

Supplementary Material

Hollingham (2026) — Newborough Warren Coastal Sand Dune Aquifer, Isle of Anglesey

Contact: martin.hollingham+nrg@gmail.com

ORCID: [0000-0003-0253-9301](https://orcid.org/0000-0003-0253-9301)

Supplementary Note S1: Interactive Map

Supplementary File S1: Interactive map of the full dipwell monitoring network (Google Maps), accessible at:

https://www.google.com/maps/d/viewer?mid=1hXLAauiMeaVsXhBR_IoUTziAtjk

Supplementary Note S2: Aquifer Geometry, Hydraulic Parameterisation and the Basis of the Spatial Flow Model

S2.1 The Borehole Data Problem

Darcy's law and any spatially explicit groundwater flow model require three spatially resolved inputs: hydraulic conductivity (K), aquifer thickness (b , from which transmissivity $T = K \cdot b$ is derived), and hydraulic gradient (∇h). For the Newborough Warren dune system, only the third of these — the hydraulic gradient — is well constrained, through the dipwell network of over 70 wells. Hydraulic conductivity and aquifer thickness are poorly constrained, because borehole logs providing lithological or geophysical evidence of the depth to the bedrock or glacial till substrate are absent across the great majority of the 700 ha dune system.

The available borehole evidence is confined to four locations documented in Betson et al. (2002), a report commissioned by the Countryside Council for Wales (Contract FC†73-05-18) covering hydrogeological investigations at the site. These boreholes provide minimum depth-to-basement constraints at four spatial points and constitute the only direct subsurface lithological evidence available for the site. They are summarised in Table S1.

Borehole	Easting (OSGB36)	Northing (OSGB36)	Sat. thickness (m)	Source
Water borehole	241,721	364,133	12.8 (min.)	Betson et al. (2002)
NH1	241,734	363,306	6.5	Betson et al. (2002)
NH2	242,024	363,107	6.5	Betson et al. (2002)
Borehole 3	242,837	363,288	3.65	Betson et al. (2002)

Table S1. Borehole constraints on saturated thickness from Betson et al. (2002). Values represent confirmed minimum depth to the low-permeability substrate (glacial till or bedrock). The Water borehole value of 12.8 m is a minimum — drilling did not reach the substrate. The eastern borehole (Borehole 3) shows progressive thinning toward the Menai Strait. These four boreholes constitute the only direct subsurface lithological evidence available for the 700 ha dune system and represent the most important data gap for any future spatially explicit groundwater model.

The tracer test reported by Betson et al. (2002) provides the single direct estimate of hydraulic conductivity for the site, yielding $K = 6.0$ m/day at the central dune plain (range 3.0–9.0 m/day from the sensitivity analysis). This value is used as the site-wide K constant where transmissivity is required for the Darcy flow direction vectors in the spatial figures (§4.9.5, Figure 44), with the range providing the basis for the exploratory uncertainty envelope noted in the figure annotations.

S2.2 The β_1 Proxy for Relative Aquifer Thickness

The physical interpretation of β_1 is the fraction of a unit rainfall pulse that reaches the water table as a head rise, integrated over the monthly timestep. In an unconfined aquifer, this is governed by the depth of the unsaturated zone and the specific yield: a thick unsaturated zone above a large storage volume attenuates the rainfall-to-head transfer, producing a low β_1 . A shallow, low-storage aquifer (thin unsaturated zone, low S_y) produces a high β_1 — rainfall arrives at the water table quickly and raises it sharply. Although β_1 cannot be inverted algebraically to yield aquifer thickness without independent knowledge of the recharge flux and specific yield, its spatial pattern is a reliable proxy for the relative contrast in aquifer thickness across the site.

The Eastern Block (C1: $\beta_1 = 4.581$; C2: $\beta_1 = 3.871$) and Western Residual cluster (C3: $\beta_1 = 3.576$) contrast is consistent with the geological interpretation of Stratford et al. (2007): the Eastern Block sits on shallow till and estuarine deposits, producing a thin, storage-limited aquifer with a flashy response; the Western Residual Cluster occupies deep, clean aeolian sand, producing a capacious, buffered aquifer with an attenuated response. The Forest cluster (C4: $\beta_1 = 2.518$) is anomalously low even relative to C3, reflecting canopy interception rather than additional aquifer depth — confirmed by the depth-dependent PET analysis (Script 15) which found $\lambda = 0$ optimal for C4, indicating that capillary disconnection at depth is not the primary control. C5 (Coastal Forest: $\beta_1 = 2.436$) shows a comparable value to C4, consistent with both clusters carrying the same Corsican pine canopy on the same deep sand substrate.

Cluster	Label	β_1	β_2	$-\beta_3$	Geological context
C1	Lake Edge	4.581	0.963	0.090	Shallow till/estuarine substrate (NH1, NH2 boreholes: 6.5 m); rapid lake exchange
C2	Dune	3.871	1.745	0.063	Shallow till substrate consistent with C1; mature open dune
C3	Western Residual	3.576	1.806	0.060	Deep aeolian sand; Water borehole ≥ 12.8 m; DEM ridge geometry
C4	Main Forest	2.518	2.504	0.021	Same deep sand substrate as C3; low β_1 reflects 24% canopy interception
C5	Coastal Forest	2.436	1.374	0.045	Pine canopy on coastal sand; geomorphological thinning toward Menai Strait

Table S2. Cluster mechanistic coefficients (from Table 3, displacement-formulation SSM) and geological context. The β_1 contrast between Eastern Block (C1/C2) and Western Residual cluster (C3) is consistent with the stratigraphic interpretation of Stratford et al. (2007). C3 and C4 share the same deep aeolian sand body, confirmed by β_1 similarity after correcting for canopy interception and by Pearson affinity persistence post-felling. Borehole 3 (3.65 m, eastern margin) confirms progressive thinning toward the coast.

S2.2.1 Aquifer Thickness Surface (Developed for the PDE Model)

During development, a two-dimensional finite-difference PDE solver was implemented as an alternative spatial framework (see S3.3). That model required a spatially continuous aquifer thickness surface to compute transmissivity $T = K \cdot b$ at each grid cell. Because only four borehole constraints were available, an IDW-interpolated thickness surface was constructed using the borehole values as hard nodes supplemented by cluster-level thickness priors derived from the β_1 proxy argument above and per-well overrides at geological boundaries. The surface was interpolated on a 50 m grid using power-1 IDW with a physical minimum of 0.3 m enforced.

The PDE model was subsequently evaluated and rejected (S3.3) because it produced near-zero site-wide differences for forest management scenarios, primarily due to dilution of the modified β coefficients during grid interpolation, weak drainage feedback at C4, and the absence of ridge-derived lateral recharge from the source terms. The thickness surface is therefore not used in any published result. The per-well equilibrium framework adopted in the main paper (Section 3.8) operates directly from the fitted β coefficients at each well and does not require a thickness parameterisation. The Darcy flow direction vectors shown in Figure 44 are normalised head-gradient vectors derived from the interpolated mean head surface and are independent of aquifer thickness.

The borehole constraints in Table S1 and the geological contrasts summarised in Table S2 remain relevant for any future spatially explicit model. Slug tests at two to three representative wells per cluster and a ground-penetrating radar survey of aquifer thickness are identified as the highest-priority data gaps for future field campaigns (§5.9).

S2.2.2 Per-Well Geological Constraints

Several wells occupy geological settings that depart from their cluster's typical substrate. These are noted here for reference, as they would serve as boundary constraints for any future spatial model.

CEH14 (241,292 E, 364,488 N), classified as C4, sits at the crest of the bedrock ridge at 14.4 m AOD with a mean water table head of 13.3 m AOD. The water table is approximately 1.1 m below ground surface here. CEH14 overlies irregular bedrock topography on the elevated ridge flank, where buried ridges impede lateral drainage — consistent with its anomalous negative β_3 under the depth-from-surface formulation and its carrying the largest water balance residual in the network (+0.031 m/month; §4.9.6, Figure 45).

Wells CEH7 (243,386 E, 363,613 N) and CEH8 (243,150 E, 363,382 N) sit at the far eastern estuarine margin of the site, where the DEM confirms ground elevation of 1.2–5.2 m AOD. Both represent coastal pinch-out positions where the aquifer thins to approximately 4.0 m. and their records are intermittent.

S2.3 Why C4 Retains Its Cluster Identity After Clearfell

A key question for the scenario analysis is whether the clearfell treatment wells (FE1, FE2, FE4, LIS1) converged toward C3 open-dune behaviour after the December 2017 felling. The Pearson affinity analysis (Section 4.3) provides a direct empirical test: all four impact wells maintained $r > 0.97$ affinity with C4 even in the post-felling period (2018–2026), with no significant shift toward C3 affinity.

This persistence is consistent with the interpretation that the Forest cluster signature at these wells reflects deep sandy substrate at high topographic elevation rather than canopy

influence alone — in which case convergence toward open-dune behaviour may not be achievable regardless of management intervention. C4 wells sit at the highest topographic positions in the dune system (mean head 9.52 m AOD), and their deep unsaturated zones and low recharge sensitivity reflect their geological position, not only their canopy cover. This conclusion is independently supported by the NW10 broadleaf comparison (§5.6.3): NW10, situated within the 1993 clearfell and 1996 broadleaf restocking block, has maintained C4 cluster affinity ($r = 0.986$) throughout the 18-year monitoring period despite its different canopy cover, confirming that substrate position rather than canopy type is the dominant control on cluster identity.

S2.5 Reproducibility

The spatial model parameterisation documented in this note is implemented in `19_spatial_groundwater.py` (scenario viewer data preparation, IDW thickness surface, per-well equilibrium calculations) and `20_spatial_figures.py` (Darcy flow direction vectors, mean head surface, water balance residual field). The β_1 proxy analysis draws on the cluster-level SSM coefficients exported by `03_state_space_model.py` (`03_03_cluster_mechanistic_coefficients.csv`). The Pearson affinity test for C4 cluster persistence post-felling is computed in `05_pearson_affinity.py`. All scripts read intermediate data produced by the main pipeline and are maintained in the canonical 00–21 sequence. The interactive scenario viewer (`scenario_viewer.html`) is a standalone browser application that reads the per-well coefficient and location CSVs exported by Script 19.

Supplementary Note S3: Scenario Modelling Framework and Limitations

S3.1 Purpose and Scope of the Scenario Analysis

An exploratory scenario analysis was developed to investigate whether the site's own SSM β coefficients could yield defensible quantitative predictions of water table response to management interventions and climate change. This note documents the scenario framework adopted (the per-well equilibrium framework described in Section 3.8), the alternative PDE-based approach that was evaluated and rejected during development, and the structural limitations that apply to any equilibrium representation of the system. The scenario outputs themselves are reported in Section 4.10 and discussed in Section 5.5.2 of the main paper; this note provides the technical reasoning behind the methodological choices and their limitations.

S3.2 The Per-Well Equilibrium Framework

The scenario framework adopted in this study is a well-level equilibrium calculation that applies scenario-specific perturbations to the SSM forcing terms (rainfall, PET, canopy interception, atmospheric draw) and computes the resulting change in equilibrium head for each reference well. Per-well Δh values are then interpolated to a 50 m grid on the British National Grid by Delaunay triangulation with linear barycentric weighting (`scipy.interpolate.griddata, method="linear"`), masked to the rectangular sea-boundary extent; this produces the elevation-based maps presented in Figures 32 and 33. The accompanying interactive scenario viewer (`scenario_viewer.html`) renders the same per-well Δh field using power-1 eight-nearest-neighbour inverse-distance weighting for responsiveness under slider-driven re-interpolation, and additionally offers a depth-below-surface view (DEM minus interpolated head, per cell) with dune-ridge cells masked — see S3.7 for the viewer-specific rendering conventions. The full per-scenario \times per-season \times per-cluster Δh output is tabulated in the supplementary file `19_scenario_summary.csv`.

The scenario head change for each well is computed directly from the fitted SSM equation as

$$\Delta h = (\beta_1 \cdot P_{eff,sc} - \beta_2 \cdot PET_{sc} - \beta_3 \cdot |h|) - (\beta_1 \cdot P_{eff,0} - \beta_2 \cdot PET_0 - \beta_3 \cdot |h|)$$

where the $\beta_3 \cdot |h|$ term cancels algebraically between scenario and baseline expressions, so β_3 does not materially affect the computed Δh . Specific yield does not appear in the equation because the fitted β coefficients already embed the aquifer storage response through the SSM fit to observed head data; the per-well WTF-derived S_y values (Section 3.7.3) are retained as supplementary display information in the interactive viewer, not as a divisor on Δh . The framework is honest about what it is: a geometric aggregator of SSM responses across the cluster structure, not a physical model of groundwater flow.

The per-well framework has two important structural features that shape the interpretation of its outputs. First, scenario parameter changes for canopy interception and β_2 apply only to the C4 and C5 Forest cluster wells, because that is where the relevant physical changes occur under forest management. The framework therefore necessarily produces a response confined to the C4/C5 zone and its immediate interpolated surroundings, with effectively no propagation into C1, C2, or wider C3 clusters. Second, climate scenario parameter changes (ΔP , ΔPET) apply to all reference wells simultaneously, so the resulting Δh field is spatially continuous and reflects the site-wide β coefficient structure. The contrast between forest management and climate scenarios in the spatial outputs is therefore a direct consequence of which wells are affected by which perturbations, not an artefact of the numerical method.

S3.3 Evaluated-and-Rejected: The 2D Steady-State PDE Model

During development, an alternative spatial framework was implemented as a two-dimensional finite-difference PDE solver based on the Helmholtz form of the steady-state groundwater flow equation:

$$\nabla \cdot (T \cdot \nabla h) - \beta_3 \cdot h = -(\beta_1 \cdot P_{eff} - \beta_2 \cdot PET)$$

with transmissivity $T = K \cdot b$ ($K = 6.0$ m/day after Betson et al., 2002; b IDW-interpolated from four borehole constraints supplemented by cluster-level priors derived from the β_1 proxy, also after Betson et al., 2002). Dirichlet head conditions ($h = 0$ m AOD) were applied at sea boundaries, and implicit Neumann no-flow conditions at the ridge. The sparse linear system of approximately 4,200 equations was solved using `scipy.sparse.linalg.spsolve` (Virtanen et al., 2020).

This approach is well-posed mathematically — the Helmholtz structure guarantees a unique solution with the given boundary conditions — and was physically consistent in that the β coefficients entering the source term are the same parameters estimated from the observed well records. The solver produced physically plausible baseline head surfaces and Darcy flux fields qualitatively similar to those obtained from the per-well framework.

However, the PDE solver produced near-zero site-wide differences for forest management scenarios (mean $\Delta h = -0.003$ to -0.008 m for full clearfell and thinning), for three structural reasons.

First, the forest management scenarios modify β coefficients only for the C4 and C5 Forest wells. These per-well β values are then IDW-interpolated to the 4,235-cell grid. Because the forest wells constitute a small minority of the network of over 70 wells, the modified β values are heavily diluted by the surrounding clusters during interpolation. The source term $\beta_1 \cdot P_{eff} - \beta_2 \cdot PET$ therefore changes by only a small amount at each cell, and the resulting change in the steady-state head solution is proportionally small.

Second, the internal drainage term $\beta_3 \cdot h$ provides a stabilising feedback in the Helmholtz equation. C4 has the lowest β_3 in the network (0.021, compared with 0.060 for C3 and 0.063 for C2; Table 3). This means that in the C4 zone the drainage feedback is weak, and a moderate source term change produces only a small equilibrium head change.

Third, and most fundamentally, the steady-state PDE does not capture the ridge-derived boundary subsidy that sustains the C4 water table. The observed mean head at C4 (9.52 m AOD) is substantially higher than the mean at C3 (6.09 m AOD), not because of C4's β coefficients but because the C4 wells are proximal to the bedrock ridge and receive lateral recharge from the ridge slope. CEH14, the well most proximal to the ridge, carries the largest persistent water balance residual in the network (+0.031 m/month; §4.9.6, Figure 45). The steady-state PDE source terms are parameterised from the SSM β coefficients alone and do not include this ridge contribution. As a result, the PDE solution underestimates the true C4 water table under both baseline and scenario conditions by a systematic amount related to the ridge contribution, and the difference between the two solutions is correspondingly small.

A fully calibrated continuous-flow model capable of representing the ridge boundary subsidy dynamically would require, at minimum: (i) slug tests at two to three representative wells per cluster to constrain the spatial distribution of K (currently constrained only by the single tracer test of Betson et al., 2002); (ii) a ground-penetrating radar survey of aquifer thickness to replace the indirect β_1 -proxy constraints with direct measurements; and (iii) instrumentation of the ridge boundary subsidy pathway sufficient to parameterise a time-varying boundary flux term. These are the same parameter-sparsity constraints that produced calibration difficulties in the Betson et al. (2002) MODFLOW model at this site.

For these reasons, the PDE approach was judged to overclaim relative to the data available and was rejected in favour of the per-well equilibrium framework. The framework represents what the SSM parameterisation supports without claiming spatial flow dynamics the monitoring data cannot constrain.

S3.4 Climate Scenarios in the Per-Well Framework

Climate scenarios applied under the per-well framework are based on the UKCP18 Regional 12 km projections for Wales (Met Office, 2018) and the CHES-SCAPE bias-corrected projections of Robinson et al. (2023). Two scenarios are carried forward into the Results (Section 4.10): UKCP18 2050s (2040–2069) applies a winter precipitation increase of +10%, summer precipitation decrease of –15%, winter PET increase of +5% and summer PET increase of +20%; UKCP18 2080s (2070–2099) applies +20%, –30%, +10% and +35% respectively for the same four terms. Both scenarios use RCP8.5 central-estimate (50th-percentile) values. Perturbations are applied per season, with winter defined as November–March and summer as May–September; the climatological baselines against which the perturbations are applied are the 2005–2026 monthly means from the RAF Valley climate record. The resulting Δh for each well is computed for each season separately, and the annual Δh is reported as the 0.5-weighted mean of the winter and summer equilibrium responses — this convention captures the physically significant seasonal asymmetry between winter recharge and summer evaporative loss that an annual-mean forcing would obscure.

Three limitations on the interpretation of these outputs apply. First, the perturbations used are central-estimate values; the UKCP18 probabilistic range at the 10th and 90th percentiles spans roughly half to twice the central perturbation for precipitation and a broader range still for PET, so the Δh values reported here are a central point within a substantially wider

envelope of plausible responses. Users of the interactive scenario viewer (`scenario_viewer.html`) can explore perturbations across the full 0.5–1.5× range of the seasonal sliders to visualise this envelope. Second, the equilibrium framework computes the steady-state response to sustained seasonal perturbations and does not resolve within-year dynamical propagation of recharge between winter and summer — for example, the effect of a wetter winter on the summer minimum through carry-over storage is captured only to the extent that the SSM β_3 drainage coefficient encodes it in the fitted mean. Third, the framework holds the SSM β coefficients themselves fixed at their historical values; it does not allow for structural changes in the β coefficients under future climate (for example, a higher β_2 under consistently warmer summers through increased vapour pressure deficit). Within these limitations, the climate scenario output is the most quantitatively defensible use of the equilibrium framework: every cluster's response is derived from its own fitted β coefficients under a uniform climate perturbation that reflects the current best estimate for the site's regional climate trajectory, and the result reflects the physical partition of rainfall and PET sensitivity that the SSM has estimated from 21 years of observation.

S3.5 Forest Management Scenarios in the Per-Well Framework

Forest management scenarios (full clearfell, 50% thinning, broadleaf conversion) apply canopy parameter changes only to the C4 and C5 Forest wells, and the framework correctly produces a response confined to the C4/C5 zone. This is a feature of the framework, not a bug: under the SSM parameterisation, canopy management directly affects only the wells beneath the canopy, and there is no structural mechanism for canopy effects to propagate into the open dune clusters beyond the hydraulic gradient already implicit in the baseline heads. The BACI monitoring record provides empirical corroboration of this structural feature: post-felling displacement is concentrated at the core impact wells, and wells outside the C4 cluster show no detectable felling-specific response once common-mode climate variability is accounted for (§4.6.4).

The clearfell scenario applies a β_2 multiplier of $\times 1.108$ at C4 and C5 wells, derived dynamically from the BACI-corrected Edge-tier post-felling β_2 ratio reported in §4.6.6 (Table 10): Edge mean ratio 1.101 minus Climate Control mean drift 0.993, plus 1.0. Canopy interception is set to 0%. The framework predicts modest annual water table rises at both forest clusters (C4: +0.028 m/month head, equivalent to +4.7 mm w.e./month; C5: +0.034 m/month head, +8.8 mm w.e./month), with a strongly asymmetric seasonal profile reflecting the dominance of summer PET in the β_2 sacrificial shielding mechanism established in §4.6.6 and §5.5.2.

The thinning scenario applies a β_2 multiplier of $\times 1.054$ (half the clearfell perturbation) with canopy interception reduced from 24% to 12%. The framework predicts approximately half the clearfell response at both forest clusters (C4: +0.014 m/month, +2.3 mm w.e./month; C5: +0.017 m/month, +4.4 mm w.e./month), as expected.

The broadleaf conversion scenario is represented by replacing the Corsican pine canopy interception fraction of 24% (Freeman, 2008) with an annual-mean deciduous interception of 15% (following the temperate deciduous meta-analysis of Komatsu et al., 2011) while holding annual-mean β_2 at baseline. No annual-mean justification exists for a year-round β_2 adjustment under broadleaf: the canopy is leafed in summer and leafless in winter, and the annual mean β_2 is expected to be comparable to pine. The framework predicts a near-neutral annual response at both forest clusters (C4: +0.010 m/month head; C5: +0.012 m/month head), masking a pronounced seasonal asymmetry: winter head-space values are the largest of any forestry scenario, but the summer response is negative at C4 (−0.007 m/month) as the growing-season transpiration penalty under mature deciduous canopy (β_2 summer multiplier $\times 1.113$) consumes the interception gain. The net annual volumetric effect

is -3.4 mm w.e./month at C4 and $+0.1$ mm at C5. The seasonal asymmetry that underlies the phenological mechanism described in §5.5.2 — larger winter recharge pulse under leafless canopy, steepened hydraulic gradient, accelerated drainage into summer — is a dynamical trajectory response that the equilibrium framework does not resolve.

For the Figure 46 hydrograph (§5.7.4), a seasonally-varying β_2 is applied to the broadleaf scenario to represent the full-LAI summer canopy and leafless winter canopy. This seasonal parameterisation is a visualisation device and is not part of the §4.10 Results; it is used specifically to communicate the phenological argument from §5.5.2 alongside the BACI-observed record. The annual-mean output in `19_scenario_summary.csv` remains the quantitatively citeable framework result.

S3.6 Implications and Future Work

The per-well equilibrium framework is a lightweight analytical device for translating the SSM β coefficients into spatial summaries of equilibrium head response. Its strengths are that it makes no claims beyond what the SSM itself supports, that it is fully reproducible from the published CSV outputs, and that it is explorable through an interactive viewer that exposes all parameter choices to the user. Its limitations are those of any equilibrium framework: it does not resolve dynamical trajectory responses, it does not capture cluster-to-cluster hydraulic coupling, and it cannot simulate transient interventions in which the time evolution of the water table matters for ecological outcome. The framework is adequate for the questions the paper addresses — whether forest management in the C4/C5 zone affects summer minima in C1 and C2, and how climate perturbations partition across the cluster structure — and the BACI monitoring record confirms the framework's structural conclusions.

Answering a different class of question, such as how a hypothetical intervention along the CEH14 ridge subsidy pathway would propagate through the aquifer, would require a calibrated continuous-flow model. For that, the critical unknowns are hydraulic conductivity (currently constrained by a single tracer test at $K = 6.0$ m/day; Betson et al., 2002) and aquifer thickness (inferred from four borehole logs across 700 ha). Slug tests at two to three representative wells per cluster and a ground-penetrating radar survey of aquifer thickness are identified as the priority field measurements required to enable a properly calibrated continuous-flow model, should future work pursue intervention scenarios where lateral hydraulic propagation is central to the management question.

Summary of scenario modelling approach

Primary framework: Per-well equilibrium with triangulation-linear grid rendering (Sections 3.8, 4.10; `scenario_viewer.html`; `19_scenario_summary.csv`)

Evaluated and rejected: 2D steady-state Helmholtz PDE solver (Section 3.8 PDE paragraph; S3.3 above)

Visualisation device: Figure 46 synthetic hydrograph with seasonally-varying broadleaf β_2 (§5.7.4; Script 21)

Empirical corroboration: BACI monitoring record (§4.6)

S3.7 Interactive Viewer: Rendering Conventions and Depth-Mode Visualisation

The companion interactive scenario viewer (`scenario_viewer.html`; Hollingham, 2026b) renders the same per-well Δh field that underpins the published Figures 32 and 33 of the main paper, but uses a distinct grid-rendering method at the browser to preserve responsiveness under slider-driven re-interpolation. The viewer implements inverse-distance weighting with power 1 and eight nearest neighbours (10 m distance floor; 5 m exact-value shortcut) rather than the triangulation-linear method used for the static figures. Power-1 IDW

with a fixed neighbour set produces a smoother surface than the naïve power-2 IDW previously used, removes the bullseye halo artefacts that arose around individual wells on the Δh surface, and is sufficiently close to the triangulation-linear rendering at well locations that the two methods agree to within the mapped colour-scale resolution across the interpolated interior. At well locations the viewer returns the well's observed value unchanged; outside the convex hull of the wells it extrapolates under the IDW kernel where the triangulation-linear method would return no data. This distinction is most visible near the south-eastern coastal margin, where the viewer renders a short extrapolated zone that the static figures correctly leave blank.

The viewer offers three complementary display modes. The first two — change in water table elevation (Δh) and absolute water table elevation (m AOD) — are elevation-based surfaces and match the quantity displayed in Figures 32 and 33 respectively. The third mode, depth below surface (m), is a viewer-specific visualisation not reproduced in the static figures. It is computed per grid cell as $d = h_DEM(E, N) - h_WT(E, N)$, where h_DEM is the bilinearly-interpolated LiDAR DEM elevation at cell centre and h_WT is the IDW-interpolated scenario water table elevation at the same cell. This per-cell subtraction is the quantity that mediates the ecological response: the Curreli et al. (2013) slack vegetation thresholds are defined in depth-below-surface terms (SD15b wet slack summer viability limit: 0.61 m; SD16 dry slack viability limit: 0.98 m), and rendering the scenario water table against the topographic surface permits direct visual identification of the slack positions where scenarios cross these thresholds. The viewer's depth-mode colour scale is anchored on these thresholds: light blue at the 0.10 m winter flooding limit, green at 0.61 m, orange at 0.98 m, and deep red at 1.50 m, with explicit black tick marks on the legend at the SD15b and SD16 threshold depths.

A dune-ridge mask is applied exclusively in depth-below-surface mode. The rationale is that the water table is a continuous surface in space, and its elevation relative to ordnance datum is well-defined at every point within the site regardless of whether the overlying ground is a ridge or a slack floor; elevation-based maps (modes one and two) therefore require no topographic masking. Depth below surface, however, is meaningful only where the ground is approximately at the well-plane elevation — at a dune ridge the interpolated water table sits several metres below ground, but this is a consequence of the ridge height, not of any hydrological feature, and rendering such cells would produce a misleading strong-red patch on the depth map that obscures the ecologically relevant slack-floor response. The viewer masks a cell in depth mode when the DEM elevation at that cell exceeds the IDW-interpolated well-plane elevation (computed by the same IDW kernel applied to well DEM ground elevations) by more than 1.0 m. This threshold is user-toggleable from the viewer's control panel; disabling it reveals the unmasked depth field, which can be useful for diagnostic inspection but should not be used for ecological interpretation.

Additional viewer conventions of note: the basemap is a LiDAR-derived greyscale hillshade rendered at approximately 35% opacity underneath the scenario surface, matching the static figure hillshade convention of `map_utils.load_dem_hillshade` with azimuth 315°, altitude 35° and vertical exaggeration 3; climate perturbations can be explored across the full 0.5–1.5× range per season, extending substantially beyond the UKCP18 2050s and 2080s central-estimate presets to allow sensitivity analysis across the probabilistic envelope; and tooltips on any well display both the baseline and scenario-perturbed depth values alongside the absolute head and Δh , providing a direct visual link between the per-well framework output and the ecological depth-below-surface interpretation.

S3.8 Reproducibility

The per-well equilibrium framework and climate/forestry scenario calculations are implemented in `19_spatial_groundwater.py`, which exports the full per-scenario × per-season × per-cluster output to `19_scenario_summary.csv` and generates the elevation-based scenario maps (Figures 32 and 33). The interactive scenario viewer (`scenario_viewer.html`) implements the browser-side IDW rendering described in S3.7 and reads the same per-well data. The synthetic mean-year hydrograph (Figure 46) is generated by `21_forestry_scenarios.py`, which reads BACI coefficients from the Script 10 clearfell suite outputs. All scenario parameters — UKCP18 perturbation fractions, canopy interception values, β_2 multipliers, and seasonal definitions — are defined as named constants at the head of each script and are fully documented in the `scenario_viewer.html` source.

Supplementary Note S4: Specific Yield Mapping

Table S4. Individual Well WTF Specific Yield Estimates — Newborough Warren 2005–2026: Event-based median specific yield (S_y) derived from the water table fluctuation (WTF) method (Healy and Cook, 2002) for 67 reference wells and 19 extended network wells. n events = number of qualifying monthly rising-limb events (criteria: $\Delta h > 5$ mm, net recharge $P - PET > 10$ mm). Q25/Q75 = interquartile range; wide IQR reflects month-to-month variability in event estimates rather than measurement error. Forest cluster (C4) wells marked Interception corrected = Yes have net recharge adjusted for 24% canopy interception (Freeman, 2008). Extended well cluster assignments derived from Pearson affinity sitewide audit (Section 4.3). *Excluded from IDW interpolation surface: CEH12 (bedrock ridge — WTF response reflects fractured rock); CEH15 (forest slack floor — slack topography dominates water table dynamics). Generated by `18_wtf_spatial.py`.

Well	Cl.	Net	n	Sy med	Sy Q25	Sy Q75	Int.corr.	Excl.
CEH11	1	Ref	50	0.223	0.167	0.322		
CEH22	1	Ref	26	0.313	0.178	0.366		
CEH23	1	Ref	34	0.192	0.121	0.275		
CEH25	1	Ref	49	0.172	0.142	0.286		
CEH26	1	Ref	39	0.223	0.130	0.315		
CEH27	1	Ref	43	0.212	0.150	0.314		
CEH5	1	Ref	52	0.192	0.139	0.323		
CEH6	1	Ref	55	0.211	0.142	0.309		
P2	1	Ext	17	0.238	0.106	0.296		
CEH10	2	Ref	50	0.245	0.119	0.328		
CEH24	2	Ref	43	0.216	0.144	0.323		
CEH28	2	Ref	35	0.227	0.146	0.374		
CEH40	2	Ref	22	0.271	0.168	0.321		
CEH41	2	Ref	24	0.222	0.132	0.289		
D10	2	Ref	37	0.235	0.156	0.369		
D15	2	Ref	36	0.198	0.143	0.308		
D17	2	Ref	42	0.251	0.175	0.338		
D38	2	Ref	47	0.216	0.144	0.307		

D41	2	Ref	37	0.227	0.154	0.377		
D43	2	Ref	45	0.226	0.158	0.340		
D44	2	Ref	45	0.237	0.179	0.355		
D5	2	Ref	40	0.259	0.140	0.382		
D6	2	Ref	48	0.267	0.158	0.415		
D7	2	Ref	42	0.213	0.128	0.333		
D8	2	Ref	36	0.219	0.140	0.367		
D9	2	Ref	38	0.211	0.153	0.354		
L7	2	Ref	47	0.212	0.139	0.299		
NW 4B	2	Ref	42	0.212	0.147	0.320		
NW13	2	Ref	36	0.207	0.127	0.266		
NW3	2	Ref	47	0.254	0.160	0.359		
NW4	2	Ref	42	0.233	0.158	0.333		
T41A	2	Ref	43	0.201	0.136	0.338		
T41B	2	Ref	41	0.212	0.163	0.397		
T41C	2	Ref	40	0.253	0.167	0.365		
T41D	2	Ref	36	0.214	0.157	0.320		
WMC1	2	Ref	35	0.232	0.135	0.350		
WMC4	2	Ref	36	0.199	0.117	0.278		
CEH12	2	Ext	8	0.366	0.338	0.380		Yes*
CEH35	2	Ext	8	0.218	0.120	0.330		
CEH37	2	Ext	28	0.197	0.128	0.308		
CEH38	2	Ext	13	0.303	0.170	0.450		
CEH7	2	Ext	22	0.209	0.144	0.318		
CEH8	2	Ext	25	0.246	0.151	0.348		
PE	2	Ext	16	0.243	0.182	0.299		
PW	2	Ext	17	0.220	0.136	0.229		
CEH 1	3	Ref	49	0.232	0.141	0.306		
CEH16	3	Ref	31	0.250	0.161	0.342		
CEH17	3	Ref	25	0.282	0.161	0.389		
CEH18	3	Ref	38	0.260	0.150	0.354		
CEH19	3	Ref	24	0.360	0.180	0.391		
CEH21	3	Ref	24	0.307	0.166	0.378		
CEH31	3	Ref	26	0.275	0.205	0.392		
CEH36	3	Ref	28	0.290	0.162	0.361		
CEH39	3	Ref	27	0.263	0.152	0.378		

CEH4	3	Ref	39	0.243	0.166	0.339		
CEH42	3	Ref	27	0.254	0.142	0.359		
CEH9	3	Ref	40	0.275	0.166	0.389		
D25	3	Ref	35	0.264	0.189	0.355		
NW1	3	Ref	47	0.173	0.118	0.282		
NW11	3	Ref	47	0.203	0.129	0.270		
NW2	3	Ref	47	0.204	0.115	0.296		
NW5	3	Ref	48	0.254	0.144	0.337		
NW6	3	Ref	41	0.237	0.188	0.309		
NW7	3	Ref	42	0.265	0.180	0.358		
NW9	3	Ref	34	0.282	0.166	0.339		
WMC2	3	Ref	38	0.234	0.176	0.348		
WMC3	3	Ref	34	0.258	0.226	0.382		
FE3	3	Ext	19	0.263	0.165	0.332		
NW8	3	Ext	14	0.233	0.140	0.357		
NW8B	3	Ext	26	0.266	0.222	0.363		
P1	3	Ext	7	0.313	0.225	0.394		
CEH13	4	Ref	38	0.237	0.143	0.333	Yes	
CEH14	4	Ref	40	0.154	0.094	0.290	Yes	
CEH2	4	Ref	42	0.182	0.116	0.252	Yes	
CEH20	4	Ref	37	0.197	0.133	0.271	Yes	
CEH30	4	Ref	34	0.223	0.138	0.293	Yes	
CEH32	4	Ref	35	0.197	0.143	0.274	Yes	
CEH33	4	Ref	35	0.221	0.148	0.286	Yes	
CEH34	4	Ref	35	0.196	0.124	0.278	Yes	
NW10	4	Ref	42	0.210	0.111	0.284	Yes	
CEH15	4	Ext	16	0.141	0.101	0.364	Yes	Yes*
FE1	4	Ext	24	0.196	0.109	0.267	Yes	
FE2	4	Ext	24	0.217	0.128	0.285	Yes	
FE4	4	Ext	20	0.228	0.160	0.281	Yes	
LIS1	4	Ext	20	0.183	0.129	0.258	Yes	
NW12	4	Ext	9	0.211	0.161	0.293	Yes	

S4.1 Reproducibility

The WTF specific yield estimates tabulated in Table S4 are computed by `17_wtf_specific_yield.py`, which identifies qualifying monthly rising-limb events, applies the canopy interception correction for C4 wells, and exports per-well S_y statistics to `17_01_wtf_sy_estimates.csv`. The spatial IDW interpolation of S_y values and the exclusion of CEH12 and CEH15 from the interpolation surface are implemented in `18_wtf_spatial.py`, which generates the S_y map figure and exports the interpolated surface. Both scripts read intermediate data from the main pipeline (`01_wells_clean.csv`, `01_climate.csv`, `01_locations.csv`, `02_cluster_stats.csv`) and are maintained in the canonical 00–21 sequence.

Supplementary Note S5: Residual-Lag Test of the Ridge-Recharge Attribution

S5.1 Purpose

The boundary-subsidy argument developed in §5.2.1 and §5.3 attributes the persistent water-balance residual at forest-margin wells — most visibly the +0.031 m/month residual at CEH14 (§4.9.6, Figure 45) — to lateral recharge derived from the northern rock ridge. This attribution rests primarily on two lines of evidence: the spatially structured pattern of positive residuals concentrated along the forest–dune boundary (Figure 45), and the water balance closure argument that no other plausible flux can account for the monthly deficit. Neither of these is a direct test of the transport mechanism itself. The residual-lag analysis described here provides an independent, falsifiable test: if ridge-derived recharge is delivered as a time-varying pulse, then water travelling from the ridge to successive dune wells must arrive at those wells with travel times that increase as a function of distance. The time structure of the SSM residuals, cross-correlated against rainfall, should therefore carry a distance-dependent lag signature. This Note documents the test and reports the null result.

S5.2 Extended State-Space Model

The single-period SSM fitted in Section 3.4.3 expresses the water table increment as a function of contemporaneous rainfall:

$$\Delta h(t) = \alpha + \beta_1 \cdot P(t) - \beta_2 \cdot PET(t) - \beta_3 \cdot h(t-1) + \varepsilon(t)$$

A preliminary analysis (Script 22) demonstrated that the residuals $\varepsilon(t)$ from this formulation carry a generic lag-1 rainfall signal at every well on the site, irrespective of cluster or position. This is the expected monthly-timestep consequence of the vadose zone: recharge takes approximately one month to propagate to the water table, so monthly-averaged water-table response lags monthly-averaged rainfall by roughly one month. This generic signal dominates the cross-correlation function at every well and masks any ridge-specific structure that might be present at longer lags.

To remove this confound, the residual-lag test uses an extended model that explicitly includes both $P(t)$ and $P(t-1)$ as regressors:

$$\Delta h(t) = \alpha + \beta_{10} \cdot P(t) + \beta_{11} \cdot P(t-1) - \beta_2 \cdot PET(t) - \beta_3 \cdot h(t-1) + \varepsilon'(t)$$

The β_{10} and β_{11} coefficients absorb the vadose-zone response between them; any remaining lag structure in $\varepsilon'(t)$ is therefore post-vadose and is the candidate ridge-transport signal. Fitted coefficients confirm that the extended parameterisation is well-identified: β_{11} is statistically significant ($p < 0.05$) at 100% of wells in every cluster, and the lagged-response fraction $\beta_{11} / (\beta_{10} + \beta_{11})$ increases monotonically from 0.55 in C1 to 0.83 in C4 — consistent with the slower drainage of forest soils. Mean R^2 improves from 0.55 (single-period) to 0.69

(extended). The extended model is used solely for this diagnostic test; the β_1 , β_2 , β_3 and α values reported in the main text remain the authoritative parameterisation for all other analyses.

S5.3 Cross-Correlation and Pre-Whitening

For each well with at least 140 months of data ($n = 57$ after excluding CEH3, CEH4, CEH7, CEH8 and CEH37 for reasons noted below), the Pearson cross-correlation $r(\varepsilon'(t), P(t - N))$ was computed for lags $N = 0$ to 12 months. Prior to cross-correlation, both the rainfall and residual series were pre-whitened using a first-order autoregressive filter: $x'(t) = x(t) - \phi \cdot x(t-1)$, where ϕ is the AR(1) coefficient of the rainfall series (+0.210 at RAF Valley). Wells whose residuals retained AR(1) structure above $|\phi| = 0.2$ were additionally pre-whitened against their own residual autocorrelation before the common filter was applied. Statistical significance was evaluated using the Bartlett 95% confidence interval for a white-noise series, $|r| \geq 1.96/\sqrt{N}$, which for typical series lengths gives a threshold of approximately 0.15.

Wells excluded from the test: CEH3 (tidal boundary — outside the SSM operational domain), CEH4 (coastal erosion drift plus post-2017 clearfell drawdown confounding any lag signal), CEH7, CEH8, CEH37 (standard upstream exclusions carried over from §4.9.4). Wells beyond the 140-month minimum were retained without further filtering.

S5.4 Hypothesis Test

The test metric is the Spearman rank correlation between each well's peak-correlation lag N^* and its Euclidean distance from the ridge reference point at (E = 241750, N = 364500, OSGB36) — a representative coordinate for the nearest point of the northern rock ridge to the dune field. Under the ridge-transport hypothesis, Spearman ρ should be significantly positive.

Of 57 wells fitted, 42 had a peak cross-correlation that exceeded the Bartlett significance threshold. Across these 42 wells, 33 showed their peak at lag 2 months and the remaining 9 distributed between lags 3 and 12 with no systematic geographic pattern. The Spearman rank correlation between peak lag and ridge distance was $\rho = -0.15$ ($p = 0.33$); a Mann-Whitney test specifically comparing C4 ridge-adjacent wells against C2+C3 wells confirmed that C4 peak lags (mean 2.78 months) were not longer than those of the geologically distant dune-body wells (mean 3.34 months; $U = 118.5$, $p = 0.74$). The hypothesis of distance-dependent ridge transport is not supported.

The magnitude of coupling tells a different story. C4 ridge-adjacent wells show substantially stronger rainfall-residual coupling at lag 2 (CEH34 $r = +0.40$; CEH13, CEH2, CEH14 all $r \approx +0.25$) than similarly close but geologically separate C3 wells (NW1 $r = +0.07$, not significant; NW10 $r = +0.15$). The strength of rainfall coupling is elevated at ridge-adjacent wells; its timing is not distance-structured.

S5.5 Interpretation

The null distance-lag result does not rule out ridge-derived recharge. It rules out one specific form of it: event-driven, pulse-delivered transport in which monthly rainfall on the ridge propagates to the dune field with travel times that scale with distance. The alternative interpretation — that ridge recharge is delivered as a near-steady baseflow, sufficiently smoothed by its transit through fractured bedrock that its month-to-month variance falls below the detection threshold of monthly water-level records — is consistent with the null. A

steady baseflow of this kind would manifest observationally as a constant positive contribution to the local water balance at wells along the flow path, which is precisely what the SSM residual α already absorbs. The two mechanisms are therefore not distinguishable from water-level data alone.

The discussion of the boundary subsidy in §5.2.1 has been adjusted to reflect this constraint: the ridge contribution is retained as a mechanism consistent with the data, but its form is characterised as steady-state rather than event-driven. Definitive resolution of the transport mechanism would require either a ridge-crest rain gauge paired with bedrock piezometers, or direct geochemical tracer work on the ridge–dune flow path. Both are identified as priority further work in §5.9.

S5.6 Reproducibility

The residual-lag test is implemented in two standalone analysis scripts, `22_residual_lag_analysis.py` and `23_ridge_recharge_lag_test.py`, which are maintained as supplementary diagnostics (Phase 11) alongside the main pipeline to preserve the reproducibility of the headline results from the canonical 00–21 sequence. Script 22 generates the residuals and AR(1) diagnostics; Script 23 fits the extended model, computes the cross-correlations, and applies the Spearman test. Both scripts read the same intermediate data produced by the main pipeline (`01_wells_clean.csv`, `01_climate.csv`, `01_locations.csv`, `02_cluster_stats.csv`) and produce independent outputs under `22_residual_lag_analysis/` and `23_ridge_recharge_lag_test/` in the project outputs tree. A plain-text summary of the hypothesis test result is written to `23_05_hypothesis_test_summary.txt` at runtime.

Supplementary Note S6: Seasonal Climatology Diagnostic of the Residual

S6.1 Purpose

Supplementary Note S5 reported a null result for the event-driven ridge-transport hypothesis. This Note addresses the other leading candidate mechanism for the water-balance residual: that the Thornthwaite PET estimate underestimates the true summer atmospheric demand on the water table because it is temperature-only and does not capture net radiation, vapour pressure deficit, wind, or surface-condition effects. Thornthwaite PET is an empirical index of atmospheric evaporative demand derived from temperature and day length; it does not define a specific reference surface. If that were the dominant explanation, the residuals should carry a systematic summer-negative signature that the fitted β_2 has not absorbed. This Note tests for that signature, along with two subsidiary diagnostics that distinguish between the remaining candidate explanations (steady ridge baseflow, nonlinear recharge response, and residual model error).

S6.2 Diagnostics

Three complementary tests are computed from the residuals $\varepsilon(t)$ produced by Script 22 for the same 57 wells that passed the ≥ 140 -month record-length filter and the site-extent exclusions detailed in S5.3.

(i) Seasonal climatology. For each well, the monthly climatology of the residual is computed as the mean of $\varepsilon(t)$ by calendar month over the record, giving a 12-point annual cycle. A sinusoidal fit $\varepsilon(m) = a_0 + a_1 \cdot \cos(2\pi m/12) + a_2 \cdot \sin(2\pi m/12)$ yields an amplitude $\sqrt{a_1^2 + a_2^2}$ and a phase (month of peak). Summer-minus-winter magnitude (JJA mean – DJF mean) is reported as a robust non-parametric proxy for the same quantity.

(ii) Independent ET proxy. A correlation of the residual against PET itself would be trivially zero because OLS fitting makes residuals orthogonal to every regressor by construction. Instead, the residual is correlated against monthly sunshine hours from the RAF Valley record — a direct radiation-based measurement that is not in the regression. If Thornthwaite underestimates summer ET in a way that β_2 has not absorbed, high-insolation months should carry disproportionately negative residuals, yielding a systematic negative Pearson r .

(iii) Spatial pattern within clusters. For C3 (Western-block open dune), the cluster is split by Euclidean distance from the ridge reference point used in S5 at a 1 km threshold (5 forest-adjacent wells, 13 warren-interior wells), and a Mann-Whitney test asks whether forest-adjacent wells have systematically smaller seasonal amplitudes than warren-interior wells — as would be expected if a steady ridge baseflow were flattening the residual at ridge-proximal locations.

S6.3 Result: the ET Hypothesis Is Not Supported

The sunshine-hours correlation is consistently negative in sign across every cluster, with network mean $r = -0.063$ and per-cluster means ranging from -0.023 (C4 Forest) to -0.079 (C3 West-Sand). None of the 57 wells exceeds the Bartlett 95% significance threshold of $|r| = 0.15$. The sign is directionally consistent with Thornthwaite slightly underestimating radiation-driven ET, but the magnitude of the bias is below the resolution of the monthly water-level data to detect.

Summer-minus-winter residuals are negative in all clusters except C1 (where they are within noise of zero), with magnitudes of order -0.014 m (C2) to -0.033 m (C4). These values are small relative to the month-to-month variability of the residuals themselves (typical well-level residual standard deviation of order 0.1 m) and should not be read as a strong summer-ET signal. Notably, C4 — the cluster where canopy interception would most plausibly produce Thornthwaite miscalibration — has the weakest sunshine correlation of any cluster, indicating that the cluster-mean β_2 has absorbed forest-specific ET behaviour reasonably well despite the simplicity of the fitted model.

The C3 within-cluster split also returns a null: forest-adjacent C3 wells have slightly larger seasonal amplitudes than warren-interior wells (0.021 m vs 0.019 m; Mann-Whitney $p = 0.88$, alternative “adjacent < interior”). This is the opposite direction from what a steady-ridge-baseflow mechanism would predict, though the difference is not statistically significant.

S6.4 What the Residuals Actually Look Like

The per-cluster seasonal climatologies reveal a structure that fits none of the candidate mechanisms cleanly: the annual cycle is bimodal, with positive residuals in January–February AND June–July, and negative troughs in April–May and September–October. The sinusoidal fit captures this poorly, which is why the fitted amplitudes appear modest relative to the visible annual range. The phase of the fitted sinusoid shifts systematically later in the year moving from C1 (mid-January) through C2, C3, to C4 (late March). No cluster has a phase in the JJA summer band.

This bimodal structure is not the signature of unmodelled summer ET, which would produce a single summer trough. It is not the signature of flat steady ridge baseflow, which would produce no annual cycle at all. And it is not a simple winter-recharge-nonlinearity signature, which would produce a single winter peak. Its most parsimonious interpretation is that a linear lumped-parameter model with time-invariant β coefficients is a rough approximation of a system whose soil-moisture storage operates nonlinearly across the annual cycle — wet-winter saturation increasing recharge efficiency beyond what the cluster-mean β_1 represents, shoulder-season soil drying reducing it below, and the whole cycle modulated by seasonal variation in vegetation water use that a constant β_2 cannot track.

The spatial pattern of seasonal amplitude nonetheless shows a coherent structure: the highest amplitudes (≥ 0.04 m) are concentrated in the north-western forest-margin cluster, with CEH14 as the network maximum (0.089 m), followed by CEH34, CEH13, and CEH2 — exactly the same spatial pattern as the water balance residual mapped in Figure 45 (§4.9.6). Whatever mechanism is producing the persistent residual is also producing the elevated seasonal variance. This is consistent with either (a) genuinely larger boundary fluxes at these wells that vary seasonally, or (b) the cluster-mean β values being a worse approximation at ridge-margin wells than at dune-interior wells. Water-level data alone cannot distinguish these.

S6.5 Combined Conclusion from S5 and S6

Taken together, Supplementary Notes S5 and S6 have excluded the two leading specific attributions of the water-balance residual:

The residual does not show the distance-dependent lag signature of event-driven ridge transport (S5).

The residual does not show the summer-negative signature of unmodelled Thornthwaite underestimation of atmospheric demand (S6).

The main text of §5.2.1 has been framed accordingly: the model explains 50–70% of month-to-month variance in Δh depending on cluster, and the unexplained portion is real and spatially structured (concentrated at the forest margin) but cannot be uniquely attributed to any single physical mechanism from water-level data alone. The fitted residual α remains the most defensible summary statistic for the steady-state boundary contribution; the underlying mechanism is most plausibly a combination of steady ridge baseflow, nonlinear soil-moisture storage dynamics not captured by the linear SSM, and minor β coefficient miscalibration, in proportions that cannot be resolved without additional measurement of the kind identified as priority further work in §5.9.

S6.6 Reproducibility

This analysis is implemented in `24_residual_seasonality.py`, which like Scripts 22 and 23 is maintained as a supplementary diagnostic (Phase 11) alongside the main pipeline to preserve the reproducibility of the headline results. The script reads the same intermediate data as Script 22 (`01_wells_clean.csv`, `01_climate.csv`, `01_locations.csv`, `02_cluster_stats.csv`) plus the raw RAF Valley climate file for sunshine hours (`RAF_Valley_Climate.csv`). All outputs are written to `24_residual_seasonality/` in the project outputs tree: a per-well climatology table (`24_residual_climatology.csv`), the per-cluster climatology panels (`24_01_climatology_panels_by_cluster.png`), the spatial amplitude map (`24_02_seasonal_amplitude_map.png`), the sunshine-correlation scatter (`24_03_sun_residual_correlation.png`), the per-cluster phase distribution (`24_04_phase_by_cluster.png`), and a plain-text summary of the diagnostic result (`24_05_diagnostic_summary.txt`).