S6}$$

$$\eta \Delta t\, c_j \le 0.99 - E_j; \frac{\Delta t}{\eta} d_j \le E_j - 0.11 \tag{S7}$$

Supplementary Methods: LP constraints and controller menu

Forecast delivery, reserve and terminal constraints

Set $\hat{p}_j = \hat{w}_j - \hat{q}_j + d_j - c_j + u_j - v_j$. Predicted curtailment is bounded by the predicted available wind, and delivery deviations beyond $\varepsilon = 0.01$ are relaxed only by their nonnegative slacks:

$$0 \le \hat{q}_j \le min\left(1, max\left(\hat{w}_j, 0\right)\right) \tag{S8}$$

$$s_j - \varepsilon - z_j^s \le \hat{p}_j \le s_j + \varepsilon + z_j^o \tag{S9}$$

For every clock hour k intersecting the planning window, let J_k contain its steps. The first partial hour receives the current state's remaining directional reserve budgets R_0^u and R_0^v; subsequent hours receive 0.05 pu h in each direction. No budget is borrowed between directions or hours:

$$\Delta t \sum_{j \in J_k} u_j \le R_k^u ; \Delta t \sum_{j \in J_k} v_j \le R_k^v \tag{S10}$$

$$r \ge E_H - 0.55 ; r \ge 0.55 - E_H ; r \ge 0 \tag{S11}$$

The initial energy is fixed to the observed pre-action value. The LP is solved by HiGHS with presolve, a 5-s limit and primal/dual feasibility tolerances of 10^{-9}. The continuous relaxation is not asserted to enforce complementarity universally: any solution with min(c,d) or min(u,v) above 2×10^{-8} is rejected. A forecast replay must match the LP objective within 2×10^{-7}. Baseline runs permit a recorded FULL fallback on solver failure but observed none; the new S diagnostic stops on failure and permits no fallback. Before-state caps in equation (S7) are imposed in addition to energy dynamics, so simultaneous relaxed charging and discharging cannot create artificial within-step headroom.

F_FULL and the exact 225-member F_BR0 menu

The first-cap menu is the lexicographically indexed Cartesian product $\{0,0.25,0.5,0.75,1\}^2 \times \{0,0.5,1\}^2$, ordered as charge, discharge, upward reserve, downward reserve. Indices run from 0 to 224; 224 is full caps. At each forecast step, define surplus max(ŵ−s,0) and deficit max(s−ŵ,0). Charge and discharge are the smaller of the cap-adjusted feasible power and surplus/deficit. Upward reserve covers max(deficit−discharge−ε,0) up to its cap; downward reserve covers max(surplus−charge−ε,0). After the first command all cap fractions are one. Resource states use these held commands, not future observed wind.

F_BR0 assesses every candidate's resulting fixed command sequence on the same 32 residual paths. Scenario-dependent feedback is confined to local curtailment. It first minimizes the number of scenarios with any two-sided failure, then total failed steps, then the mean common proxy objective, then maximizes the sum of first-cap fractions and finally chooses the smallest candidate index. Proxy ties use a 10^{-12} tolerance, and cap-sum ties 10^{-14}. This is a restricted first-cap search, whereas DMPC optimizes the continuous full sequence. These are distinct policy classes with identical physical resources.

Supplementary Methods: Reproducible paired-replay algorithm

Baseline audit at a complete origin

1. Read only the forecast, scenarios, obligations and resource state allowed by the declared availability clock. Select the deployed controller's exact first four-component command a_t.

2. For P, retain the entire origin-optimized DMPC command sequence. For B, apply a_t first and construct a FULL tail using only the forecast frozen at the origin. Neither tail uses later forecast revisions.

3. Replay each held sequence from the same pre-action state on the same realized wind and original obligation window. Reset reserve budgets only at the same clock-hour boundaries. Use actual wind only in local curtailment and event scoring.

4. Reconstruct D from the saved chronological deployment commands for the identical window. Check that its first command and starting state equal those used for B and P. No commands from another origin's counterfactual state are spliced into an otherwise different deployment.

5. Mark the directional shortfall event if any active step crosses the tolerance. Exclude an origin only if its full window leaves the same contiguous episode; keep its executed step in operational accounting. Save both directions of discordance, target-matched probabilities and the resource ledger.

Additional fixed-original-endpoint continuation S

1. Fix the same origin, endpoint, obligation suffix, state and exact first command as above. Copy the archived first command without reoptimizing it.

2. For every later active step, update only to the arrived forecast at that step and the resource state produced by S. Solve the unchanged DMPC on the remaining original window. Do not add another delivery hour when the wall clock rolls over.

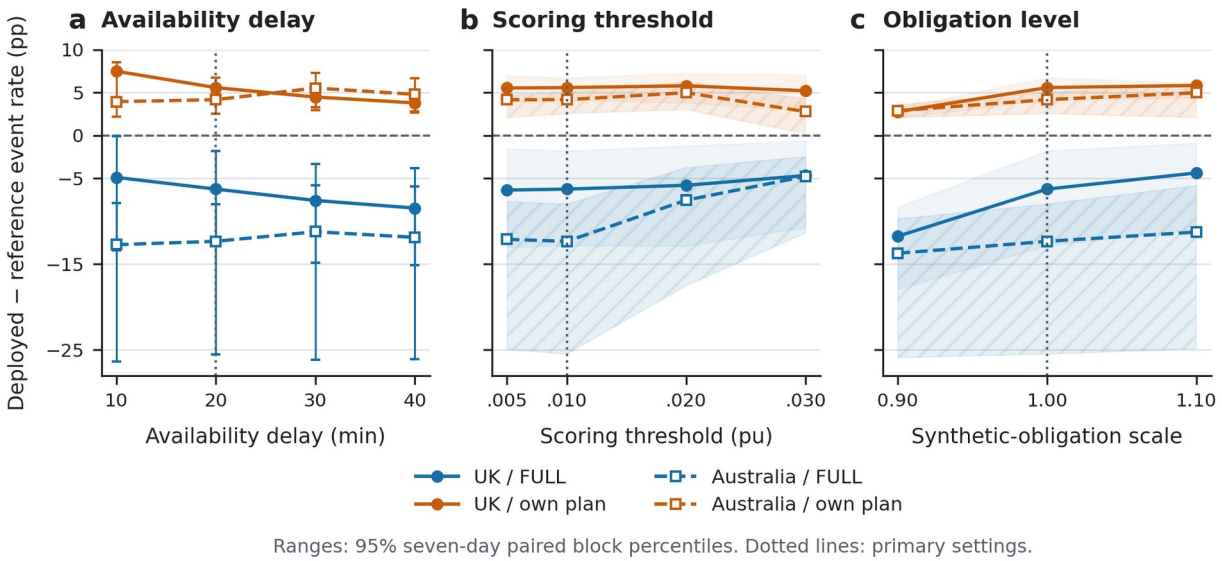
3. Execute only that newly selected first command using the common physical executor. Reject failure or invalid commands rather than silently dropping the origin.

4. Independently replay the final S command sequence with the scalar reference physics and compare every ledger value and event label. Check D/S command equality before their endpoint definitions first differ. Retain all original complete origins.

Release and verification scope

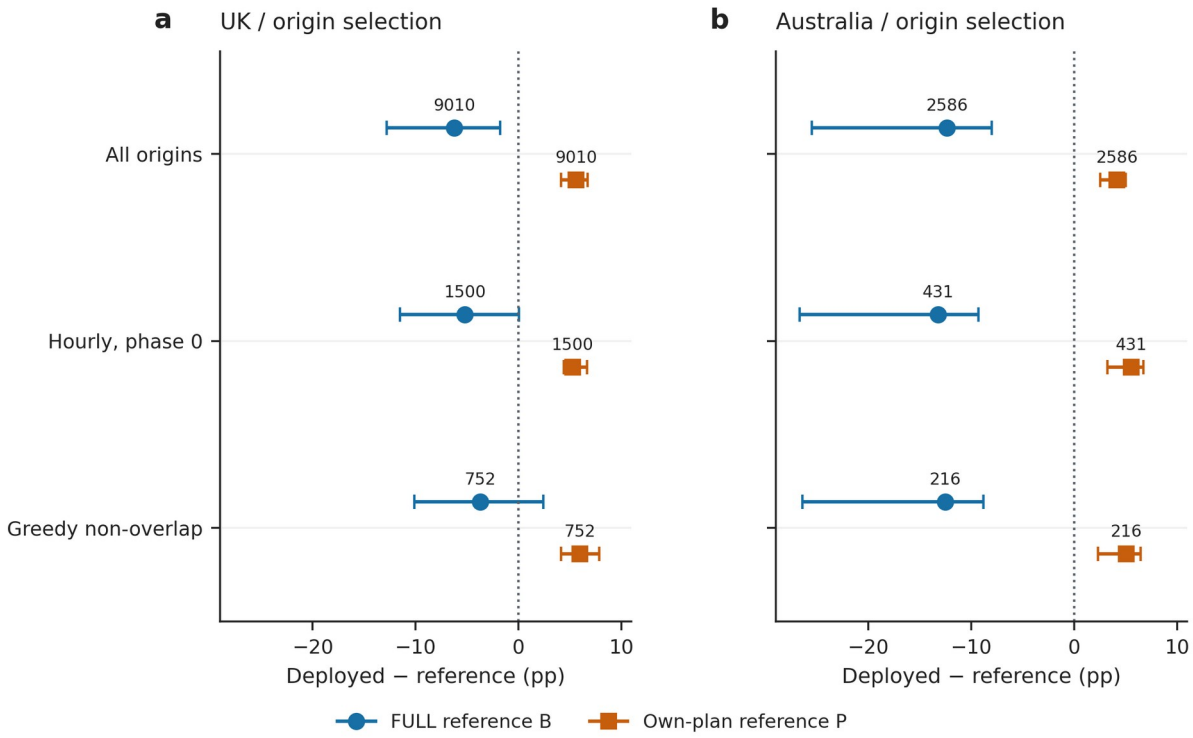
Supplementary Data 1 and Data 2 are unchanged source releases. Supplementary Data 3 contains the added drivers, pre-execution diagnostic specifications, frozen new outputs, CSV tables, figure sources and SHA-256 records. Those specifications were written after the primary findings were known and are not independent preregistrations. One portable command reruns numerical analyses from saved outputs; separate commands rerun the new controllers. The new work did not recollect AEMO data, refit forecasts, rerun the complete original latency/resource grids or create independent human verification. Original author-provided metadata and declarations require author review before submission.

Supplementary Figure S1



Supplementary Figure S1. Archived post-freeze sensitivity to availability gates, scoring tolerances and obligation levels. This is the original main Figure 4, moved here without changing values. All bands and bars are exploratory ordinary-block percentiles, subject to the qualifications in Tables S14–S17; point-estimate stability is not uniform exclusion of zero across all origin designs.

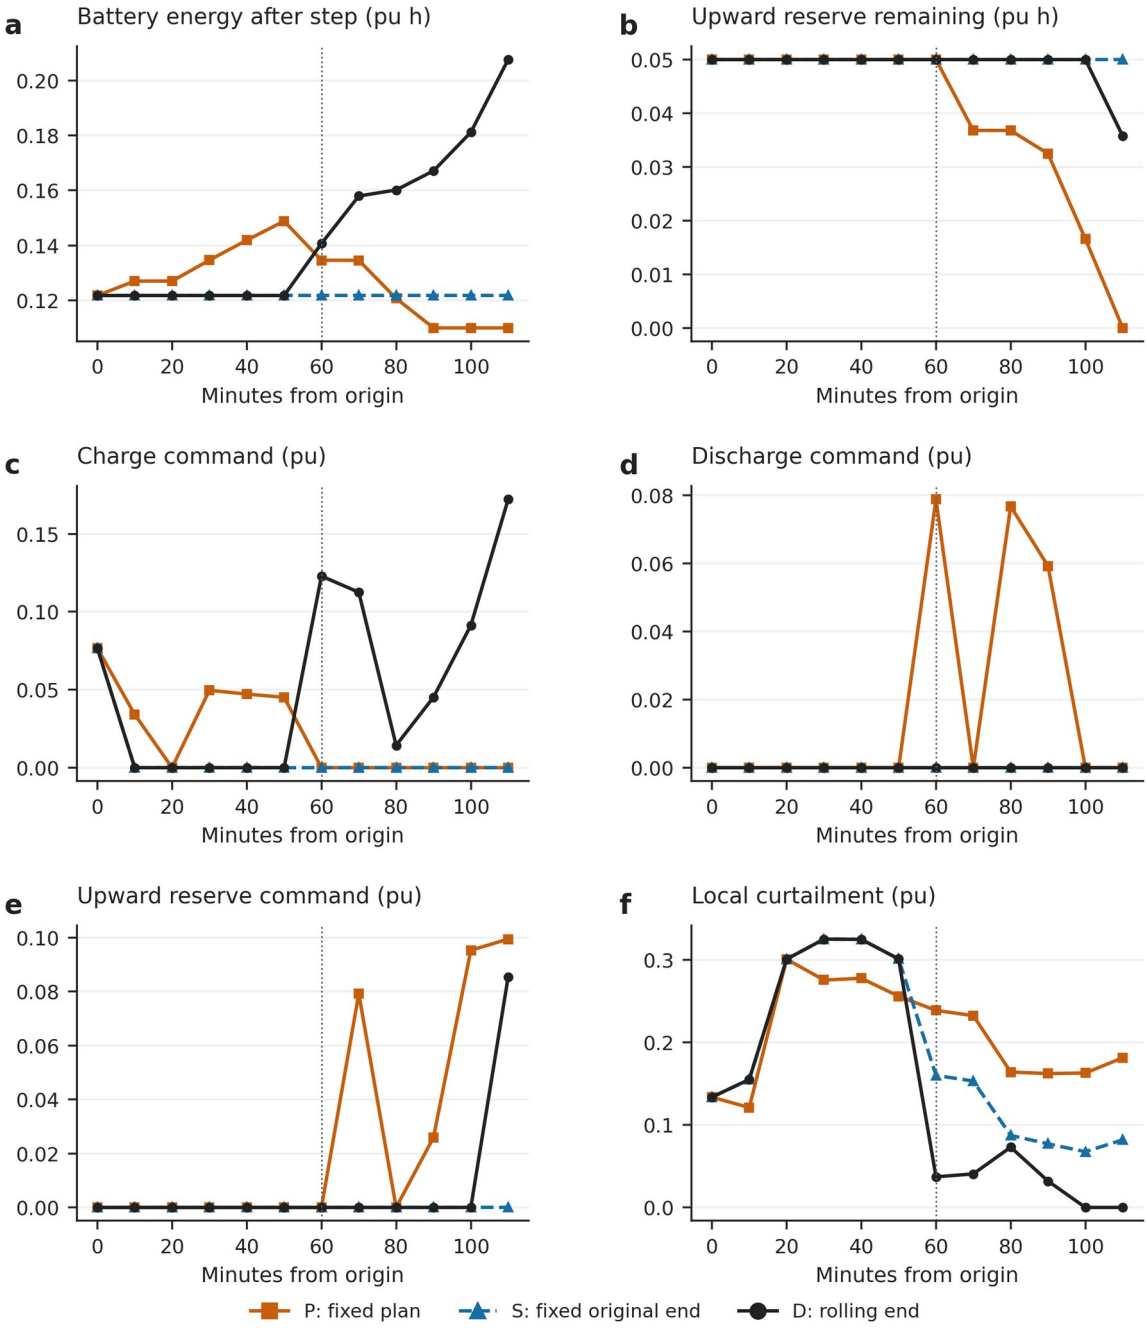
Supplementary Figure S2



Labels: sample counts. Ranges are exploratory; different origin designs change phase/window composition.

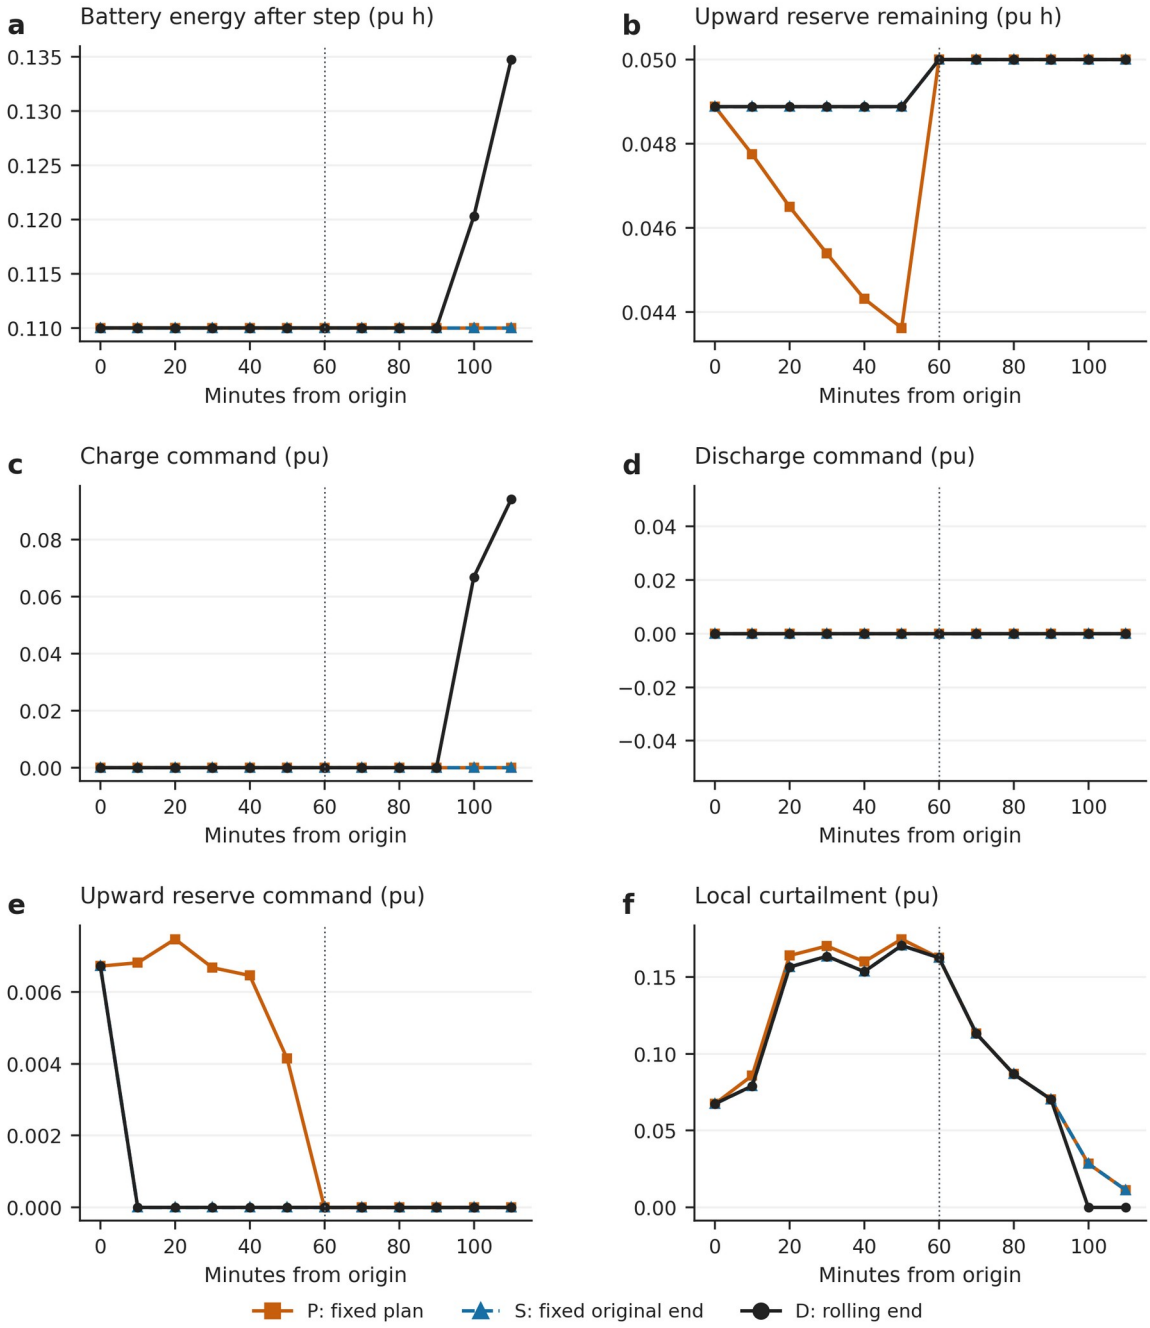
Supplementary Figure S2. All three origin-selection designs for each cohort and both DMPC references. Labels identify sample sizes. UK FULL-reference ranges cross zero for hourly phase-zero and non-overlapping selection. These sampling designs differ in window and phase composition as well as sample count.

Supplementary Figure S4



Supplementary Figure S4. UK resource/command trajectories for the chronologically earliest D-only-versus-S example (14 October 2021, 20:00 UTC). The original first command and all commands before the 60-min rollover agree between D and S. Charging introduced by rolling deployment after rollover contributes to the illustrated shortfall. P, S and B pass the scored window; D fails. This outcome-defined case is not a representative sample.

Supplementary Figure S5



Supplementary Figure S5. Australian resource/command trajectories for the chronologically earliest D-only-versus-S example (11 September 2026, 05:00 UTC). Rolling deployment charges late in the original window while S does not. All D/S commands agree before rollover. P, S and B pass the scored window; D fails. The deterministic selection is illustrative and does not estimate how often the mechanism occurs.