Hydrogeological Dynamics, Behavioural Clustering and Management Intervention Analysis at Newborough Warren Coastal Sand Dune Aquifer, Wales

A state-space modelling and empirical analysis of groundwater behaviour, forest management, dune scraping, hydrological boundaries and critical rainfall forecasting

M. Hollingham | March 2026
ORCID: 0000-0003-0253-9301
Contact : martin.hollingham+nrg@gmail.com

Abstract

Coastal sand dune aquifers are ecologically sensitive systems whose groundwater regimes determine the persistence of legally protected dune slack habitats. Newborough Warren, a 1,300-hectare Special Area of Conservation (SAC) on the Isle of Anglesey, Wales, has undergone sustained water table decline over recent decades, threatening the humid dune slack communities for which the site is designated. This study presents a low-cost, reproducible, data-driven analytical pipeline of 27 steps across eleven phases, requiring only monthly dipwell readings and publicly available climate data from the nearest meteorological station. The pipeline is applied to a network of 89 wells monitored over 2005–2026, integrating hierarchical clustering, Pearson affinity analysis, a three-term state-space model (SSM) parametrizing recharge, atmospheric draw and head-dependent drainage, comparative model benchmarking, water table fluctuation specific yield estimation, dune scraping and clearfell BACI analysis with three-counterfactual ANCOVA climate correction, cluster-level water balance decomposition, per-well spatial coefficient mapping, critical rainfall threshold (P_flood) forecasting, seasonal prediction equations, and a per-well scenario framework for climate and forest management projections. The approach is designed for replication at comparable coastal dune sites using minimal equipment and no specialist modelling infrastructure.

Hierarchical clustering (Ward's linkage, k = 5) identifies five hydrogeological units spanning a behavioural spectrum from shallow pan to deep sponge to canopy sponge. The Eastern Block (C1 Lake Edge, C2 Dune) responds rapidly to rainfall with high recharge sensitivity and fast drainage. The Western Residual (C3) acts as a deep buffer (LCSC = 28.0%). Two Forest clusters (C4 Main Forest, C5 Coastal Forest) share elevated LCSC values (39.7%, 41.0%) reflecting canopy interception of approximately 24% of rainfall (Freeman, 2008), but differ in substrate setting: C4 overlies irregular bedrock topography on the elevated ridge flank (atmospheric draw coefficient $\beta_2 = 2.504$) while C5 occupies deeper, more uniform coastal sand ($\beta_2 = 1.374$). Per-well analysis within the forest zone shows that elevation alone explains 95% of β_2 variance, indicating substrate as a major contributor to the atmospheric draw coefficient, with depth-coupling and other factors also acting. WTF-derived specific yields converge across all clusters once corrected for interception, indicating that the

plantation's distinctive β-coefficient signature derives largely from canopy boundary conditions at the surface rather than from differences in subsurface storage architecture.

The SSM outperforms the traditional linear model in iterative forecasting (median NSE improvement +0.62; 65 of 66 wells with positive NSE vs 42 of 66 for TLM), demonstrating that explicit parameterisation of head-dependent drainage memory is essential for long-range prediction in this aquifer. The framework requires only monthly rainfall and temperature from a single climate station and is operationally transferable to other coastal dune aquifers. The water balance closes to within 2.5% at all five clusters under a displacement formulation, though the residual field is spatially structured with positive residuals along the northern forest margin whose mechanistic attribution is unresolved. Seasonal prediction equations establish that antecedent summer minimum depth, not winter rainfall, is the dominant predictor of winter flooding potential. The ecological significance is acute: the entire gradient from SD15b wet slack through SD16 dry slack to the SD16 recovery threshold spans only 37 cm of summer minimum depth (Curreli et al., 2013), so all management interventions must be evaluated on their summer minimum effect.

Dune scraping at CEH36 produced a paired summer minimum shift of +195 mm ($p = 0.004$) — more than half the SD15b–SD16 ecological gradient — confirmed by three independent methods (raw BACI +131 mm, synthetic control +135 mm, SSM forward residual +83 mm). The benefit persisted through the post-felling period while the paired control deteriorated, and the operational zone for this intervention is the C1/C2 area where P_flood thresholds remain achievable. CEH18 and CEH21, scraped in October 2023, showed limited responses that did not survive correction for background drift — consistent with progressive coastal boundary retreat overwhelming the intervention at the more coastal sites. No detectable drawdown was found at wells uphill of the scrape, though the observation window is too short to rule out longer-timescale propagation. Because scraping operates through the same physical mechanism as coastal erosion — lowering ground elevation to increase the drainage gradient — its spatial placement determines whether it complements or exacerbates the existing coastal drainage regime.

A five-tier, 17-well BACI experiment with three counterfactual controls tested the December 2017 clearfell of 8.4 ha of Corsican pine. Against the Forest control — the most direct test, comparing felled wells against unfelled forest sharing the same canopy and substrate — the ANCOVA-BACI detected a significant clearfell step of +93 mm at the Impact well ($p = 0.019$) and +153 mm at the Edge tier ($p < 0.001$), confirmed by synthetic extension of within-compartment FE wells (+103 mm, $p = 0.011$). The clearfell signal is concentrated in the mean monthly water level rather than in summer minima at the felled wells, where felled and unfelled wells tracked each other. The coefficient decomposition reveals a site-wide decline in recharge efficiency (β_1) across all five tiers of the BACI network, independent of management intervention and consistent with changing within-month rainfall distribution. This β_1 decline is the strongest candidate mechanism for the observed summer minimum deterioration and operates independently of canopy cover, affecting felled and unfelled wells equally.

The scenario framework evaluates clearfell, 50% thinning and broadleaf conversion alongside UKCP18 climate projections. Clearfell and thinning produce small positive net annual responses at the forest clusters (clearfell: C4 +4.7 mm w.e./month, C5 +8.8 mm; thinning: +2.3 and +4.4 mm respectively). Broadleaf conversion is near-neutral (C4 −3.4

mm; C5 +0.1 mm): the interception reduction from 24% to 15% during the recharge season is offset by the growing-season transpiration penalty. Forest management perturbations do not propagate measurably to the open dune clusters where ecological need is greatest. The forestry-scenario magnitudes are substantially smaller than the climate-driven changes (C4 dry-scenario −15.2 mm; wet-scenario +7.5 mm), indicating that climate trajectory dominates over canopy management at this site.

Network-wide summer minimum depths have declined at all five clusters, with C1 and C5 statistically significant, C2 marginal, and C3 and C4 non-significant. C5 Coastal Forest shows the steepest decline in both seasons and is the only cluster with a significant winter decline. A network-scale gradient regression of per-well summer-minimum slopes against perpendicular distance to the eroding western shoreline attributes approximately 42% of C5's exceptional decline to a coastal-retreat signal (fitted coast-edge anomaly −31 mm yr^{-1}, inland reach 869 m), with the remainder consistent with the two other candidate mechanisms operating in parallel: amplified head response under thin saturated thickness at the coastal margin, and progressive canopy maturation. C1, C2 and C4 lie beyond the fitted inland reach. The deterioration is concentrated in the summer season, consistent with the observed warming trend (post-2013 period +0.94°C above baseline) and increasing atmospheric evaporative demand. Trend line extrapolations indicate C1 summer minima approaching the SD16 viability threshold around 2030–2032.

The P_flood threshold equations, seasonal prediction equations and scenario framework are implemented as interactive tools hosted on the Newborough Warren Groundwater Research project page. The monitoring evidence characterises canopy removal as a winter-recharge intervention: it produces a detectable mean water table recovery at the Forest clusters, but its effect on summer minimum depth — the metric that determines dune slack viability — is set by the same site-wide β_1 decline that affects unfelled forest and open dune wells alike. The operational domain for direct hydrological intervention lies in targeted scraping in the C1/C2 zone where P_flood thresholds remain achievable, against a climate trajectory that is the binding constraint on all outcomes. The intervention window is measured in one to two decades.

Table of Contents

1. Introduction

Sand dune systems are among the most ecologically dynamic and internationally important coastal habitats in northwest Europe (Dargie, 1995). Their characteristic dune slack communities — low-lying inter-dune hollows that flood seasonally where the water table intersects the surface — support a concentration of nationally rare plant species including creeping willow (*Salix repens*), and a diverse suite of wetland-associated invertebrates. The persistence of these communities is fundamentally dependent on a shallow, responsive groundwater table that rises into the slack floor each winter.

Newborough Warren on the Isle of Anglesey (National Grid Reference SH 406 636) is one of the largest and most ecologically important coastal dune systems in Wales. Designated as a Special Area of Conservation primarily for its humid dune slack and creeping willow communities, the site comprises approximately 1,300 hectares of late-glacial blown sand resting on weakly permeable glacial till (Stratford et al., 2007). The northern 700 hectares were afforested with Corsican pine (*Pinus nigra* var. *laricio*) between 1948 and 1965.

A well-documented trend of water table decline has been observed at the site over the past six decades (Ranwell, 1959; Jennings, 1990). Conservation management since 2018 has included an experimental clearfell of approximately 8 ha within the plantation, and topographical dune scraping at selected slack sites. These interventions provide natural experiments by which the hydrological consequences of management can be assessed empirically. Despite the importance of the site and the scale of management investment, no systematic network-wide quantitative analysis has previously been undertaken.

Restoration of dune slack hydrology is widely assumed to be achievable through topographical intervention provided sufficient winter rainfall occurs. However, if summer water table minima are declining at a rate that progressively increases the recharge deficit that must be overcome each winter, there exists an intervention window — a period during which topographical management can still realistically restore flooding conditions — beyond which no realistic winter rainfall total can compensate for the accumulated summer deficit. Quantifying the rate of closure of this window, and identifying the site-specific mechanisms that drive it, is therefore an urgent practical and scientific priority. This requires not only a long baseline record of groundwater behaviour but a mechanistic understanding of how recharge, evapotranspiration and drainage interact across the hydrogeologically distinct sub-units of the dune system.

This study presents such an analysis. The specific objectives are to: (1) classify the dipwell network into hydrogeologically coherent groups by hierarchical clustering and extend membership to the full monitoring network via Pearson affinity analysis; (2) characterize mechanistic groundwater signatures using a state-space model benchmarked against a traditional regression approach; (3) map and interpret unmeasured boundary subsidies; (4) quantify the effects of dune scraping using a two-tier hierarchical CUSUM-BACI framework that separates localized treatment signals from background system-wide deterioration; (5) conduct a three-counterfactual ANCOVA-BACI assessment of the December 2017 plantation clearance across a 17-well five-tier experimental network, partitioning intervention signals from climate variability, scraping confounds and coastal erosion gradients; and (6) derive operational critical rainfall threshold equations for conservation planning.

1.1 The Newborough Warren Analytical Architecture

To achieve a rigorous characterisation of the aquifer's hydrogeological dynamics, this study employs a nine-tiered analytical framework that scales from individual well diagnostics to network-wide predictive modelling. A further objective is to demonstrate that the analytical framework is replicable at comparable dune slack systems using only manual dipwell records and publicly available climate station data, without specialist survey equipment, telemetry or laboratory analysis.

- **Hydrogeological Clustering:** The 66-well reference network is partitioned into hydrogeologically distinct zones using agglomerative hierarchical clustering with Ward's minimum variance linkage (Section 3.2). The Pearson correlation distance isolates temporal hydrograph dynamics from absolute water table elevation, and the optimal cluster number ($k = 5$) is selected using a dual-metric approach combining Ward's merge distances and mean silhouette coefficient. The resulting five clusters provide the spatial framework for all subsequent analysis.
- **Extended Network Classification and Pattern Matching:** The primary hydrogeological classification is extended from the 66-well reference network to the full 89-well monitoring array using Pearson Affinity Analysis (Section 3.3). Unlike the mechanistic models, this phase utilizes pattern-matching on raw observed time series without model fitting. It identifies core cluster members and pinpoints Multi-Cluster Affinity (MCA) flags at structural transition zones where local topographic drivers are superseded by system-wide climatic forcing. The wells in the extended network are not used in the rest of the analysis.
- **Mechanistic Modelling and Storage Characterization:** At the cluster scale ($n=5$), the State-Space Model (SSM) is used to derive physical coefficients (β_1, β_2, β_3) directly from cluster-average time series (Section 3.4.1). These coefficients characterize the fundamental contrast between the "shallow pan" and "deep sponge" blocks, providing the basis for both the Lumped Catchment Storage Coefficients (LCSC; Section 3.4.1) and the operational P_flood threshold equations (Section 3.6.3) used for management forecasting. Per-well coefficients are interpolated across the site (Section 3.4.3), providing a spatially continuous expression of recharge sensitivity, atmospheric draw, and drainage rate that independently confirms the cluster boundaries and delineates the interception imprint of the plantation canopy.
- **Predictive Validation and Model Benchmarking:** The predictive skill of the SSM is formally validated through matched comparison against a Traditional Linear Model (TLM) at every individual well ($n=66$) in the reference network (Section 3.4.2). This benchmarking uses iterative Nash-Sutcliffe Efficiency (NSE) and one-step R^2 to assess predictive skill across the network. Spatial mapping of the performance gains (ΔR^2, ΔNSE) demonstrates where the explicit parameterisation of storage decay is indispensable for accurate multi-step forecasting.
- **Intervention Analyses:** The hydrological consequences of two management interventions are quantified using complementary BACI frameworks (Sections 3.5.4–3.5.5):

  - *Dune Scraping:* A two-tier hierarchical CUSUM-BACI framework across five focal wells separates localized scraping signals from background system-wide deterioration. Tier 1 validates local control wells against a regional mean; Tier 2 differences each treatment well against its matched control. Era-specific storage decay coefficients ($-\beta_3$) are estimated independently for each management period to assess whether scraping has altered subsurface drainage mechanics.
  - *Plantation Clearfell:* A three-counterfactual ANCOVA-BACI experiment across 17 wells in a five-tier hierarchical design — Impact, Edge, Forest Control, Coastal Control and Climate Control — assesses the hydrological consequences of the December 2017 clearfell. The ANCOVA model includes a cumulative water balance covariate, a distance-weighted scraping covariate, and an easting × time coastal erosion interaction. The model is run against three control centroids (Forest, Climate, Combined) to bracket the clearfell effect.

- **Robustness and Confound Assessment:** Independent robustness analyses are conducted for both intervention analyses to test whether the detected signals are genuine rather than artefacts of control well limitations or climate confounds (Sections 3.5.4–3.5.5). For the scraping analysis, the reliability of CEH4 as a stable long-term control is tested using two independent methods: a synthetic control constructed as an optimally-weighted combination of 11 donor wells calibrated on the pre-scraping period, and an SSM forward residual computed by calibrating the model on the pre-scraping period and running it forward iteratively. Convergence between the raw BACI estimate, the synthetic control gap and the SSM residual confirms whether the scraping signal at CEH36 is detectable independently of CEH4's own coastal erosion and felling pulse trajectory; an additional propagation analysis tests for distance-decay drawdown at uphill wells. For the clearfell analysis, the primary three-counterfactual ANCOVA design (Forest, Climate and Combined controls) is itself a robustness structure, with the Forest control providing the most direct test against position-matched unfelled wells. Four further supplementary analyses test the ANCOVA result: summer minima analysis with mixed-effects modelling; coefficient decomposition comparing Before and After SSM fits; SSM residual normalisation against the Forest Control drift; and a synthetic control constructed from optimally weighted donor wells. A spatial transect analysis tests whether post-felling step changes show a distance-dependent gradient consistent with a localized clearfell effect.

- **Climate Trajectory and Threshold Exceedance Analysis:** Annual summer minima and winter maxima are extracted for C1–C3 and projected to 2040 against the eco-hydrological viability thresholds established by Curreli et al. (2013), quantifying the rate of closure of the ecological intervention window and the proportion of monitored years in which flooding thresholds are exceeded across the principal structural domains (Section 3.6.4).

- **Water Balance Decomposition:** Cluster-level SSM coefficients are used to decompose the mean monthly water balance into its component fluxes — recharge, atmospheric draw, gravity drainage — expressed both as head-space quantities (m/month) and as indicative volumetric fluxes (mm/year) using assumed specific yield (Section 3.7.2). This tier quantifies the structural water balance at each cluster, independently of land cover or management history.

- **Water Table Fluctuation Specific Yield Estimation:** Empirical upper bounds on specific yield are derived independently from the dipwell record using the water table fluctuation method (Healy and Cook, 2002), applying an interception correction for the Forest cluster following Freeman (2008). The resulting cluster-specific Sy estimates (Section 3.7.3) bracket the volumetric water balance from above, confirming that the assumed values used in the primary decomposition are conservative lower bounds, and confirm broadly uniform aquifer storage properties across the site irrespective of surface land cover.

2. Study Site

Newborough Warren occupies the south-western tip of the Isle of Anglesey in north-west Wales (Figure 1). The dune field is underlain by weakly permeable glacial till forming the effective base of the sand aquifer. A bedrock ridge runs south-west to north-east and forms the northern boundary of the study area, its western end entering the sea. A link to an interactive map of the long term monitoring network is in the supplementary materials section. The hydrological study area, defined by the extent of the DEM-derived flow network (Figure 1: modelled drainage lines are derived from topographic analysis of the digital elevation model, and represent theoretical flow paths rather than mapped surface watercourses), covers approximately 1,172 ha (~4.0 km × 3.0 km). The 89-dipwell monitoring network spans a ~720 ha core zone (~3.3 km × 2.2 km) within this, lying entirely on the south-eastern flank of the ridge. The dune system on the north-western flank of the ridge, draining toward the Malltraeth estuary, is hydrologically separate from the main Newborough Warren aquifer; monitoring wells located on that flank do not have sufficient record length for inclusion in the clustering and state-space analyses reported here.

Lateral recharge descending from the ridge into the dune sands contributes to the water balance residuals detected in the state-space model. These residuals are most pronounced at the forest wells along the northern margin of the study area, which lie closest to the ridge; lake-edge wells (C1) show no residual contribution, consistent with their water balance being dominated by direct exchange with Llyn Rhos-Ddu rather than ridge-derived recharge. The dune aquifer south-east of the ridge is divided into an eastern block, where shallower sand over till and estuarine deposits produces a more responsive, heterogeneous water table regime, and a western block, where deep sand accumulation over a broader platform permits significant groundwater storage and more buffered seasonal fluctuations (Stratford et al., 2007).

The aquifer discharges principally across the south-western foreshore (Caernarfon Bay) and the southern and eastern coast (Menai Strait) as tides recede. A small lake within the warren, Llyn Rhos-Ddu, acts as a near-fixed head body and is the dominant local drainage sink for the immediately adjacent C1 Lake Edge wells; it receives groundwater throughout the year as the regional water table maintains its head above the lake surface, and supplements aquifer recharge during the driest months (Stratford et al., 2007). The full monitoring network at the site comprises 117 dipwells which have been measured at various times (Figure 1), of which 66 have sufficient record length for inclusion in the core analyses.

The northern 700 hectares are managed plantation dominated by mature Corsican pine. A partial clearance of approximately 8.4 ha (just over 1% of the plantation) began in December 2017, removing trees across a defined felling zone at the plantation–open dune transition (Figure 2). Dune scraping has been undertaken at selected western slack locations: CEH36 in April 2015, and CEH18 and CEH21 in October 2023.

The dune slack communities at Newborough Warren are classified under the National Vegetation Classification as SD15b (wet slack, *Salix repens–Calliergon cuspidatum* community) and SD16 (dry slack, *Salix repens–Holcus lanatus* community). Curreli et al. (2013) (see also Davy et al., 2006) established the water table depth requirements for these communities from a site-specific vegetation survey across Newborough Warren and four other Welsh dune systems: SD15b viability requires the mean annual summer minimum water table to remain shallower than 0.61 m below the ground surface and the winter maximum to reach within 0.10 m of the surface; SD16 communities tolerate summer minima to 0.98 m and require winter maxima within 0.25 m of the surface. These thresholds provide the ecological reference frame for all hydrological assessments in this study. Wells whose mean summer minimum exceeds 0.98 m are classified as ecologically critical — beyond the range of SD16 viability. For the purposes of assessing management options, two intermediate bands are defined: wells between 0.61 m and 0.75 m below ground are likely to be recoverable to SD15b viability by a standard dune scraping event of approximately 0.14 m depth (consistent with the BACI-measured benefit at CEH36; Section 4.5); wells between 0.98 m and 1.20 m are likely to be recoverable to SD16 dry slack viability by a comparable intervention. These 'recoverable' classifications should be treated as indicative — achievable excavation depth varies with substrate, access and licence conditions, and the empirical basis is a single scraping event at one well. Wells deeper than 1.20 m are beyond the reach of a standard single scraping event under average antecedent conditions."

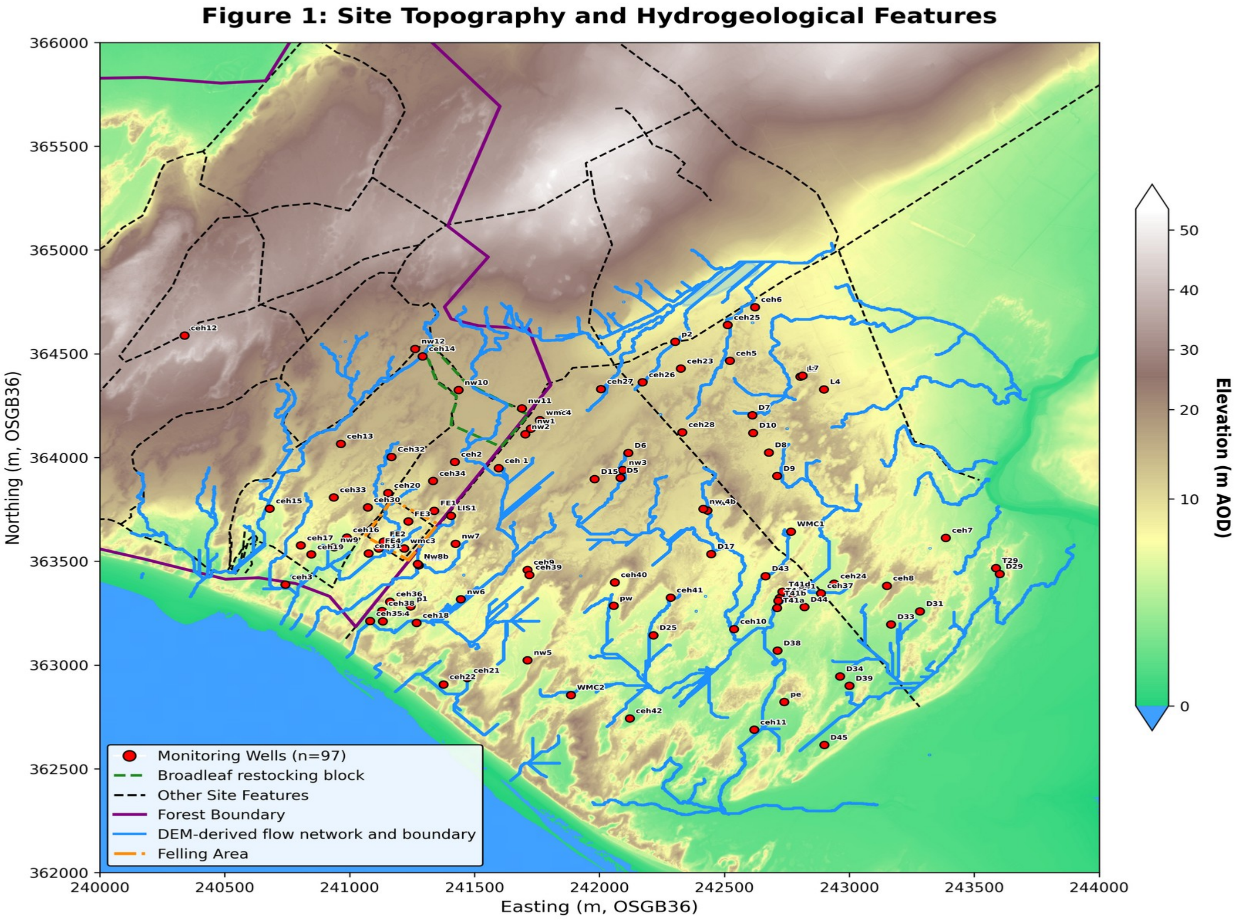

Figure 1. Site topography and hydrogeological features. Digital elevation model of Newborough Warren showing the full 117-well monitoring network, rock ridge, forest boundary, DEM-derived flow network (modelled drainage lines from topographic convergence analysis; not mapped surface watercourses) and watershed boundary. (12_01_dem_site_overview.png) (LiDAR data: © NRW & OS; see Section 10 for full attribution).

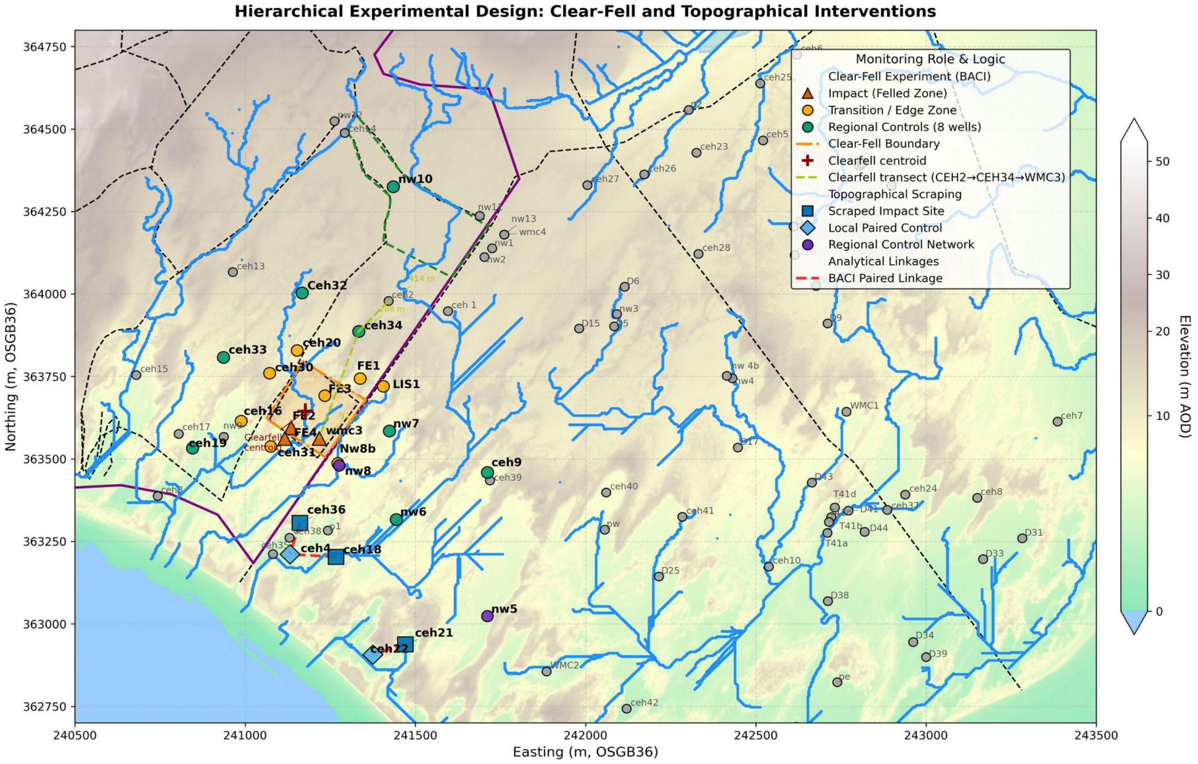

__Figure 2.__ Hierarchical experimental design. Spatial layout of the monitoring network showing the five-tier BACI structure: Impact (WMC3), Edge (CEH31, CEH20, CEH30, CEH16), Forest Control (CEH32, CEH34, CEH33, NW10, CEH2), Coastal Control (CEH19, CEH17), Climate Control (CEH9, NW7, NW6, NW5, WMC2), scraped sites (CEH36, CEH18, CEH21/22), and felling boundary. (LiDAR data: © NRW & OS; see Section 10 for full attribution).

3. Methods

3.1 Data Sources and Preparation

All analyses were implemented as a sequential, numbered Python pipeline of 27 scripts, of which 23 constitute the core analytical workflow and 3 provide supplementary diagnostics. Each script reads from and writes to a shared outputs directory, and scripts are numbered to reflect execution order (00 through 24, including script 11b). The pipeline is documented in the accompanying README and all outputs are reproducible by running the published scripts in sequence. The companion browser-based modelling suite, which operationalises the P_flood forecasting equations and scenario tools derived from this pipeline, is also hosted at https://newbroman.github.io/Newborough_Hydrology/. Where script names are cited in the text below, they refer to files in the pipeline repository.

3.1.1 Groundwater Records

Groundwater level data were obtained from a network of 89 dipwells. Field measurements record depth to water below the top of each standpipe as a positive number in metres; water standing above the standpipe top is recorded as a negative value. Records span April 2005 to February 2026. Raw records were quality-controlled by removing values exceeding 4.0 m below standpipe top as physically implausible for this dune system, and assigned to the appropriate calendar month by field convention: readings taken on or before the 15th of a month are assigned to the previous month (reflecting the field practice of measuring at the start of the month to characterize the preceding month's conditions), while readings on the 16th or later are assigned to the current month. In general most readings were taken one or two days of the month’s end. Where multiple readings fall within the same month after assignment, they are averaged to produce a single monthly mean which applies to a few readings in 2005-06. Gaps of up to three consecutive months were filled by time-based linear interpolation; longer gaps were retained as missing values. All monthly data are indexed by year and month (e.g. July 2007); where date stamps appear in tabular outputs with a day suffix (e.g. 2007-07-01), the day component is a formatting artefact and does not denote a specific day — the record represents the complete calendar month. Monthly rainfall and Thornthwaite PET totals from RAF Valley were paired to the same calendar month as the water level they accompany. The monthly change in water level (Δh) is thus explained by the climate forcing that acted during that month, with drainage driven by the water table position at the end of the preceding month.

A reference network of 66 wells, each with more than 100 months of valid record extending to February 2026, formed the basis of the clustering and state-space modelling analyses. Eight wells meeting the record-length criterion were excluded from the reference network prior to analysis. Wells within the clearfell management footprint (FE1, FE2, FE3, FE4 and LIS1; Section 4.6) were held out on the grounds that removal of canopy interception following the 2017 felling might have introduced a non-stationary upward drift in their water tables that the time-invariant state-space formulation is not designed to represent. These wells were subsequently excluded from the clearfell BACI analysis (Section 4.6) on the grounds of insufficient pre-felling baseline (Section 3.5.5) and contribute to the extended-network Pearson affinity classification (Section 4.3). One lake-surface well (Llyn Rhos-ddu) was excluded because it reads a near-fixed head body rather than a water table and therefore violates the storage-decay assumptions of the state-space model. Two further wells (CEH3, CEH22) were excluded on clustering grounds described in Section 3.2.3 as they showed drift attributable to coastal erosion. Seven of the eight excluded wells remain in the extended network for per-well analysis. Llyn Rhos-ddu is excluded from both networks — its lake-stage signal is not interpretable as a groundwater response. CEH36, which received scraping management at a single point during the monitoring record (Section 4.5), was retained as a reference-network member on the grounds that scraping produced a one-time baseline shift rather than ongoing trend, and its post-intervention response remains consistent with its cluster mates. Wells not meeting the reference-network criteria for reasons of insufficient record length were assigned to the extended network used in the Pearson affinity analysis to extend the hydrogeological classification beyond the reference network boundary. Dipwells are positioned across a range of geomorphological contexts including dune slack floors, slack margins, and inter-slack hollows; most but not all wells are sited within slacks. Notable exceptions include L1, which is positioned adjacent to a fenced grazing experiment on slightly elevated ground, and NW1, where the immediate surroundings have since been excavated.

Ground surface elevations for all dipwells were derived from the Natural Resources Wales LiDAR Digital Terrain Model (2 m resolution, captured March 2023; Natural Resources Wales, 2023), obtained from the DataMapWales repository (https://datamap.gov.wales/maps/lidar-viewer/). The accuracy of the DEM-derived elevations was assessed by comparison against independently surveyed DGPS ground elevations at 61 wells ($R^2 = 0.976$, RMSE = 0.477 m, mean bias = +0.264 m). The bulk of this RMSE is attributable to six outlier wells where discrepancies exceeded 0.5 m; five are located within the Corsican pine plantation, where satellite signal attenuation under dense evergreen canopy is a known source of DGPS positional error, and the sixth is an open-ground well whose DGPS-recorded elevation matches that of a plantation well 245 m distant, strongly suggesting a recording error. Excluding these six wells, the RMSE falls to 0.231 m with a mean bias of +0.203 m, consistent with the NRW stated DTM accuracy of ±0.15 m RMSE on open ground plus a modest systematic offset attributable to residual vegetation returns. The DEM is therefore adopted as the primary elevation source throughout this study.

Water table elevation in metres above Ordnance Datum (maOD) was derived using surveyed pipe-top elevations (Pipe_Top_Elev) and field-measured upstand heights (u, the height of the pipe top above ground surface), as maOD = Pipe_Top_Elev + raw_reading, where raw readings follow the negative depth-below-pipe-top convention. Upstand heights ranged from −0.14 m (CEH28, a sunken casing) to +0.71 m (CEH2). The maOD head surfaces are used in the spatial groundwater analysis (Section 3.8) and the broadleaf comparison (Section 4.6.8); the temporal clustering, SSM parameterisation, and BACI analyses use depth-below-pipe-top records throughout for consistency with the original monitoring convention.

For the cluster average hydrograph construction (Section 3.4.1), each well's depth-to-water series was corrected to a common ground-surface datum by subtracting its surveyed upstand height before averaging across cluster members, ensuring that between-well differences in pipe-top elevation do not bias the cluster mean. This correction was not applied to individual well SSM fits, where the pipe-top offset cancels in the first-difference operation, nor to the clustering and Pearson affinity analyses, where z-score standardisation renders absolute depth irrelevant. Network summary statistics and well characterisation figures are presented in Section 4.1.2 and Figure 5.

3.1.2 Climate Data

Monthly total precipitation and potential evapotranspiration were obtained from the RAF Valley Met Office climate station (SH 30691 75549), approximately 16 km north of the site. The station record extends from December 1930 to February 2026, with a mean annual rainfall of 890 mm over the well-record period (calculated from complete calendar years 2006–2025; the partial years at the start and end of the monitoring record are excluded from this summary statistic). A single missing monthly rainfall value in the historic record (June 1941) was retained as missing rather than substituted; it falls outside the well-record period and has no influence on any analysis.

Potential evapotranspiration was calculated using the Thornthwaite and Mather (1957) method with a day-length and month-length correction, applied to mean monthly air temperature derived as the mean of daily maximum and minimum. The method proceeds in three steps. First, the annual heat index I is computed as the sum of twelve monthly contributions:

$$I=\sum_{m=1}^{12}\left(\frac{T_m}{5}\right)^{1.514}$$

where T_m is mean monthly temperature clipped to zero for months at or below 0°C. Second, unadjusted monthly PET is calculated as:

$$PET_{unadj}=16\times\left(\frac{10\,T}{I}\right)^{alpha}$$

where: $alpha=6.75\times10^{-7}I^3-7.71\times10^{-5}I^2+1.792\times10^{-2}I+0.49239$

For months where mean temperature exceeds 26.5°C a high-temperature correction is applied $PET_{unadj}=-415.85+32.24\,T-0.43\,T^2$; this threshold is not reached at RAF Valley under the historic record but is retained for formal completeness. Months with mean temperature at or below zero are assigned PET = 0. Third, the unadjusted value is scaled by a day-length and month-length correction factor:

$$PET=PET_{unadj}\times K\text{, where }K=\left(\frac{N}{12}\right)\times\left(\frac{d}{30}\right)$$

where N is the maximum possible sunshine duration for the month (hours), calculated from site latitude (53.25°N, RAF Valley) using standard solar geometry, and d is the number of days in the month. This correction accounts for the fact that the original Thornthwaite equation was calibrated for a standard 30-day month and 12-hour photoperiod.

Thornthwaite PET is known to overestimate evaporative demand in humid maritime climates relative to physically-based methods such as Penman-Monteith (Donohue et al., 2010), and it was selected here because it requires only mean air temperature, consistent with the low-cost replication objective described in Section 3.9. Although forest recharge at Newborough is estimated to be 100–200 mm yr^{-1} lower than open dune recharge (Stratford, 2006), these absolute differences do not affect the comparative conclusions of this study. Because the same PET series is applied uniformly across all clusters and time periods, systematic biases cancel in cluster-to-cluster and before-versus-after contrasts.

A centred cumulative water balance series (CWB; alternatively termed the centred cumulative water balance anomaly elsewhere in this report) was computed for the well-record period by summing monthly net balance (P − PET), with a mean-centring correction applied to remove the systematic long-term offset between precipitation and Thornthwaite PET inherent to this maritime climate; the corrected series oscillates around zero and directly reflects periods of cumulative surplus and deficit relative to the long-term mean. CWB is referenced as a covariate in the BACI ANCOVA models in Section 3.5.5 and in the climate sensitivity analysis in Section 4.6.3. The long-run detrending mean was 19.46 mm $month^{-1}$. The correspondence between this cumulative balance series and the unsmoothed network mean water level was assessed by linear regression at zero lag ($R^2 = 0.515$), confirming that interannual water table variability is substantially explained by antecedent climate forcing and providing the empirical motivation for the state-space modelling framework described in Section 3.4.

Baseline climate figures and well network characterisation statistics (record length, mean water level, seasonal amplitude, and monthly box plots) were produced by `00_climate_summary.py`*, with data preparation, quality control, gap-filling and upstand correction implemented in* `01_data_prep.py`*.*

3.2 Hierarchical Clustering

To characterize the spatial structure of the aquifer and identify distinct zones of hydrological behaviour, agglomerative hierarchical clustering was applied to the continuous dipwell monitoring network.

3.2.1 Data Pre-processing and the Pearson Correlation Distance

Because the absolute elevation of the water table varies significantly across the site due to topography, analyzing raw head data would artificially separate functionally identical slacks. To isolate the temporal dynamics—the "shape" of the hydrograph—each well's monthly time series was used directly in the correlation matrix; the correlation-based distance metric inherently removes mean and variance differences between series, making explicit z-score standardisation unnecessary. (Note that explicit z-score standardisation is applied at the cluster average construction stage of the subsequent Pearson affinity analysis, where it ensures that wells with different absolute ranges contribute equally to the cluster average time series.)

Pairwise dissimilarity between the standardized time series was then computed using the Pearson correlation-based distance $d_{i,j}$:

$$d_{i,j} = 1 - r_{i,j}$$

where $r_{i,j}$ represents the Pearson correlation coefficient calculated over contemporaneous, overlapping temporal records for wells i and j. This specific metric translates the statistical correlation directly into the physical behaviour of the aquifer:

- **Synchronized Slacks (r=1):** Where two dipwells rise and fall in perfect harmony, the distance metric becomes zero (d=0). The algorithm recognizes identical hydrogeological drivers and groups them tightly.
- **Unrelated Slacks (r≈0):** Where one well responds to flashy winter storms and another is heavily buffered, the correlation approaches zero, yielding a distance of one (d=1), effectively pushing them apart in the clustering space.
- **Opposing Slacks (r=−1):** Where wells exhibit inverse temporal behaviour, the distance maximizes at two (d=2).

By utilizing the Pearson correlation distance, the absolute topographic elevation of the sand dunes is stripped from the analysis. The algorithm is forced to group wells purely on the timing and shape of their temporal hydrographs, allowing deep aquifer wells to cluster together even if they are situated at vastly different physical elevations.

3.2.2 Agglomeration Strategy

The dissimilarity matrix was processed using agglomerative hierarchical clustering, employing Ward’s minimum variance linkage (Ward, 1963). At the initiation of the algorithm, each dipwell is treated as an individual cluster. The algorithm then iteratively merges the two clusters that result in the smallest possible increase in the total within-cluster variance, quantified mathematically as the Error Sum of Squares (ESS).

For any given cluster C containing n wells, the ESS is defined as the sum of the squared distances between each well's standardized time series (X_i) and the mean temporal signature (centroid, $\bar{x}_C$) of that cluster:

$$ESS_C = \sum_{i \in C} \left\| x_i - \bar{x}_C \right\|^2$$

When evaluating the potential merge of two clusters, A and B, the algorithm calculates the resulting increase in total variance (*ΔESS)*. This merge cost is computed algebraically as:

$$\Delta ESS(A, B) = ESS_{A \cup B} - \left(ESS_A + ESS_B \right)$$

Which simplifies to the objective function:

where n_A and n_B represent the number of wells in clusters A and B, respectively.

This specific algebraic formulation is hydrogeologically essential for mapping a continuous dune aquifer. Because the spatial groundwater gradient transitions smoothly from the bedrock ridge to the foreshore, single-linkage (nearest-neighbour) algorithms frequently suffer from "chaining"—sequentially merging adjacent wells until the entire site collapses into a single, uninterpretable cluster.

Ward's mathematical objective function actively resists this chaining effect. The $\frac{n_A n_B}{n_A + n_B}$term acts as a strict mathematical penalty against merging disproportionately large clusters or appending single outliers to established groups. Instead, it forces the algorithm to prioritize merging smaller, tightly correlated groups first, ensuring that the hydrographic signatures within any newly formed group remain highly synchronized. By penalizing variance, Ward's linkage imposes discrete analytical boundaries onto a continuous physical system, delineating the continuous sand matrix into compact, functionally distinct hydrogeological catchment zones (e.g., the rapid-draining Western Block and the near-perennial Eastern Block).

3.2.3 Cluster Determination and Robustness Assessment

The optimal number of clusters (k) was evaluated iteratively for k = 2 to k = 10 using three complementary metrics: the elbow method applied to Ward's merge distances, the mean silhouette coefficient (Yuan and Yang, 2019), and the Calinski–Harabasz index (Caliński and Harabasz, 1974). Silhouette and Calinski–Harabasz both peaked at k = 2, reflecting a fundamental binary split between the shallow western and deeper eastern aquifer responses. However, selecting k = 2 provides insufficient resolution to distinguish the geomorphological units relevant to conservation management. Rather than selecting k on a single metric, the partition was evaluated across a range of candidate values using bootstrap stability analysis (Hennig, 2007) to quantify how sensitive cluster assignments are to the particular set of wells observed.

For each candidate k in {4, 5, 6, 7}, the 66-well reference network was resampled with replacement 1000 times. At each resample, Ward's linkage was re-applied and the resulting assignments compared with the reference (full-sample) partition. For every pair of wells, the fraction of resamples in which both wells were assigned to the same cluster was recorded as a pairwise co-assignment probability; the median of these probabilities across a well's reference-cluster partners provides a per-well stability score. A well scoring 1.0 co-assigns with the same neighbours in every resample; lower scores indicate sensitivity to the composition of the network.

At k = 5, four of the five clusters exhibited median per-well stability exceeding 0.93: the Main Forest cluster (stability 1.00), the Coastal Forest cluster (0.99), the Lake Edge cluster (0.98), and the Dune cluster (0.93). The fifth cluster, Western Residual, exhibited moderate stability (0.50), reflecting the landscape heterogeneity of this group rather than a failure of the partition — at higher k this group fragments into sub-clusters with individually poor stability, confirming that k = 5 represents the finest resolution at which Ward's linkage produces behaviourally coherent groups. The final solution of k = 5 was therefore selected as the finest partition at which the majority of clusters are demonstrably robust under resampling, while providing sufficient resolution to distinguish the principal geomorphological units.

The silhouette coefficient at k = 5 is 0.39. This moderate value is consistent with the physical reality of a contiguous, highly permeable sand aquifer where cluster boundaries are inherently gradational. The bootstrap stability analysis provides a more informative assessment of partition quality than the silhouette score alone, because it directly quantifies whether individual well assignments are reproducible rather than measuring only average cluster separation.

Two wells — CEH3 and CEH22 — were excluded from the reference network prior to clustering because Ward's linkage consistently identified them as singleton outliers at every k from 4 to 9, resistant to grouping with any behavioural cluster. Both are low-elevation coastal wells (ground elevation 3.3–4.5 m) whose correlation structure is consistent with tidal-signal contamination superimposed on the climate-driven response that the clustering and subsequent state-space modelling are designed to characterize. Their exclusion follows the same principle applied to the FE/LIS clearfell wells and Llyn Rhos-ddu (Section 3.1.1): wells whose behaviour violates the stationarity or single-mechanism assumptions of the analytical framework are retained in the extended network for per-well analysis but excluded from the reference partition. Including either well distorts the Ward's tree at lower k values — CEH3 suppresses the Lake–Dune separation at k = 4, and CEH22 appears as a persistent singleton from k = 5 onward — so their exclusion is both physically motivated and algorithmically necessary.

A converse case is CEH11, which is assigned to C1 (Lake Edge) on the basis of hydrograph similarity but is geographically isolated from the other six C1 members. The remaining C1 wells occupy ground at 8.6–10.9 m AOD around Llyn Rhos-ddu in the northern part of the site; CEH11 sits at 3.6 m AOD on the south-eastern coast near the Menai Strait, approximately 1.6 km from the nearest C1 well (CEH27). Its damped, low-amplitude seasonal hydrograph — the signature that drives its C1 assignment — likely reflects proximity to the Menai Strait tidal boundary rather than to Llyn Rhos-ddu: a different constant-head source producing a similar buffered water-table response. The cluster assignment is retained because the statistical grouping is hydrograph-based and CEH11's behaviour is genuinely C1-like, but the different boundary condition means its mechanistic coefficients should not be interpreted alongside the lake-edge wells without qualification. This distinction is relevant to the β_2–elevation analysis in Section 4.3, where CEH11's inclusion extends the C1 elevation range from 8.6–10.9 m down to 3.6 m without contributing to the lake-proximity gradient that governs the other members.

3.2.4 Seasonal Amplitude as a Secondary Cluster Descriptor

The Pearson correlation distance described in Section 3.2.1 is scale-invariant: because Pearson correlation is unchanged by linear rescaling, two wells whose hydrographs rise and fall in synchrony will cluster together regardless of whether one exhibits a seasonal range of 0.5 m and the other 2.0 m. This property is desirable for identifying shared drainage mechanisms, but it means that wells within a single cluster may differ substantially in response magnitude — reflecting local differences in specific yield, proximity to hydrogeological boundaries, or depth to an impermeable layer. To characterize each cluster on this second, orthogonal axis, a suite of seasonal amplitude descriptors was computed as a supplementary characterisation alongside the pattern-coherence statistics reported in Section 4.2.

The primary amplitude metric is the difference between the 90th and 10th percentiles of each well's monthly water-level series (p90 − p10), expressed in metres. This inter-decile range captures a representative seasonal swing while remaining robust to single-month outliers. It was computed for each well over three windows: the full record, the pre-2018 period (prior to the December 2017 clearfell event in the FE compartment), and the post-2018 period. Wells with fewer than 24 monthly observations in a given window were excluded from that window's calculation. Cluster-level statistics are reported as the median of per-well amplitudes within each cluster, preserving the visibility of within-cluster spread that would be suppressed by computing a single amplitude on the cluster-mean hydrograph.

Because the post-2018 window contains a higher proportion of drought summers than the pre-2018 window (two of eight summers versus one of thirteen, identified empirically from the RAF Valley Jun–Sep rainfall record against the 1931–2017 baseline), a climate-normalised variant was also computed. Monthly observations from empirically identified drought summers (2005, 2018, 2022; threshold: Jun–Sep rainfall total $\geq 1\sigma$ below the long-term mean of 260 mm) were removed from both windows before recomputing p90 − p10. This normalisation is reported alongside the raw values as a sensitivity check; where the two agree in direction and approximate magnitude, the raw value is taken as primary. The normalisation predominantly affects the post-2018 window, since only wells whose records extend back to 2005 — principally the NW-series — are modified in the pre-2018 window.

Amplitude results are presented in Section 4.2 as a secondary descriptor alongside pattern-coherence and bootstrap stability. Where the ratio of maximum to minimum per-well amplitude within a cluster approaches or exceeds 2×, the cluster is noted as amplitude-heterogeneous; this is relevant to the interpretation of cluster-mean β coefficients from the state-space model (Section3.4), which represent an average response magnitude across wells whose individual magnitudes may differ meaningfully.

Clustering, silhouette validation, and cluster assignment were implemented in 02_clustering.py. Cluster map was produced by 04_cluster_visualisations.py.

3.3 Checking Cluster Allocation with Pearson Affinity Analysis

While hierarchical clustering imposes discrete boundaries, the physical reality of a highly transmissive sand aquifer is one of continuous hydraulic gradients. To evaluate the rigidity of the Ward's linkage assignments and quantify the gradational nature of the aquifer, a secondary Pearson Affinity Analysis was conducted to measure the probabilistic membership of each well to all identified clusters (Rao and Srinivas, 2006).

3.3.1 Cluster Membership and Affinity Scoring

First, a cluster average hydrograph time series, $\bar{X}_k(t)$, was generated for each of the five identified clusters (k=1,…,5; shown in Figure 8) by calculating the mean standardized response of all wells definitively assigned to that cluster by the Ward's algorithm.

The temporal affinity of every individual well i to each cluster centroid k was then quantified using the Pearson correlation coefficient ($r_{i,k}$) over contemporaneous overlapping records of length N:

$$r_{i,k}=\frac{\sum_{t=1}^{N}\left(x_i(t)-\bar{x}_i\right)\left(\bar{X}_k(t)-\bar{\bar{X}}_k\right)}{\sqrt{\sum_{t=1}^{N}\left(x_i(t)-\bar{x}_i\right)^2\sum_{t=1}^{N}\left(\bar{X}_k(t)-\bar{\bar{X}}_k\right)^2}}$$

3.3.2 Classification Criteria

To categorize the spatial transition of the groundwater dynamics, let r_{1st} represent a well's highest correlation with any cluster centroid, and r_{2nd} represent its second-highest correlation. The affinity margin, $\Delta r = r_{1st} - r_{2nd}$, serves as a mathematical index of well hydrograph behaviour. Based on these metrics, each well in the 66-well reference network was classified into one of four structural categories. The same affinity procedure was subsequently applied to the extended network of wells with shorter or earlier records allowing the hydrogeological classification to be extended beyond the reference network boundary(Section 3.3.3).

- **Core ($\Delta r \geq 0.05$):** The well exhibits a dominant, unambiguous affinity to a single cluster, and its highest correlation matches its spatial assignment from the Ward's linkage. These represent the definitive hydrological signature of their respective geomorphological zones.
- **Fuzzy ($\Delta r < 0.05$):** The well exhibits highly competing affinities between its primary and secondary clusters. This mathematical ambivalence identifies physical ecotones—transitional zones where the groundwater regime shifts smoothly between two distinct structural blocks (e.g., the boundary between the Western Residual and the Forest margin).
- **Spy:** The well's highest temporal affinity (r_{1st}) belongs to a different cluster than the one assigned by the Ward's spatial linkage. This divergence highlights localized hydrogeological anomalies where sub-surface stratigraphy overrides broad spatial proximity (Hannah et al., 2000).

- **Unclassified:** The temporal record is of insufficient length or overlap to generate statistically significant correlation coefficients.

Multi-Cluster Affinity (MCA) Finally, Multi-Cluster Affinity (MCA) flags were raised for any well where three or more distinct cluster centroids achieved a correlation of r>0.90. An MCA flag mathematically identifies highly generic, highly responsive boundary wells that lack a specialized local signature, acting instead as structural nodes governed by system-wide climatic forcing rather than local topography.

3.3.3 Extended Network Classification

To extend the hydrogeological classification beyond the 66-well reference network, the Pearson affinity procedure was applied to all wells with a minimum of 24 months of record that did not meet the reference network criteria. The reference network cluster averages derived in Section 3.3.2 were used as fixed classification templates; these were not recalculated from the extended wells. Prior to correlation, all wells — both reference and extended — were z-score standardized jointly as a single combined dataset, ensuring that extended wells with shorter or earlier records were standardized relative to the same temporal distribution as the reference network rather than independently.

Extended wells were classified using the same affinity margin threshold (Δr = 0.05) and MCA criteria (r > 0.90 with three or more centroids) as the reference network audit. Because extended wells carry no prior Ward's cluster assignment, the Spy classification is not applicable; extended wells are assigned as either Core or Fuzzy based solely on the strength of their best-match affinity.

Affinity classification of the 66-well reference network was implemented in 05_pearson_affinity.py; extension to the full 117-well monitoring array was implemented in* 06_pearson_extended.py.*

3.4 Using Models to derive the Aquifer Architecture

Predicting how a dune aquifer responds to rainfall and evaporation over time requires a model that captures not just the immediate effect of each month's weather, but also the aquifer's memory of its own state. A simple approach — the Traditional Linear Model (TLM) used as a benchmark in Section 3.4.2 — treats the aquifer as a bucket that fills with rainfall and empties at a fixed background rate, regardless of how full it is (von Asmuth et al., 2002). This works reasonably well when the model is continuously corrected with observed dipwell readings, but fails when run autonomously over longer periods because it has no mechanism to slow drainage as the water table drops or accelerate it when the water table is high (Hypolite et al., 2021; Young, 2011). The state-space model (SSM) captures this physical reality by adding a head-dependent drainage term that scales drainage rate to the current water table depth, consistent with the functional relationship between groundwater head and lateral discharge demonstrated in shallow sandy aquifers (Gumuła-Kawęcka et al., 2021; Manis et al., 2022). This single addition transforms the model from a statistical curve-fit into a physically constrained representation of aquifer behaviour, allowing it to simulate water table dynamics autonomously over multi-year periods without accumulating systematic error (Knotters and van Walsum, 1997). The baseline formulation is:

$$h(t)=\beta_1\cdot P(t)-\beta_2\cdot PET(t)-\beta_3\cdot h(t-1)$$

where h(t) is depth-to-water-table; P(t) and PET(t) are monthly rainfall and potential evapotranspiration; β_1 is the recharge sensitivity (per-month water table rise per unit rainfall); β_2 is the atmospheric draw (per-month deepening per unit PET); and β_3 is the storage decay coefficient governing the rate at which antecedent water table depth is dissipated by lateral drainage. The sign convention used throughout this report — including the $-\beta_3$ notation that appears in the equation and in subsequent figures and tables — is described in Section 3.4.1.

The drainage term is formulated using displacement above a reference drainage base rather than raw depth below the ground surface. Displacement is defined as h_disp = D + h_depth, where D is the drainage datum depth (metres below ground surface) and h_depth is the depth-to-water in the negative convention. This reformulation ensures that β_3 is positive (Darcy-consistent) across all five clusters; fitting on raw depth produces negative β_3 for three of five clusters, indicating that the depth-below-surface coordinate does not correctly represent the hydraulic gradient driving lateral drainage. The datum D was selected by sweeping reference depths from 0.5 to 8.0 m in 0.1 m steps, refitting the cluster-centroid SSM at each depth and recording β_3, its p-value, R^2 and AIC (Burnham and Anderson, 2002) for all five clusters. The selection criterion was the minimum depth at which all five clusters simultaneously produce positive and statistically significant ($p < 0.05$) β_3. The C4 Main Forest cluster — which has the deepest mean water table in the network and is therefore closest to the drainage base — is the binding constraint: C1, C2, C3 and C5 reach significance at datums of 0.5–1.3 m, but C4 requires 3.7 m. This value is therefore adopted as the network-wide datum to ensure consistent positive displacement across all clusters. As more monitoring data have accumulated since the original calibration, the empirical minimum required to satisfy the criterion has shifted shallower (current pipeline value 1.7 m); 3.7 m is retained as the canonical value to maintain comparability with all results derived under the displacement formulation. R^2 and AIC are insensitive to the datum across the viable range, confirming that the choice does not distort model fit. The displacement formulation does not affect β_1 or β_2, which are invariant to the datum constant, nor Δh, in which the constant cancels. The per-well spatial distribution of optimal datum depths and the R^2 cost of the uniform assumption are presented in Section 4.9.1 (Figure 39).

The SSM formulation adopted here sits within an established lineage of transfer-function-noise (TFN) and linear reservoir models for groundwater head series, in which monthly or daily head is expressed as a linear response to climatic forcing with an autoregressive memory term (von Asmuth et al., 2002; Knotters and Bierkens, 2000; Peterson and Western, 2014; Peters et al., 2003). Recent open-source implementations — notably Pastas (Collenteur et al., 2019) and HydroSight (Peterson and Western, 2014) — provide impulse-response calibration at individual wells and have been applied to shallow unconfined aquifers across a range of climatic settings (Obergfell et al., 2019; Bakker and Schaars, 2019). The formulation used in this study departs from these implementations in three ways developed in Section 3.4.1: it is fitted at cluster-average scale rather than individual-well scale; it retains the storage decay coefficient explicitly rather than absorbing it into a convolution kernel; and it constrains the intercept to zero. These choices enable the Lumped Catchment Storage Coefficient (LCSC) to be derived directly from the fitted parameters and inverted algebraically into the P_flood threshold equations reported in Section 3.6.3.

3.4.1 Modelling the Cluster water levels

The state-space model defined in Section 3.4 was applied at the cluster level to derive mechanistic coefficients representative of each hydrogeological zone. For each of the five clusters, a mean monthly time series was computed by averaging the upstand-corrected records of all member wells (Section 3.1.1), producing a single representative hydrograph of seasonal water table behaviour. The SSM was fitted to this average hydrograph using the full available record rather than the 100-month window applied to individual wells (Section 3.4.3), yielding a single set of coefficients (β_1, β_2, β_3) representative of the cluster as a whole. Cluster-scale fitting reduces the influence of individual well idiosyncrasies and local topographic noise, producing a parsimonious mechanistic characterisation of each hydrogeological zone.

Sign convention. β_3 is defined and reported as a positive quantity throughout this study. The negative sign appears explicitly in the model equation as $-\beta_3 \cdot h(t-1)$ to indicate that storage decay acts as a loss term: a deeper antecedent water table generates greater drainage, tending to restore equilibrium. Where the notation $-\beta_3$ appears in subsequent figures, tables and analytical sections (including the BACI, scraping and cluster mechanistic analyses), it refers to this loss term and should be read as the magnitude of the drainage decay coefficient. An increase in $-\beta_3$ therefore indicates accelerated drainage; a decrease indicates retarded drainage or greater storage retention.

Intercept constraint. The cluster regression was fitted with the intercept strictly forced to zero. This enforces a strict physical mass balance: in the absence of external climatic forcing (P = 0, PET = 0), the aquifer cannot spontaneously generate or lose water beyond the passive drainage decay dictated by its own antecedent storage state ($-\beta_3 \cdot$h_t−1). The constraint also enables direct algebraic inversion into the P_flood threshold equations (Section 3.6.3).

Physical interpretation of β_3. The storage decay coefficient provides a parsimonious empirical representation of head-dependent lateral drainage that does not require explicit knowledge of hydraulic conductivity, aquifer geometry or hydraulic gradient — parameters that are either unmeasured or spatially variable across a 700 ha dune system where borehole logs are absent across the majority of the study area. β_3 is functionally equivalent to a linearised Darcy flux under the assumption that the hydraulic gradient scales proportionally with water table height above the drainage base, which is a reasonable approximation for a shallow unconfined dune aquifer draining to a fixed boundary (Knotters and van Walsum, 1997).

Lumped Catchment Storage Coefficient. The LCSC is defined as the reciprocal of the recharge sensitivity, expressed as a percentage (LCSC = $100/\beta_1$). Expressing it as LCSC% converts the recharge sensitivity into physically intuitive units — the millimetres of rainfall required to raise the water table by 10 cm — and scales in the same direction as storage capacity, making it directly comparable to the empirical specific yield concept used by site managers. A higher LCSC therefore indicates a more buffered, less responsive aquifer: a shallow aquifer over impermeable till produces a large water table rise from a small rainfall event, a high β_1, and a low LCSC, whilst a deep sand aquifer absorbs rain across a thick unsaturated zone before the water table rises appreciably, producing a low β_1 and a high LCSC.

3.4.2 State-Space vs. Traditional Linear Models for Diagnostic Validation of Aquifer Architecture

Model Configurations . Both models were independently fitted to the most recent 100 months of data to predict the monthly change in water table elevation (Δh_t) using the same climatic forcing data, differing only in their structural representation of the aquifer:

1. **Traditional Linear Model (TLM)** treats the aquifer as a passive linear store driven purely by atmospheric inputs, with a fitted intercept term capturing background drift and unmeasured forcing, but no mechanism representing the aquifer's memory of its own state (von Asmuth et al., 2002):

$$\Delta h_t = \beta_1 P_t - \beta_2 PET_t + c$$

2. **State-Space Model (SSM):** The physically constrained mass-balance framework defined in Section 3.4, fitted through the origin (c = 0), with antecedent head h(t−1) included as a feedback term that causes drainage to accelerate when the water table is high and slow when it is low

$$\Delta h_t = \beta_1 P_t - \beta_2 PET_t - \beta_3 h_{t-1}$$

Forecasting Modes: Diagnostic vs. Iterative

Model performance was evaluated under two distinct operational modes to separate baseline statistical fit from true forecasting stability:

- **One-Step Diagnostic Fit:** The model predicts the change in head for month **t** using the *observed* field measurement of h(t−1) from the prior month. This tests the fundamental validity of the regression coefficients without compounding error.
- **Iterative Forecasting Stability:** The model predicts the change in head for month **t** using its own simulated head from the previous timestep (h_sim), relying entirely on its own simulated state over the 100-month evaluation window without periodic data assimilation; this tests resistance to cumulative drift over multi-year autonomous simulation.

Performance Metrics. Diagnostic fit (one-step) was evaluated using the Coefficient of Determination (R^2). Iterative forecasting stability was evaluated using R^2, the Root Mean Square Error (RMSE), and the Nash-Sutcliffe Efficiency (NSE) (Nash and Sutcliffe, 1970; Krause et al., 2005).

The RMSE is mathematically defined as:

$$R^2 = \left(\frac{\sum \left(h_{obs} - \bar{h}_{obs}\right)\left(h_{\square} - \bar{h}_{\square}\right)}{\sqrt{\sum \left(h_{obs} - \bar{h}_{obs}\right)^2 \sum \left(h_{\square} - \bar{h}_{\square}\right)^2}} \right)^2$$

The NSE is mathematically defined as:

$$NSE = 1 - \frac{\sum \left(h_{obs,t} - h_{,t}\right)^2}{\sum \left(h_{obs,t} - \bar{h}_{obs}\right)^2}$$

The NSE is the most informative metric for continuous hydrological forecasting. An NSE=1.0 indicates perfect predictive agreement, an NSE=0 indicates the model is only as accurate as using the long-term observed mean, and an NSE<0 indicates that the model is actively worse than using the observed mean, typically due to unchecked cumulative mass-balance drift.

Spatial Diagnostics of Subsurface Architecture using SSM and TLM comparison. Plotting the comparative performance metrics, defined as the improvement in variance explanation:

$$\Delta R^2 = R^2_{SSM} - R^2_{TLM}$$

and the improvement in iterative forecasting stability:

$$\Delta NSE = NSE_{SSM} - NSE_{TLM}$$

serves as more than a statistical validation; it acts as a spatial diagnostic tool for the underlying geological architecture (Beven, 2012; Gupta et al., 2008). Similar multi-methodological approaches have been employed to constrain recharge estimates and reduce uncertainty in shallow sandy systems by evaluating the sensitivity of model boundaries to observed head variations. Because the fundamental mathematical difference between the two models is the inclusion of the head-dependent internal storage decay coefficient ($-\beta_3 \cdot h(t-1)$), the spatial distribution of model improvement physically delineates the subsurface structural boundaries of the aquifer:

- **High ΔR^2 and ΔNSE zones:** Regions where the SSM most substantially outperforms the TLM are those where drainage is most strongly head-dependent at the monthly timescale (Fetter, 2001; Robins et al., 2013). In these zones the proportional drainage rate varies considerably between high and low water table states, making the $-\beta_3$ term essential for accurate autonomous simulation. The Western Block's deep sand buffer moderates this dynamic, and canopy interception in the Forest cluster smooths it further, producing a gradient of decreasing SSM gain from east to west that mirrors the structural gradient identified by clustering.
- **Low ΔR^2 and ΔNSE zones:** Conversely, areas exhibiting small or marginally negative improvement between the two models reflect zones where drainage is least strongly head-dependent. A small or zero SSM gain in these zones does not indicate model failure but confirms that the TLM's linear approximation is adequate where drainage dynamics are sufficiently buffered or smoothed (Gumuła-Kawęcka et al., 2021). The two C4 ridge-flank wells where the SSM underperforms TLM (CEH14 strongly, CEH13 mildly) are the same wells that carry the network's largest water balance residuals (Section 4.9.6, Figure 45): both findings point to a hydrological setting at the bedrock-ridge interface where the standard rainfall–PET–drainage framework characterizes the system less completely than at typical aquifer locations.

By mapping the success or failure of the storage decay coefficient -β_3 , this comparative methodology physically draws the structural boundary lines between deep sand reserves, perched slacks, and complex discharge zones without requiring extensive invasive coring.

SSM versus traditional linear model benchmarking across all 66 reference wells was implemented in 08_model_benchmarking.py.

3.4.3 Spatial Coefficient Mapping

To visualise the spatial structure of aquifer behaviour across the site, per-well SSM coefficients (β_1, β_2, β_3) and model fit quality (R^2) were interpolated onto a regular 25 m grid using inverse-distance weighting (IDW, power = 2) and rendered over a 1 m LiDAR DEM hillshade. The per-well coefficients are derived from the displacement-formulation SSM (Section 3.4) fitted independently at each of the 66 reference wells, using the same 3.7 m drainage datum as the cluster-centroid model. The β_3 drainage surface was plotted on a logarithmic scale to resolve the order-of-magnitude contrast between the fast-draining lake-edge wells and the near-stagnant forest interior. The resulting coefficient atlas (Figure 40a–d) maps the model fit quality, recharge sensitivity, atmospheric draw, and drainage rate across the full dune field, providing a spatially continuous expression of the mechanistic contrasts identified at cluster level in Table 3. *Spatial coefficient mapping was implemented in 07_spatial_coefficients.py.*

3.4.4 Within-Forest Spatial Coefficient Analysis

To test whether within-forest coefficient variation is driven by substrate properties or canopy differences, per-well SSM coefficients for the 14 forest zone wells (9 in C4, 5 in C5) were regressed against three spatial predictors: DEM ground elevation, Euclidean distance from the ridge crest (E = 241,750, N = 364,500), and Easting as a proxy for distance from the south-eastern coast (Menai Strait). Single-predictor Pearson correlations were computed for each coefficient against each spatial variable, and multiple regression was used to test whether distance from the ridge added explanatory power beyond elevation alone. The C4/C5 boundary was assessed for continuity or discontinuity using t-tests on elevation, β_1, β_2 and β_3 between the two clusters. *Results are reported in Section 4.9.4 (Table 17). The analysis is implemented in 10c_forest_zone_analysis.py.*

3.5 Anthropogenic Interventions

3.5.1 The Before-After Control-Impact (BACI) Design

To evaluate the hydrological outcomes of anthropogenic interventions at the study site, a Before-After Control-Impact (BACI) experimental design was used (Stewart-Oaten et al., 1986). This framework is designed to isolate the effect of a specific localized disturbance from broader natural variability, such as seasonal weather patterns or regional climatic drift.

The fundamental mechanism of the BACI approach involves contrasting an impact zone (where the intervention occurred) with a reference control zone (which remained undisturbed) across designated pre-intervention and post-intervention periods. Rather than simply measuring absolute changes in water levels, the model calculates the difference (Δ) between the impact and control sites. The actual effect of the intervention is subsequently quantified by measuring the shift in this difference after the event.

For highly complex environments where multiple interventions overlap, the standard BACI method was developed into a Hierarchical Nested Control design to ensure robust baseline validation. By using multiple control sites to defend against localized hydrogeological anomalies, this study employs a "Beyond BACI" approach (Underwood, 1992).

General Before-After Control-Impact (BACI) Calculation. The standard BACI effect size isolates the intervention signal from natural background variation. The fundamental equation calculates the difference in means (μ) between the Impact (I) and Control (C) sites across the Before (B) and After (A) periods:

$$BACI_{\text{effect}} = \left(\mu_{I,A} - \mu_{C,A}\right) - \left(\mu_{I,B} - \mu_{C,B}\right)$$

Where:

- $\mu_{I,A}$ is the mean water level at the impact site post-intervention.
- $\mu_{C,A}$ is the mean water level at the control site post-intervention.
- $\mu_{I,B}$ and $\mu_{C,B}$ represent the respective means prior to the intervention.

Cumulative Sum (CUSUM) Statistical Verification. To validate the precise temporal onset and permanence of hydrological shifts identified by the BACI framework, a Cumulative Sum (CUSUM) sequential analysis was integrated into the analytical framework (Page, 1954). While the BACI contrast quantifies the magnitude of an intervention's effect, the CUSUM statistic is used to identify the exact temporal structural break — or regime shift — within the time series data. Era-specific estimation of the storage decay coefficient $-\beta_3$ is additionally reported across treatment and control wells; changes that are confined to the intervention footprint would be consistent with a genuine drainage response, whilst changes extending uniformly to unmanipulated control wells indicate a network-wide signal attributable to climate or SSM fitting artefacts rather than the intervention itself.

Mathematical Framework. The CUSUM method operates by calculating a running total of the deviations of individual data points from an established baseline mean. In the context of this study, the input variables (x_i) are the sequential differences between the impact and control wells (Δ), and the baseline (μ_baseline) is the mean of these differences during the pre-intervention period. The cumulative sum at time t is formalized as:

$$C_t = \sum_{i=1}^{t} \left(x_i - \mu_{\text{baseline}}\right)$$

Interpretation of the CUSUM Trajectory. The resulting CUSUM trajectory provides a clear visualization of system stability:

- **Baseline Stability:** A stable relationship between the impact and control sites yields random fluctuations around the baseline mean, producing a roughly horizontal CUSUM trajectory.
- **Structural Break Detection:** A sustained shift in the hydrological relationship — such as that caused by canopy removal or slack scraping — results in directional deviations that accumulate rapidly, manifesting as a distinct change in the gradient of the CUSUM line.
- **Temporal Verification:** The exact point of inflection, where the trajectory breaks from its horizontal trend, marks the initiation of the intervention's physical effect.

By aligning these inflection points with the documented dates of the December 2017 clearfell and the 2015 and 2023 scraping interventions, the analysis confirms that the observed divergence is a direct consequence of the site works rather than gradual environmental drift. To isolate the localized effects of topographical and silvicultural interventions from background climatic drift, all CUSUM trajectories were calculated on the BACI delta (impact minus control) rather than raw water table depth. This ensures that any identified structural breaks represent a true divergence from the undisturbed regional baseline.

Where an impact site has experienced more than one intervention over the monitoring period, or where a secondary disturbance confounds the primary signal, the record is subdivided into discrete management eras defined by the dates of known interventions. Each era is treated as a distinct analytical period with its own mean water level, inter-era BACI delta and $-\beta_3$ estimate. This allows the analysis to separate the signal of a specific intervention from preceding baseline conditions and from subsequent confounding events — for example, isolating the effect of dune scraping before the onset of a plantation clearance pulse that would otherwise inflate or mask the scraping response. Era boundaries are defined a priori from documented intervention dates rather than from the data, preventing circular reasoning in the identification of structural breaks.

3.5.2 Hierarchical Nested Control Application

This framework operates across two distinct analytical tiers:

- **Tier 1 (Background Drift Validation):** Local control wells are first evaluated against a broader regional mean. This step verifies that the designated control sites are behaving consistently with macro-environmental trends and are not subject to anomalous localized influences, such as unmapped coastal drawdown. This tier verifies that the local control (C_local) accurately tracks the broader regional mean (C_region). The temporal difference (Δ_drift) at any time t is calculated as:

  $$\Delta_{\text{drift},t} = C_{\text{local},t} - C_{\text{region},t}$$

  A structurally stable Δ_drift,t over the baseline timeline validates the local control for use in Tier 2.

- **Tier 2 (Intervention Signal Isolation):** Once validated, these local controls are paired directly with specific impact wells. This pairing subtracts the regional noise and secondary intervention pulses, isolating the pure signal of the targeted experiment. The Tier 2 BACI effect calculates the net benefit of the intervention by contrasting the impact well (I) against the validated local control (C_local) across multiple distinct eras (E):

  $$\text{Net Benefit}_{E1 \rightarrow E2} = \left(\mu_{I,E2} - \mu_{C_{\text{local}},E2}\right) - \left(\mu_{I,E1} - \mu_{C_{\text{local}},E1}\right)$$

  This formulation systematically removes regional climatic noise and secondary pulses, standardizing the output to isolate the pure scraping signal.

3.5.3 Change in the (-β_3) Storage Decay Coefficient

Both the clearfell and scraping interventions may cause changes to subsurface drainage mechanics, assessed through era-specific estimation of the storage decay coefficient $-\beta_3$. To reduce the confounding influence of inter-era variation in rainfall and evapotranspiration, a two-step procedure was applied. First, β_1 and β_2 were estimated by fitting the full SSM to the complete well record.

$$\Delta h = \beta_1 P_m - \beta_2 PET + \beta_3 h_{prev} + \epsilon$$

The climate forcing contribution was then removed from each monthly observation to yield a drainage residual series (D_residual), representing the component of water table change attributable to head-dependent drainage after removing the estimated rainfall and evapotranspiration contributions:

$$D_{residual} = \Delta h - \beta_1 P_m + \beta_2 PET$$

Second, for each management era separately, $-\beta_3$ was estimated by regressing this drainage residual against −h_prev with a constant term:

$$D_{residual} = -\beta_3 \cdot \left(-h_{prev}\right) + c + \epsilon$$

This procedure reduces the influence of differential climate forcing between eras, but does not eliminate all sources of artefact: where the water table undergoes a sustained directional shift between eras — as occurs following both clearfell and scraping — the h_prev term itself tracks a non-stationary trajectory, and the model may interpret systematic head changes as an elevated drainage signal. Era-specific $-\beta_3$ estimates should therefore be interpreted cautiously, and spatial extent is used as the primary diagnostic: changes confined to the intervention footprint are consistent with a genuine drainage response, whilst changes extending uniformly to unmanipulated control wells point to a network-wide artefact rather than an intervention effect.

3.5.4 Dune Scraping Intervention Analysis

A two-tier hierarchical BACI framework was applied across five focal wells. Three wells served as scraping treatment sites — CEH36 (central Warren, scraped April 2015), CEH18 (western boundary, scraped October 2023) and CEH21 (coastal margin, scraped October 2023) — and two as local controls: CEH4 (central Warren, unmanipulated) and CEH22 (coastal margin, unmanipulated). A regional mean was computed as the arithmetic mean of five C3 climate control wells (CEH9, NW7, NW6, NW5, WMC2) to characterize background system-wide trends. NW8 and NW8B were excluded on the grounds of data quality.

Because the three treatment wells span different management histories, each was decomposed into site-specific eras defined by intervention dates:

- CEH36 and its paired control CEH4 were decomposed into three eras: Baseline — pre-April 2015; Pure Scraping — April 2015 to November 2017; Felling Pulse — December 2017 onwards, capturing the confounding influence of the plantation clearance.

- CEH18 and CEH21 were decomposed into three eras: Baseline — pre-December 2017; Felling Pulse — December 2017 to October 2023; After Scraping — October 2023 onwards.

Three limitations of the control wells and treatment records bear directly on interpretation and are stated here to avoid repetition in the results.

First, CEH4 is subject to two independent sources of instability: progressive ground surface lowering from coastal erosion, and sustained water table deterioration attributable to the post-2017 felling pulse confirmed by its Tier 1 CUSUM trajectory. Net benefit estimates at CEH36 during the Pure Scraping era should be treated as upper bounds; an additional upward bias may arise from contemporaneous foredune build-up at the site contributing to a rising water table independently of the scraping intervention. Net benefit estimates at CEH18 over the After Scraping era are similarly subject to caution given CEH4's post-felling deterioration. The robustness of the CEH36 scraping signal is tested independently in Section 4.5 using a synthetic control constructed from donor wells calibrated on the pre-scraping period and by SSM forward residual analysis (*09e_robustness.py*).

Second, CEH22 experienced accelerating deterioration throughout the monitoring period that, as the network-wide $-\beta_3$ analysis demonstrates, reflects a system-wide signal rather than a process localized to the coastal margin. CEH22 does not constitute a stable Tier 1 baseline; net benefits are therefore not reported for CEH21, and raw Tier 2 BACI shifts are reported for completeness only.

Third, CEH18 and CEH21 were scraped in October 2023, providing only two full post-scraping summers by the end of 2025. The short post-intervention record limits the statistical power of pre- vs post-scraping comparisons at these wells and precludes any conclusion about whether the scraping benefit observed at CEH36 over a longer timeframe will be replicated at the secondary sites.

Fourth, the pre-scraping baseline length differs substantially across the focal wells: CEH4 (paired control, installed May 2006) provides 8.9 years; CEH36 (central treatment well, installed January 2011) provides 4.2 years; CEH18 (installed September 2007) provides 7.6 years pre-2015; and CEH21 (installed July 2010) provides 4.8 years pre-2015. CEH36's shorter pre-scraping baseline does not invalidate the analysis: the scraping methodology does not require estimation of a long-baseline parametric covariate structure (in contrast to the clearfell ANCOVA, Section 3.5.5, which estimates a cumulative-water-balance slope from pre-felling data). CUSUM, Welch's t-test on era means, synthetic-control calibration and SSM forward residual all operate effectively with the available pre-event window, and the robustness analysis (09e, Section 4.5.2) tests the resulting BACI step against three independent methods to confirm convergence.

Fifth, the scraping analysis treats CEH4's progressive ground-surface lowering from coastal erosion as a confound documented in the era-by-era comparisons rather than as a covariate explicitly modelled in the BACI design. CEH4 sits ~340 m inland of the southern foreshore, and the characteristic propagation timescale (L^2Sy/K, Section 5.4.3) for coastal-erosion effects to reach inland wells through the aquifer is multi-year to multi-decadal. Most of the coastal-erosion propagation envelope therefore lies outside the analysis window, with the visible deterioration at CEH4 reflecting both the integrated effect of historical coastal retreat and a contribution from the post-2017 felling pulse. The Tier 1 CUSUM analysis at CEH4 captures this combined drift but cannot decompose it into its component drivers; the synthetic control and SSM forward residual analyses (09e) provide the necessary alternative counterfactuals.

Paired BACI and β_3 testing (09a). Tier 1 CUSUM analysis evaluated each local control well against the regional mean to verify that the controls were tracking system-wide background trends consistently. Tier 2 paired CUSUM analysis differenced each treatment well against its matched local control to isolate the scraping signal from regional climatic drift. Net benefit was computed as the BACI shift in the treatment-minus-control difference between consecutive eras as defined in Section 3.5.1. The displacement-formulation SSM (Section 3.4, with drainage datum D = 3.7 m) was re-fitted independently for each era at treatment and control wells, and β_3 compared across eras with 95% confidence intervals, interpreted as described in Section 3.5.3. Results are presented in Section 4.5.1 (background drift), Section 4.5.2 (paired BACI shifts) and Section 4.5.3 (drainage coefficient response).

Scraping propagation analysis (09b). A separate analysis tested whether the CEH36 scraping event — which created a topographic low acting as a local drain — produced a detectable drawdown at wells uphill of the scraped site. Groundwater flow toward the new drain would steepen the local hydraulic gradient, lowering the water table at neighbouring wells. The SSM was fitted independently to pre-scraping (start of record to March 2015) and post-scraping (April 2015 to November 2017) windows for each of 10 wells located north and northwest of CEH36, positioned along the expected groundwater flow path into the forest interior. Seven distant C3 wells (NW1, NW2, NW11, NW13, WMC4, D25, WMC2), located 850–1100 m from CEH36 and well beyond any expected scraping influence, served as BACI controls. Per-well coefficient shifts ($\Delta\beta_1$, $\Delta\beta_2$, $\Delta\beta_3$) were corrected by subtracting the control centroid shift, isolating any scraping-specific perturbation from background climate drift. A distance-decay regression tested whether the BACI-corrected $\Delta\beta_3$ diminished with distance from CEH36, as would be expected if the scraping drain were pulling water from the uphill aquifer. Results are presented in Section 4.5.5.

Summer minima analysis (09c). The annual summer minimum depth (June–September) was computed for each of the five focal wells. For each well, the gap between the well's summer minimum and both the climate control centroid and the paired control summer minimum was computed annually, and the pre- vs post-scraping shift tested by Welch's t-test. This mirrors the methodology applied to the clearfell intervention (Section 3.5.5, 10d) and addresses the ecological question directly: does scraping improve the critical summer low that determines dune slack habitat viability? Results are presented in Section 4.5.4.

Scenario comparison (09d). To contextualise the observed scraping benefit at CEH36, five alternative scenarios were computed at the same well using its own SSM coefficients (β_1 = 2.68, β_2 = 1.08, β_3 = 0.046) and its own per-well WTF-derived specific yield (Sy = 0.36, from Section 4.2.4). Three forestry scenarios — clearfell, 50% thinning, and broadleaf conversion — are hypothetical, asking "what if CEH36 had pine canopy and was then managed?" These apply interception and atmospheric draw perturbations from the forestry analysis (Section 3.5.5) to CEH36's SSM. Two climate scenarios (dry and wet) apply UKCP18 RCP8.5 2050s central estimates. Monthly volumetric water table change (mm water equivalent per month) was computed from the SSM equilibrium response. Summer minimum depth change (mm) was computed from the observed paired BACI summer minimum shift (CEH36 vs CEH4, +165 mm) for scraping, and from SSM equilibrium predictions scaled by an empirical summer amplification factor for the alternative scenarios. All parameters are read from upstream pipeline outputs (Scripts 01, 03, 17). Non-scraping scenario perturbations (clearfell, thinning, broadleaf, climate) are computed by a shared function (`compute_scenario_bars()` in `scraping_common.py`) that is also called by Script 21, ensuring consistency between the site-level and cluster-level scenario comparisons. Results are presented in Section 4.5.6.

Robustness analysis (09e). Three independent estimates of the CEH36 Pure Scraping era step change were computed to test whether the BACI result is robust to the choice of counterfactual: (1) the raw paired BACI (CEH36 minus CEH4); (2) a synthetic control constructed from a weighted composite of donor wells optimised to match the pre-scraping trajectory; and (3) an SSM forward residual, where the SSM was calibrated on the pre-scraping baseline and run forward using observed climate forcing, with the residual (observed minus predicted) measuring the non-climatic component of the water table change. Method convergence across the three estimates supports the inference that the scraping benefit is not an artefact of CEH4's own progressive deepening. Results are presented in Section 4.5.2.

The scraping analysis was implemented as a modular suite (run_09_scraping.py: 09a–09e) with shared data loading and well definitions in scraping_common.py.

3.5.5 Clearfell BACI experiment

A five-tier hierarchical BACI experiment was conducted across 17 wells to assess the hydrological consequences of the December 2017 clearfell. Wells were assigned to tiers defined a priori from spatial proximity to the felled compartment and cluster membership, with all wells required to have a minimum of eight years' pre-felling baseline to ensure stable coefficient estimation:

- **Impact:** WMC3 (45 m from felling compartment centroid E=241210, N=363607, inside the cleared area; 8.4 years pre-felling record).
- **Edge:** CEH31, CEH20, CEH30, CEH16 (152–229 m, adjacent to the felling boundary; 7.4–10.3 years pre-felling).
- **Forest control:** CEH32, CEH34, CEH33, NW10, CEH2 (C4 Main Forest interior, unaffected by felling, same canopy cover; 306–428 m; 7.3–11.6 years pre-felling).
- **Coastal control:** CEH19, CEH17 (C5 Coastal Forest; separated from C4 on the basis of a distinct β_2 regime and non-overlapping elevation range — Section 4.2; 9.7–10.3 years pre-felling).
- **Climate control:** CEH9, NW7, NW6, NW5, WMC2 (C3 Western Block, no forest canopy; 571–1056 m; 8.8–12.6 years pre-felling).

Wells with insufficient pre-felling baseline were excluded from the primary analysis. The candidate-well criterion adopted here was monitoring continuity from the 2010 network expansion onwards, providing a minimum of seven years of pre-felling record and at least four years of pre-scraping record before the April 2015 scraping. This baseline length is required for two reasons. First, the April 2015 scraping at CEH36 (Section 4.5) introduced a level shift in the BACI gap at the same date, requiring a scraping covariate in the ANCOVA (Section 4.6.1) to avoid biasing the clearfell estimate; reliable estimation of this covariate requires that included wells have monitoring data spanning both pre- and post-scraping periods so that the April 2015 step is identifiable separately from the December 2017 felling step. Second, multiple full annual cycles are needed to constrain the seasonal cycle and the cumulative water balance slope independently of the felling step. FE1, FE2 (both installed August 2015) and NW8B (July 2015) were installed after the 2010 expansion and within four months of the scraping; their entire pre-felling baseline sits within the post-scraping era. FE3 (April 2017), FE4 (November 2017) and LIS1 (November 2017) provide eight months or fewer of pre-felling record. Following general BACI guidance that pre-intervention monitoring should span sufficient years to characterize inter-annual variability and resolve overlapping perturbations (Smokorowski and Randall, 2017), these wells were excluded. CEH36 was excluded as the scraping treatment site itself: its post-scraping water table record does not represent undisturbed pre-intervention conditions.

The separation of controls into three counterfactual pools — Forest, Climate and Combined — allows the clearfell question to be asked three ways: (a) did the felled wells diverge from unfelled C4 wells sharing the same canopy, substrate and elevation range, the most position-matched comparison; (b) did they diverge from the climate trajectory defined by unforested C3 wells in the Western Block; and (c) against a blended counterfactual pooling all 12 control wells (C4 Forest, C5 Coastal Forest and C3 Climate). The Coastal Control tier (C5: CEH19, CEH17) is not used as a primary counterfactual in the three-pool ANCOVA — its distinct β_2 regime and non-overlapping elevation range from C4 mean it is not position-matched to the felled wells — but it appears in the Combined pool and is analysed separately as a coastal-erosion diagnostic in the summer minima and coefficient decomposition analyses (Sections 4.6.4 and 4.6.6).

The BACI gap for each target tier (Impact, Edge) was calculated as the target centroid mean water level minus the control centroid mean, producing a monthly time series of (target − control) values for each of the three control definitions. This monthly BACI gap is distinct from the SSM displacement (Section 3.4), which is the head above the drainage datum used internally as a predictor in the SSM β_3 term; the BACI gap is the dependent variable in the clearfell ANCOVA.

To separate the intervention signal from background climate variability and from the prior scraping event, an ANCOVA-BACI model (10a; Underwood, 1992; Gotelli & Ellison, 2004; Osenberg et al., 1994) was fitted to each monthly BACI gap series:

$$BACI_{gap}=\alpha+\beta_1\cdot CWB+\beta_2\cdot D_{scrape}(t)+\beta_3\cdot D_{fell}(t)+\beta_4\cdot(CWB\times D_{fell})+\beta_5\cdot Easting(t)+\varepsilon$$

where CWB is the centred cumulative water balance anomaly (cumulated P − PET departure from the long-run mean); D_scrape(t) is a distance-weighted scraping covariate applying an exponential decay from CEH36 with length scale $\lambda = 300$ m, activated at the April 2015 scraping date; D_fell(t) is a step function for the December 2017 clearfell; CWB × D_fell is an interaction term testing whether climate sensitivity changed after felling; and Easting(t) is an easting × time interaction term that absorbs the progressive coastal erosion gradient across the well network.

The BACI displacement series (target centroid minus control centroid) contains several overlapping signals: the scraping step, the clearfell step, climate-driven fluctuations (absorbed by CWB), and a slow linear drift caused by coastal erosion progressively deepening the more coastal target wells relative to inland controls. If this drift is not explicitly modelled, the regression distributes the unmodelled trend across whichever coefficients best reduce the residual, biasing all estimates including the clearfell step. The easting × time term is therefore included for any control run in which the combined target-and-control well set spans an easting range exceeding 200 m; this criterion is met for all three control definitions (Forest control easting range 499 m, Climate 1,056 m, Combined 1,056 m). This linear gradient captures average differential deterioration across the observation window but does not represent time-lagged inland propagation of the coastal-erosion signal: at the felling-affected wells, which sit several hundred metres inland of the coast, the characteristic propagation timescale ($L^2 S_y/K$, Section 5.4.3) for coastal-erosion effects to reach the wells is multi-year to multi-decadal, with most of the propagation envelope lying outside the 2018–2026 analysis window. The linear gradient should therefore be read as a first-order spatial control rather than a complete model of the coastal-erosion confound; longer-record analyses or explicit boundary-condition modelling would be required to characterize the time-lagged component. An additional step term for the October 2023 re-scraping was tested in each model and retained where it improved fit ($\Delta AIC < -2$). The scraping distance-decay length scale was fixed at $\lambda = 300$ m on physical grounds (Section 4.5); sensitivity to $\lambda = 200$ m and $\lambda = 500$ m was tested and reported.

The inclusion of the scraping covariate is empirically motivated: a significant level shift in the BACI gap at the April 2015 scraping date was detected independently in the scraping analysis (Section 4.5). The propagation analysis (Section 4.5, 09b) found no detectable coefficient perturbation at uphill wells, indicating that the level shift at WMC3 may reflect a coincident climate signal rather than hydraulic propagation from the scraped site; the covariate is retained to avoid biasing the clearfell estimate regardless of the underlying mechanism.

This three-counterfactual design produces six primary ANCOVA results (three controls × two target tiers), each with its own clearfell step estimate, confidence interval, and model diagnostics.

- **Spatial step-change mapping (10b).** The post-felling step at each reference well — relative to its pre-felling mean and corrected by subtracting the appropriate control centroid step — was computed and rendered as climate-corrected spatial maps for the felling and scraping eras (Figures 22, 23). These maps provide a visual diagnostic of whether the step changes show a spatial gradient consistent with a localized intervention effect.
- **The within-forest spatial coefficient analysis (10c**; Section 3.4.4, results in Section 4.9.4) underpins the C4-only composition of the Forest control: it establishes that elevation alone explains 95% of within-forest β_2 variance, so position-matched controls rather than canopy-matched controls are the appropriate counterfactual for testing canopy management effects.

Four supplementary analyses examined the temporal dimension of the clearfell signal:

- **Summer minima analysis (10d).** The annual summer minimum depth (June–September) was computed for each well in the 17-well network. For each well, the gap between the well's summer minimum and the control centroid summer minimum was computed annually, and the pre- vs post-felling shift in this gap tested by Welch's t-test. Tier-mean shifts were computed for both forest and climate control centroids. A mixed-effects model with random intercept per well and fixed effects for post-felling and post-scraping indicators was fitted to each tier to provide a pooled clearfell step estimate with proper uncertainty (Results are presented in Section 4.6.6):

$$gap = \beta_0 + \beta_1 \cdot D_{fell} + \beta_2 \cdot D_{scrape} + u_{well} + \varepsilon$$

- **Coefficient decomposition (10e).** The SSM was fitted independently to Before (pre-scraping to felling) and After (felling to end of record) windows for each well. Per-well coefficient shifts ($\Delta\beta_1$, $\Delta\beta_2$, $\Delta\beta_3$) were computed, and the predicted clearfell effect estimated as:

  $$\Delta h_{pred} = \Delta\beta_1 \cdot mean_P - \Delta\beta_2 \cdot mean_{PET} - \Delta\beta_3 \cdot mean_{h_{disp}}$$

  where the means are long-run annual averages. Comparison of Δh_predicted with the observed ANCOVA step quantifies the fraction of the clearfell effect attributable to SSM coefficient changes versus other sources. Results are presented in Section 4.6.6.

- **SSM residual normalisation (10f).** The SSM was calibrated on the pre-felling record for each well and run forward iteratively using observed climate forcing. Post-felling residuals (observed minus predicted) were normalised by subtracting the forest control mean residual, providing a climate-independent test of clearfell-specific displacement. Results are presented in Section 4.6.7.

- **Synthetic control (10f).** A synthetic counterfactual for the Impact and Edge tiers was constructed from an optimally weighted combination of six donor wells calibrated on the pre-intervention period. The gap between observed and synthetic post-felling provides an estimate of the intervention effect without relying on the ANCOVA parametric assumptions. Results are presented in Section 4.6.7.

- **A clearfell transect analysis (10g)** tested whether the post-felling response showed a spatial gradient consistent with a localized clearfell effect, comparing era-mean depths across a transect from the plantation interior to the felling core. A genuine clearfell response would produce a gradient strongest at WMC3 and diminishing with distance; a spatially uniform response would indicate a common climate signal. Results are presented in Section 4.6.5.

  Additional diagnostics (10g) included NW10 broadleaf trend analysis (testing whether the broadleaf-canopy well diverges from the pine composite over time), rolling-window SSM coefficient tracking, and drainage diagnostic plots. The NW10 broadleaf comparison is treated as a single-well diagnostic with its limitations discussed in Section 6.3; the rolling-window diagnostics support the temporal stability checks underpinning Section 4.6.7.

- **Synthetic extension of FE wells (10h).** FE1 and FE2, which sit inside or at the immediate edge of the clearfell compartment, were excluded from the primary ANCOVA-BACI because their records begin in July 2015 — too late for a pre-scraping baseline. To test whether these wells confirm the WMC3 result, their records were extended backwards using OLS donor regression. The donors were the three Forest Control wells CEH34, CEH2, and CEH33, selected on the basis of proximity, shared canopy conditions, and pre-clearfell correlation ($r > 0.99$ for both FE wells). The regression was calibrated on the pre-clearfell overlap window (July 2015 – November 2017, $n = 29$ months) and applied to hindcast FE well levels back to August 2010, providing a synthetic pre-scraping baseline.

The synthetic hindcast was spliced with the actual FE observations to create a stable-composition impact centroid spanning all three eras (pre-scraping, post-scraping, post-felling). Three impact centroid variants were tested: (A) WMC3 + FE1 + FE2, (B) WMC3 + FE2, and (C) WMC3 alone. Each was run through the same ANCOVA framework as the primary analysis (Section 3.5.5), against the Forest, Climate, and Combined control definitions. FE1 is caveated as sitting approximately 20 m outside the clearfell boundary in standing forest. Results are presented in Section 4.6.7.

The clearfell BACI experiment was implemented as a modular suite (run_10_clearfell.py: 10a–10h) with shared data loading and spatial functions in clearfell_common.py.

3.6 Threshold Forecasting

To translate the observed spatial dynamics into predictive management tools, operational forecasting equations were derived using a multi-tiered mathematical framework. This approach combines continuous mass-balance modelling with seasonal prediction equations to quantify the aquifer's climatic vulnerability.

3.6.1 Cluster-Level Mechanistic State-Space Equations

The cluster-level SSM coefficients derived in Section 3.4.1 were used to construct operational P_flood threshold equations for each of the five principal clusters. The model was fitted with the intercept strictly forced to zero, ensuring that the coefficients represent a strict physical mass balance and providing the basis for algebraic inversion in Section 3.6.3.

3.6.2 Seasonal Prediction Equations

While the SSM operates on a continuous monthly basis, conservation management typically requires inter-seasonal risk forecasting — specifically, predicting winter flood peaks and summer drought minima from antecedent conditions. To address this, autoregressive prediction equations were fitted at a hydrological-year resolution for the three primary structural domains: Eastern Block, Western Block, and Forest clusters.

Two equations are fitted for each domain. The winter prediction equation predicts the peak winter water table (h_peak) from cumulative winter precipitation and the antecedent summer minimum:

$$h_{peak} = a_1 P_{winter} + b_2 h_{min} + c$$

The summer prediction equation predicts the summer minimum water table (h_min) from cumulative summer precipitation and the antecedent winter peak:

$$h_{min} = a_1 P_{summer} + b_2 h_{maxwinter} + c$$

where P_winter is cumulative precipitation over October–March, P_summer is cumulative precipitation over April–September, h_min is the preceding summer minimum water table depth, h_max_winter is the preceding winter maximum water table depth, a is the fitted regression coefficient on the seasonal precipitation predictor, b is the fitted coefficient on the antecedent-state predictor, and c is a fitted intercept. The seasonal equations are AR(1)-style empirical regressions at hydrological-year resolution, providing operational forecasting at coarser temporal resolution than the monthly SSM described in Section 3.4. Data are organised by hydrological year commencing 1 October.

The seasonal equations are AR(1)-style empirical regressions at hydrological-year resolution, structurally simpler than the monthly SSM. Head-dependent drainage (the $-\beta_3 \cdot h(t-1)$ term in the SSM) is not fitted as a separate parameter at the seasonal scale: with one observation per hydrological year, drainage and the antecedent-state term are aliased, so drainage is absorbed implicitly into b and the fitted intercept c rather than appearing as a separate coefficient.

The antecedent state terms (h_min in the winter equation, h_max in the summer equation) explicitly formalise the memory of the dune aquifer system — the degree to which a winter flood peak is conditioned by the depth of the preceding summer drought, and vice versa. This parameterisation allows managers to construct simple conditional forecasts: given a known or projected winter rainfall total and the current summer water table position, the expected winter flood maximum can be estimated directly from the fitted equation without running the full SSM.

3.6.3 Algebraic Derivation of the Critical Rainfall Threshold (P_flood)

The most operationally significant metric for dune slack management is the critical rainfall threshold (P_flood) — the cumulative winter rainfall depth required to raise the water table from its summer minimum to the elevation of the slack floor (h_target, where depth to water table equals zero).

This threshold was derived by iterating the cluster-level state-space equation (Section 3.4.1) forward from the September summer minimum (h_0) over the months from October to each cluster's historical peak-of-record month. The horizon is therefore cluster-specific and reflects the observed timing of winter peak water table: five months (October–February) for C1 and C2, and six months (October–March) for C3, C4 and C5. Within each monthly step the mass-balance equation is applied with its fitted cluster coefficients(β_1, β_2, β_3 from Section 3.4.1; values given in Table 3) :

$$h_t = \left(1-\beta_3\right) \cdot h_{t-1} + \beta_1 \cdot P_t - \beta_2 \cdot E_t$$

Monthly precipitation is expressed as P_t = λ·P̄_t, where P̄_t is the long-term monthly mean rainfall from the RAF Valley reference climatology and λ is a scaling multiplier representing the departure of a given winter from the long-term mean. Monthly potential evapotranspiration is taken directly from the climatology (E_t = Ē_t) without scaling, reflecting the fact that winter PET varies considerably less year-to-year than winter rainfall and is not meaningfully coupled to it at this latitude. This replaces the winter PET ≈ 0 simplification used in earlier provisional versions of this derivation, which neglected approximately 30% of the winter atmospheric demand at RAF Valley and biased the threshold low at long forecast horizons.

Letting $\alpha = 1 - \beta_3$, iteration of the monthly recurrence over n steps yields the closed form:

$$h_n = h_0 \cdot \alpha^n + \beta_1 \cdot \sum_{i=1}^{n} \alpha^{n-i} \cdot \lambda \cdot \bar{P}_i - \beta_2 \cdot \sum_{i=1}^{n} \alpha^{n-i} \cdot \bar{E}_i$$

Setting h_n = h_target and solving for λ:

$$\lambda = \frac{h_{target} - h_0 \cdot \alpha^n + \beta_2 \cdot S_E}{\beta_1 \cdot S_P}$$

where S_P = $\Sigma_{i=1}^{n}$ α^(n−i)·P̄_i and S_E = $\Sigma_{i=1}^{n}$ α^(n−i)·Ē_i are the drainage-weighted rainfall and PET sums over the horizon. The critical rainfall threshold follows directly:

$$P_{flood} = \lambda \cdot \sum_{i=1}^{n} \bar{P}_i$$

Because the weighting α^(n−i) places greater emphasis on early-winter rainfall than on late-winter rainfall, the formulation encodes the physical fact that October recharge has more months to contribute to the peak than March recharge does, modulated by the cluster-specific drainage coefficient β_3. Three natural classes of well emerge from this framework: those whose required λ is less than unity (P_flood reachable under sub-average winters), those requiring 1.0 ≤ λ < 2.0 (reachable only in wetter-than-average winters of varying rarity), and those for which λ ≥ 2.0 (structurally unreachable under any observed winter in the RAF Valley record, equivalent to annual rainfall totals unprecedented at this site). This classification is applied spatially across the dipwell network in Section 4.

This algebraic inversion provides site managers with a dynamic, site-specific forecasting tool. By entering a measured summer minimum dipwell reading (h_0) into the cluster's equation, a manager obtains both the cumulative winter rainfall (mm) required to raise that well to the slack-floor threshold, and its expression as a multiple of climatological winter rainfall (λ) — a direct measure of ecological flooding feasibility in the subsequent recharge season.

3.6.4 Long-Term Climate Trajectory and Threshold Exceedance Analysis

To assess the long-term trajectory of slack hydrological viability, annual summer minima and winter maxima were extracted for all five clusters from the cluster-centroid time series generated during the clustering analysis. C4 and C5, both forested clusters whose water tables sit below the slack viability thresholds, are included for completeness to characterize the full network trajectory and to support future analysis of forest management effects on seasonal extremes. The hydrological year was defined as 1 October to 30 September, consistent with Section 3.6.2, and a minimum of three valid monthly observations per season was required for a year to be included.

For summer minima, ordinary least squares linear regression was fitted to the annual series for each cluster over the 2004–2026 observation window and extrapolated to 2040, with 95% confidence intervals derived analytically. Summer trajectories were plotted against the eco-hydrological viability thresholds established by Curreli et al. (2013): −0.61 m for wet slack communities (SD15b) and −0.98 m for dry slack communities (SD16).

For winter maxima, no reliable projection method was identified — direct linear regression on annual winter maxima produced R^2 values below 0.03 across all clusters, and transfer function projection via the seasonal models of Section 3.6.2 produced physically implausible trajectories for the eastern clusters due to extrapolation beyond the fitted range. The winter analysis therefore presents the observed record only. Threshold exceedance frequency was calculated for each cluster as the proportion of hydrological years in which the winter maximum exceeded the wet slack flooding threshold of −0.10 m (SD15b) and the dry slack threshold of −0.25 m (SD16), providing a direct measure of current ecological flooding adequacy for each structural domain.

Critical rainfall threshold equations, peak flood prediction equations, and summer drought prediction equations were implemented in 11_forecasting_thresholds.py. Climate trajectory projections and threshold exceedance analysis were implemented in 14_climate_projections.py.

3.7 Characterising Water Table Losses and Long-Term Trajectories

3.7.1 Depth-Dependent Evapotranspiration Model

To test whether the atmospheric draw coefficient varies significantly with water table depth — and whether a depth-dependent formulation is warranted over the standard fixed-β_2 SSM — an extended model was implemented in which β_2 was replaced by the depth-dependent term $\beta_2 \cdot \exp(-\lambda \cdot d)$, where d is the depth of the water table below ground surface (metres) and λ is a decay parameter (m^{-1}) controlling the rate at which evapotranspiration influence declines as the capillary fringe retreats from the root zone. Depth below ground surface was computed at each timestep as $d = -h + \bar{u}$, where h is the water table position (negative convention) and $\bar{u}$ is the cluster mean upstand height. Because the modified predictor $-\exp(-\lambda \cdot d) \cdot \mathrm{PET}$ remains linear in β_2 for any fixed λ, the three coefficients β_1, β_2 and β_3 were recovered by no-intercept OLS at each value of λ. A systematic grid search over $\lambda \in [0, 6]\ m^{-1}$ at intervals of $0.05\ m^{-1}$ was conducted independently for each cluster, with the optimal λ selected by maximising iterative simulation NSE. Setting $\lambda = 0$ recovers the standard SSM exactly, providing a direct baseline for comparison.

Despite the statistical improvement that the grid search produces (Section 4.4), the depth-dependent formulation introduces a nonlinearity that would propagate through the threshold equations (Section 4.7), forestry scenarios (Section 4.9), and forecasting tools. The standard SSM's fixed β_2 provides an effective depth-averaged parameterisation that preserves the closed-form analytical framework on which these operational outputs depend. The depth-dependent formulation was therefore not adopted for the primary analysis; the standard fixed-β_2 SSM is retained throughout. The fixed β_2 should be understood as absorbing a real depth-coupling effect whose magnitude varies across clusters; its diagnostic significance is discussed further in Section 5.6.1.

The depth-dependent PET grid search and benchmark comparison were implemented in 15_depth_dependent_pet.py.*

3.7.2 Mean Monthly Water Balance Decomposition

Water balance decomposition of the kind applied here — partitioning monthly head change into recharge, evaporative draw, drainage and residual — has a long history in the coastal dune hydrology literature, particularly in the Dutch and North Sea coastal systems where managed water tables and dune slack conservation drive analogous restoration questions (Stuyfzand, 1993; Grootjans et al., 2002, 2017). The decomposition applied in this study derives the component fluxes directly from the fitted SSM coefficients, allowing direct comparison of fluxes across clusters subject to identical climate forcing using only the dipwell and climate records already available for the site.

To quantify and compare the hydrological inputs and losses operating on each hydrogeological cluster, a mean monthly water balance decomposition was applied to the five principal clusters (C1–C5) using cluster-level SSM coefficients derived in Section 3.4.1. All terms are expressed in metres of water table change per month (m/month), the natural unit of the state-space model and directly comparable across clusters receiving the same climate forcing.

The decomposition isolates three components from the state-space equation $\Delta h = \beta_1 \cdot P - \beta_2 \cdot PET - \beta_3 \cdot h\_prev$. Mean monthly recharge was computed as $\beta_1 \times \bar{P}$, where $\bar{P}$ is the mean monthly precipitation over the monitoring period from the RAF Valley record. Atmospheric draw was computed as $\beta_2 \times \overline{PET}$, where $\overline{PET}$ is the mean monthly Thornthwaite PET. Gravity drainage was computed as $\beta_3 \times \bar{h}\_disp$, where $\bar{h}\_disp$ is the mean displacement above the drainage datum ($\bar{h}\_disp = D + \bar{h}$, with $D = 3.7$ m; see Section 3.4) for each cluster derived from the cluster-average hydrograph. Canopy interception is not included as a separate loss term — it is implicitly captured in the forested clusters' (C4 and C5) lower β_1 and lower β_2, and adding it explicitly would double-count losses already embedded in the fitted coefficients. Total losses are therefore the sum of atmospheric draw and drainage only. The net monthly balance — recharge minus total losses — quantifies the water balance residual required from ridge-derived lateral recharge to maintain long-run water table equilibrium or errors.

An indicative volumetric conversion is presented alongside the primary head-space decomposition (Table 4b; Figure 9a), using assumed specific yield values: C1 = 0.08 (shallow sand over glacial till) and C2–C5 = 0.12 (deeper clean dune sand), consistent with the literature range for unconfined sandy aquifer materials (Fetter, 2001). Empirical upper bounds on Sy derived from the WTF method are presented in Section 3.7.3 and Table 4c, and used to bracket the volumetric estimates from above; the partition uncertainty range between the two methods is shown as a hatched band in Figure 9b. For the forested clusters (C4 and C5), canopy interception is shown as an explicit segment on the inputs bar of Figure 9a, computed as $0.24 \times \bar{P}$ following Freeman (2008), reducing the boundary subsidy rather than being additive.

The mean monthly water balance decomposition and volumetric conversion were produced by 16_water_bal.py (outputs: 16_water_bal_table.csv, 16_water_bal_vol_table.csv).

3.7.3 Specific Yield Estimation: Water Table Fluctuation Method

To provide an empirical upper bound on the specific yield (Sy) values used in the volumetric water balance conversion (Section 3.7.2; Table 4b), an independent estimation was performed using the water table fluctuation (WTF) method of Healy and Cook (2002):

$$Sy = R / \Delta h$$

where R is the net recharge reaching the water table (m) and Δh is the corresponding water table rise (m). The method was applied to the cluster-mean hydrographs derived in Section 3.4.1, using monthly net recharge computed as $P - PET$. Months were retained where both the water table rose ($\Delta h > 5$ mm) and net recharge was positive ($P - PET > 10$ mm); event-level estimates were further constrained to the physically plausible range $0.01 < Sy < 0.50$ to exclude months where the suppressed or negative Δh signal produced implausibly high individual Sy values. Monthly data cannot isolate individual storm events, and the WTF method at monthly resolution is known to overestimate effective Sy because it cannot separate true gravity drainage from capillary fringe release (Healy and Cook, 2002; Scanlon et al., 2002). The resulting estimates should therefore be interpreted as upper bounds on Sy rather than point estimates of the long-term effective value.

For the Forest clusters (C4 and C5), a correction was applied for canopy interception prior to computing net recharge:

$$R\_effective = (1 - 0.24) \times P - PET$$

following the site-specific interception fraction of Freeman (2008). Only the precipitation term is reduced. The PET term is retained unchanged because Thornthwaite PET is an energy-based atmospheric demand derived from temperature and day length, independent of land cover; it represents the capacity of the atmosphere to receive water vapour rather than the realised evaporative flux from any particular surface. Reducing PET to account for canopy-evaporated intercepted water would conflate atmospheric demand with actual evapotranspiration and remove the same quantity from both sides of the water balance. Without the interception correction, months in which canopy suppression of Δh is most pronounced produce individual Sy estimates above the plausibility threshold and are excluded from the event pool; applying the correction recovers these months (increasing the event count from 50 to 61 at C4 and from 34 to 50 at C5) and lowers the event-median Sy from 0.317 to 0.243 at C4 and from 0.360 to 0.320 at C5, as the recovered months carry below-median Sy values (see Section 4.2.4 and Table 4c).

Results of the cluster-level and well-level WTF analyses are presented in Section 4.2.4. *The cluster-level analysis was implemented in 17_wtf_specific_yield.py and the well-level spatial analysis in 18_wtf_spatial.py.*

3.7.4 Network-Scale Coastal-Retreat Gradient Regression

To partition each cluster's observed summer-minimum decline into a climate-driven component and a coastal-retreat-driven component, and to provide an independent check on the BACI ANCOVA's easting × time covariate (Section 3.5.5), a non-linear panel regression was fitted to the full 21-year monthly record. Each well's perpendicular distance to the eroding Caernarfon Bay High Water Mark was computed once out-of-pipeline from the OS Open Map Local TidalBoundary shapefile, restricted to the 15 km west-facing Newborough shoreline; the Menai Strait and Llanddwyn Island were excluded as non-eroding boundaries. Per-well distances are supplied to the pipeline as a versioned data input (data/well_distance_to_coast.csv) with full provenance documented in data/COASTLINE_PROVENANCE.md.

Two functional forms were fitted: a Dupuit–Forchheimer steady-state strip-aquifer form, $\delta(d) = \max(\delta\_0 \cdot (1 - d/L), 0) + c$, and an exponential decay form, $\delta(d) = \delta\_0 \cdot \exp(-d/L) + c$, where d is distance from the coast, δ_0 is the coast-edge water-table deepening rate above climate, L is the inland reach, and c is the far-field climate background. Both were fitted by profile non-linear least squares, with well and month fixed effects and a cumulative-water-balance covariate absorbed by within-well demeaning. Model selection between the two functional forms used the Akaike Information Criterion. Three nested specifications were fitted to test robustness to forest-cover inclusion: the full network (72 wells, all clusters minus the BACI clearfell-zone wells), a forest-free specification dropping all C4 and C5 wells, and a C3-only specification using the 17 non-forested C3 wells alone. The forest-free linear-with-cutoff fit is taken as the headline; full methodological detail including identifiability arguments and z-test consistency checks between specifications is provided in Supplementary Material.

The headline fit is then applied to two downstream questions: per-cluster decomposition of the Script 14 summer-minimum slope into gradient and climate components (Section 4.8.1), and corroboration of the BACI ANCOVA's easting × time coefficient against the model's predicted differential coastal-retreat signal between each impact zone and each control tier (Section 4.6.3).

The analysis was implemented in 25_coastal_gradient.py

3.8 Computational Methods and Software Implementation

Data Pre-processing and Time-Series Management All data pre-processing, statistical modelling, and spatial mapping were executed within a Python programming environment. Time-series data manipulation, interpolation of missing temporal records, and standardisation were performed utilizing the pandas and NumPy libraries.

Statistical and Predictive Modelling The core parameterization of the aquifer mechanics was executed using the *statsmodels* module. For the continuous state-space modelling (Section 3.4), the Ordinary Least Squares (OLS) algorithm was constrained to fit through the origin (omitting the intercept constant) to enforce physical mass balance. Conversely, seasonal prediction equations predicting peak flood and summer drought minimums used unconstrained OLS regressions. The *statsmodels* architecture was further used to extract 95% confidence intervals and p-values for the drainage coefficients ($-\beta_3$) to statistically verify the impact of the clearfelling intervention. Iterative forecasting metrics, including the Nash-Sutcliffe Efficiency (NSE) and Root Mean Square Error (RMSE), were computed via *NumPy* array operations to benchmark the state-space model against traditional regression approaches.

Geospatial Processing and Cartographic Visualization The integration of the topographical data with the monitoring network was managed through a dedicated spatial stack. *GeoPandas* and *Fiona* were used to parse external KML feature boundaries (e.g., the clearfell perimeter and hydrological boundaries) and standardize all spatial geometries to the British National Grid coordinate reference system (EPSG:27700).

The high-resolution Digital Elevation Model (DEM) was ingested and masked using the *rasterio* library. To visualize the localized dune topography effectively, the DEM arrays were rendered using a dual-slope normalization technique (`TwoSlopeNorm`) provided by *matplotlib*, anchoring the visual transition at specific elevation thresholds to highlight critical hydrological features. Contextual basemaps were fetched via *contextily* when necessary, and the label placement and occlusion avoidance was handled using the adjustText library.

Spatial Groundwater Analysis. A well-level scenario framework was implemented to translate the per-well SSM parameterisation into site-wide spatial summaries of head change (Δh) under prescribed climate and forest management perturbations. For each well, the scenario head change relative to baseline was computed directly from the per-well SSM equation as:-

$$\Delta h = \left(\beta_1 \cdot P_{eff,sc} - \beta_{2,sc} \cdot PET_{sc} - \beta_3 \cdot |h|\right) - \left(\beta_1 \cdot P_{eff,0} - \beta_2 \cdot PET_0 - \beta_3 \cdot |h|\right)$$

where P_eff applies the relevant canopy interception factor at forest wells (C4 and C5; 24% under the Corsican pine baseline, following Freeman, 2008; 15% annual mean under broadleaf conversion following Komatsu et al. (2011); 12% under 50% thinning; 0% under full clearfell), β_2,sc applies a scenario-specific multiplier at forest wells (×1.108 for full clearfell, derived dynamically from the BACI-corrected Edge-tier post-felling β_2 ratio reported in Section 4.6.6, Table 11 — Edge mean ratio 1.101 minus Climate Control mean drift 0.993, plus 1.0). The Edge tier is used as a conservative upper bound: although Edge wells were not directly felled, they recorded the largest tier-mean β_2 increase of any near-felling tier, while the directly-felled Impact well (WMC3) showed a smaller ratio of ×1.041 not statistically distinguishable from the Climate Control drift. Adopting the Edge ratio therefore yields a precautionary upper bound on canopy-removal-attributable β_2 change. The thinning multiplier of ×1.054 is set as half the clearfell perturbation; β_2 is held at baseline under broadleaf conversion, with no annual-mean justification for a year-round β_2 increase under deciduous canopy. The same clearfell multiplier is applied at both C4 and C5 forest wells; no felled C5 wells were available in the BACI design, so the C5 application represents extrapolation from the C4 near-felling signal rather than direct measurement. The $\beta_3 \cdot |h|$ term cancels between the scenario and baseline expressions and does not enter the Δh calculation. Sy 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 specific yield values (Section 3.7.3) are retained as supplementary display information in the interactive scenario viewer, not as a divisor on Δh.

Six scenarios were evaluated: baseline; UKCP18 2050s (winter P +10%, summer P −15%, winter PET +5%, summer PET +20%, RCP8.5 central estimate for Wales, informed by the UKCP18 Regional 12 km ensemble of Met Office (2018) and the CHESS-SCAPE bias-corrected projections of Robinson et al. (2023)); UKCP18 2080s (winter P +20%, summer P −30%, winter PET +10%, summer PET +35%, same sources); full clearfell; 50% thinning; and broadleaf conversion. Climate scenarios were evaluated per season, with the annual Δh reported as the 0.5-weighted mean of the winter and summer equilibrium responses. The per-well Δh field was interpolated to a site-wide 50 m grid on the British National Grid (EPSG:27700) using 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 33 and 34. Two wells were excluded from the interpolation: CEH12 (bedrock-anchored, geologically distinct from the dune aquifer) and CEH15 (forest slack edge, transitional between Forest and Western Residual clusters). The hydraulic conductivity reported for the site (K = 6.0 m/day; Betson et al., 2002) is retained as directional context for the Darcy flow vectors shown in Figure 44 and does not enter the Δh calculation. The accompanying interactive scenario viewer (Hollingham, 2026b; scenario_viewer.html) renders the same per-well Δh field using inverse-distance-weighted interpolation (power 1, eight nearest neighbours) for responsiveness under slider-driven re-interpolation, and offers a complementary depth-below-surface view obtained by subtracting the IDW-interpolated head field from the high-resolution DEM on a per-cell basis and masking inter-dune ridge cells (where the DEM exceeds the well-plane IDW elevation by more than 1.0 m) to confine the visualisation to hydrologically meaningful slack positions; the viewer and the static figures therefore report matching per-well Δh values but use distinct interpolation methods and present complementary quantities at the grid-rendering stage.

A two-dimensional finite-difference PDE model was initially implemented during development as an alternative spatial framework. Using the Helmholtz form of the steady-state groundwater flow equation (Bear, 1972; Freeze and Cherry, 1979) :

$$\nabla\cdot(T\cdot\nabla h) - \beta_3\cdot h = -(\beta_1\cdot P\_eff - \beta_2\cdot PET)$$

with transmissivity $T = K\cdot b$ ($K = 6.0$ m/day after Betson et al. (2002); b IDW-interpolated from four borehole constraints after Bristow, 2003), Dirichlet head conditions at sea boundaries, and implicit Neumann no-flow at the ridge, the sparse linear system of approximately 4,200 equations was solved using scipy.sparse.linalg.spsolve (Virtanen et al., 2020). The solver produced physically plausible baseline head surfaces and Darcy flux fields, but forest management scenarios evaluated through the PDE produced near-zero site-wide differences (mean Δh = −0.003 to −0.008 m for full clearfell and thinning). This reflects three structural limitations: the β coefficient changes affect only the forest wells (C4 and C5), which when IDW-interpolated to the grid are diluted by the surrounding clusters; the weak drainage feedback at C4 ($\beta_3 = 0.021$, the lowest cluster centroid value in the network) means a moderate source term change produces only a small equilibrium head change; and lateral recharge from the ridge is not parameterised in the source terms. Combined with the sparse constraints on K (a single tracer test) and b (four borehole logs across 700 ha), the PDE approach was judged to overclaim relative to the data available — echoing the calibration difficulties that constrained Betson et al. (2002) MODFLOW model at this site. The well-level IDW framework described above was therefore adopted as the primary spatial analysis, on the basis that it retains the physical content that the SSM parameterisation supports without claiming spatial flow dynamics the monitoring data cannot constrain. All scenario outputs are provided through a self-contained interactive viewer allowing site-wide and cluster-level Δh responses to be explored under any combination of the five defined scenarios and their underlying parameter sliders (scenario_viewer.html; available at https://newbroman.github.io/Newborough_Hydrology/scenario_viewer.html).

Spatial groundwater analysis and scenario mapping were implemented in 19_spatial_groundwater.py. Publication-quality spatial figures were produced by 20_spatial_figures.py.

3.9 Data requirements and replicability

A deliberate objective of this study was to demonstrate the hydrological insight achievable from minimal, low-cost monitoring infrastructure. All primary data inputs consist of manual dipwell measurements, recorded with a standard electrical contact meter or tape measure, and publicly available monthly climate summaries from the Met Office historic station dataset. No specialist survey equipment, telemetry, or laboratory analysis was required. Potential evapotranspiration was estimated using the Thornthwaite and Mather (1957) method specifically because it requires only mean monthly air temperature — a variable recorded at every Met Office station and at the majority of automatic weather stations worldwide. Methods requiring wind speed, humidity, or radiation data were deliberately avoided as these would limit replicability to sites with more comprehensive climate instrumentation.

The analytical framework presented here could be deployed at any comparable lowland wetland, dune slack, or shallow unconfined aquifer system with a dipwell network and access to a nearby climate station. The principal constraint on analysis quality is record length rather than measurement sophistication; the state-space modelling and BACI analyses both benefit substantially from the long baseline available here, but meaningful cluster separation and mechanistic coefficient estimation are likely achievable with records of eight to ten years or more.

4. Results

4.1 Climate and Well Network Characterisation

4.1.1 Climate Record

The RAF Valley climate station provides a continuous monthly record from December 1930 to February 2026 (Figure 3a) (Section 3.1.2). Over the well-record period (April 2005 to February 2026), mean annual precipitation was 890 mm and mean annual Thornthwaite PET was 652 mm, giving a mean P/PET ratio of 1.36 (Table 1). The wettest year on record within this period was 2008 (1,040 mm), followed by 2012 (1,022 mm); the driest complete year was 2010 (736 mm). Four years recorded a P/PET ratio below 1.2 — indicating near-balance between precipitation and evaporative demand — namely 2007, 2010, 2011 and 2022, with 2022 recording the lowest ratio of the monitoring period at 1.12. There is no meaningful difference in mean annual precipitation between the pre-felling (2006–2017: 887 mm) and post-felling (2018–2025: 893 mm) periods, confirming that the post-2018 hydrological signal detected in the BACI analysis (Section 4.6) is not driven by a concurrent change in rainfall amount.

PET shows strong seasonality but negligible inter-annual trend across the record, with the 12-month rolling mean remaining stable at approximately 50–55 mm/month throughout (Figure 3b). Consequently, inter-annual variability in the water balance is driven almost entirely by precipitation rather than evaporative demand.

The levelled cumulative water balance series (Figure 3c), expressed relative to the study-period mean of 19.46 mm/month, oscillates with a clear seasonal signal — rising sharply each winter as precipitation exceeds PET and falling each summer — while tracking the multi-year rhythm of wetter and drier phases. No sustained multi-year drift away from the baseline is evident, confirming that the site maintains a consistent long-run surplus throughout the monitoring period. The water table clearly follows this movement (Figure 3d). Linear regression of the cumulative balance against the unsmoothed network mean water level at zero lag yielded $R^2 = 0.515$, confirming that approximately half the variance in network mean water table depth is explained by antecedent climate forcing alone. The residual variance reflects the spatial heterogeneity of the network, local boundary subsidies, and land cover effects explored in subsequent sections. This close correspondence between cumulative climate forcing and water table behaviour provided the empirical motivation for the state-space modelling framework described in Section 3.4.

The seasonal structure of the climate record has direct implications for aquifer behaviour. Mean monthly rainfall peaks in October–November (around 96–100 mm/month) and falls to its minimum in April–May (around 50–52 mm/month; Figure 3a, whilst Thornthwaite (1957) PET peaks in July–August (around 100–120 mm/month) and is negligible from November to February (Figure 3b). The seasonal opposition between peak PET and minimum rainfall concentrates atmospheric evaporative demand on the aquifer during the period of lowest recharge input, driving summer water table drawdown that must be overcome by the following winter's recharge before any slack flooding can occur. This seasonal asymmetry — not the annual rainfall total — is the primary control on dune slack hydrology and is reflected throughout the subsequent analyses. Over the monitoring period, mean annual winter rainfall (October–March) was 521 mm and mean annual summer rainfall (April–September) was 369 mm, together summing to the 890 mm annual mean.

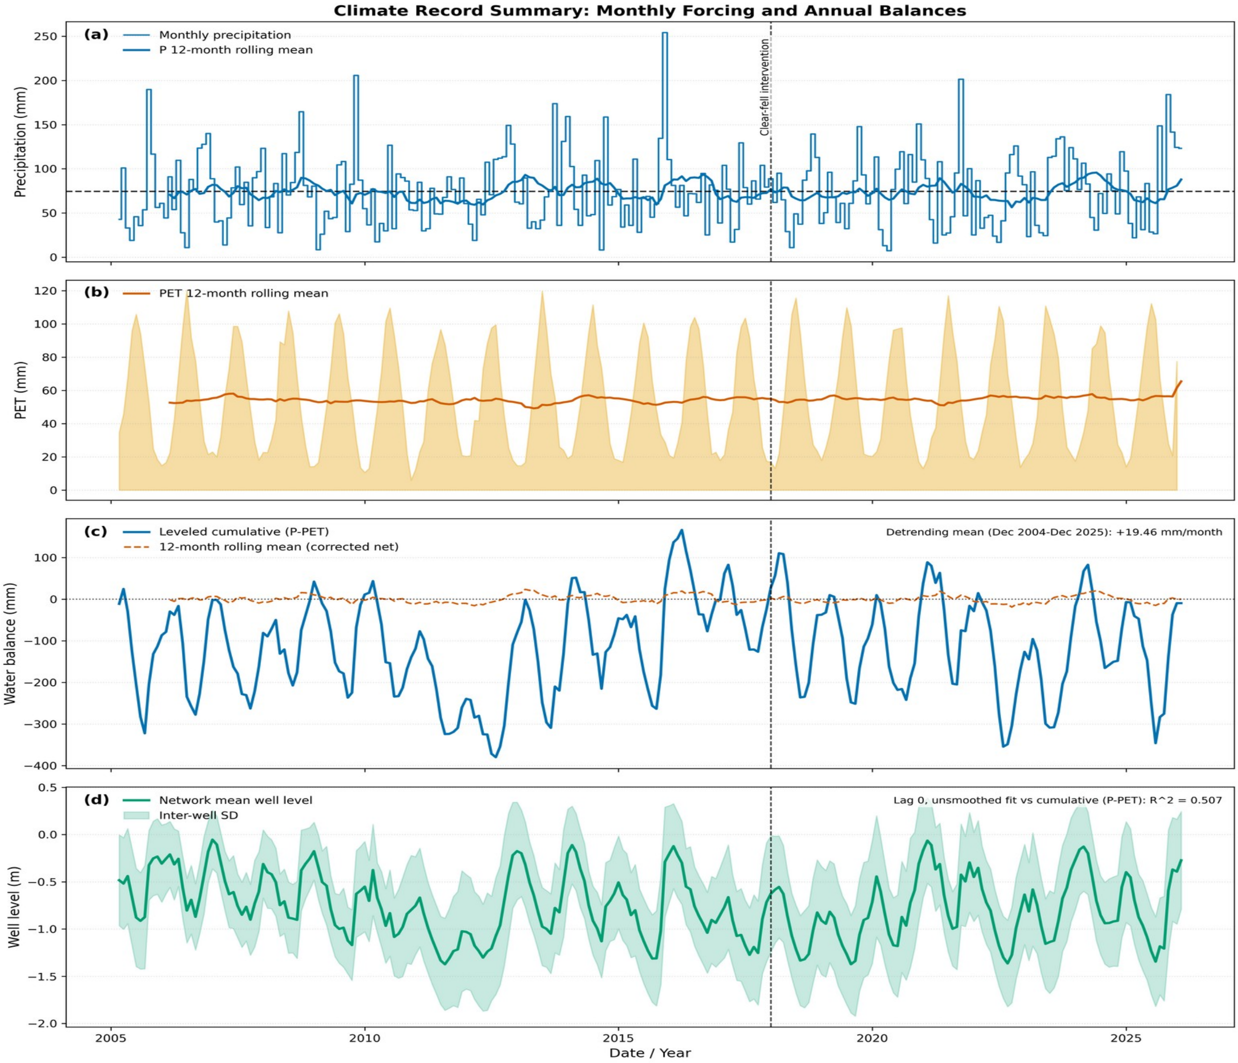

Figure 3. *Climate record summary for the well-monitoring period (April 2005 – February 2026), derived from RAF Valley Met Office station. Panel (a): monthly precipitation (mm) with long-term monthly mean indicated by dashed line and 12-month rolling mean as solid line. Panel (b): monthly Thornthwaite PET (mm) as filled area with 12-month rolling mean overlaid. Panel (c): levelled cumulative water balance (P − PET, mm), corrected by subtracting the long-term mean monthly net balance to centre the series on zero; the dashed orange line shows the 12-month rolling mean of the corrected net balance. The intercept correction value is shown inset. Panel (d): network mean monthly water level (m below pipe top) across the 66-well reference network, with inter-well standard deviation as shaded envelope; the R² inset gives the coefficient of determination from linear regression of the cumulative balance against the unsmoothed network mean at zero lag. Vertical dashed line marks the December 2017 clearfell intervention in all panels*. *(00_01_climate_timeseries_short.png)*

Year	Annual P mm	Annual PET mm	Months complete	P PET ratio
2005	650.3	577.3	9	—
2006	954	655.4	12	1.46
2007	773.7	653.7	12	1.18
2008	1039.7	648.7	12	1.6
2009	871.3	647	12	1.35
2010	736.2	628.8	12	1.17
2011	760.6	647.5	12	1.17
2012	1021.7	637	12	1.6
2013	844	633.6	12	1.33
2014	943	662.5	12	1.42
2015	974	631.2	12	1.54
2016	858.4	652	12	1.32
2017	871	659.4	12	1.32
2018	873.2	653	12	1.34
2019	832	654.4	12	1.27
2020	944.2	656.6	12	1.44
2021	906.6	649	12	1.4
2022	747.2	668.2	12	1.12
2023	986.2	672.9	12	1.47
2024	911.2	657.8	12	1.39
2025	943	675.5	12	1.4
2026	247.2	77.4	1	—
Long-term mean	890.6	652.2		1.36

Table 1. *Annual climate summary for RAF Valley station over the well-monitoring period. Annual precipitation (P) and Thornthwaite PET are in mm. P/PET ratio below 1.0 indicates evaporative demand exceeds precipitation. Months_complete gives the number of months with valid data in that year. Partial years (2005: April–December only; 2026: January–February only) are included for completeness but no annual P/PET ratio is calculated for these years (shown as —). Long-term mean row is calculated from complete years 2006–2025 only. (00_01_annual_climate_summary_short.csv).*

The RAF Valley record also shows a statistically significant warming trend in mean summer maximum temperature over the full 95-year record (slope +0.014°C yr^{-1}, $R^2 = 0.142$, $p < 0.001$; Figure 4). The post-2013 period has run 0.94°C above the pre-2013 baseline mean of 17.89°C, with six of the thirteen summers between 2013 and 2025 exceeding this baseline by more than 1°C (2013, 2018, 2021, 2022, 2023 and 2025). The most extreme anomalies occurred in 2022 (+1.77°C) and 2023 (+1.64°C). The year 2022 combined a +1.77°C temperature anomaly with only 747 mm annual rainfall — the most severe dual-stress year in the monitoring period. Annual rainfall shows no significant trend over the full record ($p = 0.498$).

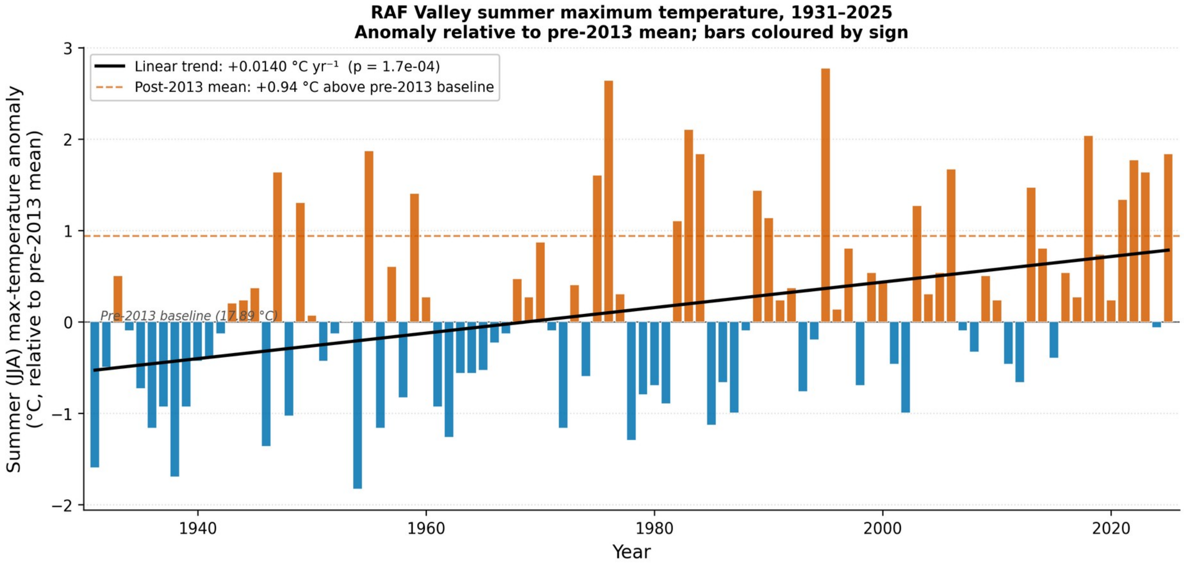

Figure 4. *Summer maximum temperature anomaly at RAF Valley, 1931–2025, expressed as departure from the pre-2013 mean (17.89°C). Bars coloured by sign: orange = above baseline, blue = below. Solid black line: OLS linear trend (+0.014°C yr^{-1}, $p < 0.001$). Dashed orange line: post-2013 mean anomaly (+0.94°C above baseline). (00_03_summer_warming_trend.png)*

4.1.2 Well Network

The 66-well reference network (Section 3.1.1) spans a median record length of 150 months (12.5 years), with individual records ranging from 100 to 201 months (Figure 5, upper left). The network mean water level across all wells and months was 0.74 m below pipe top, with individual well means ranging from 0.04 m above pipe top to 2.07 m below, reflecting the considerable topographic and hydrogeological diversity of the site (Figure 5, upper right). The median seasonal amplitude — defined as the difference between mean August and mean February water levels — was 0.59 m, with the water table deepest in late summer and shallowest in late winter, consistent with the expected seasonal recharge cycle for a maritime dune system (Figure 5, lower right). The network mean time series (Figure 5, lower left) shows a clear seasonal cycle throughout the record, with the inter-well standard deviation envelope widening visibly after the December 2017 clearfell intervention, reflecting the divergence of the impact zone wells from the broader network that is explored in detail in Section 4.6.

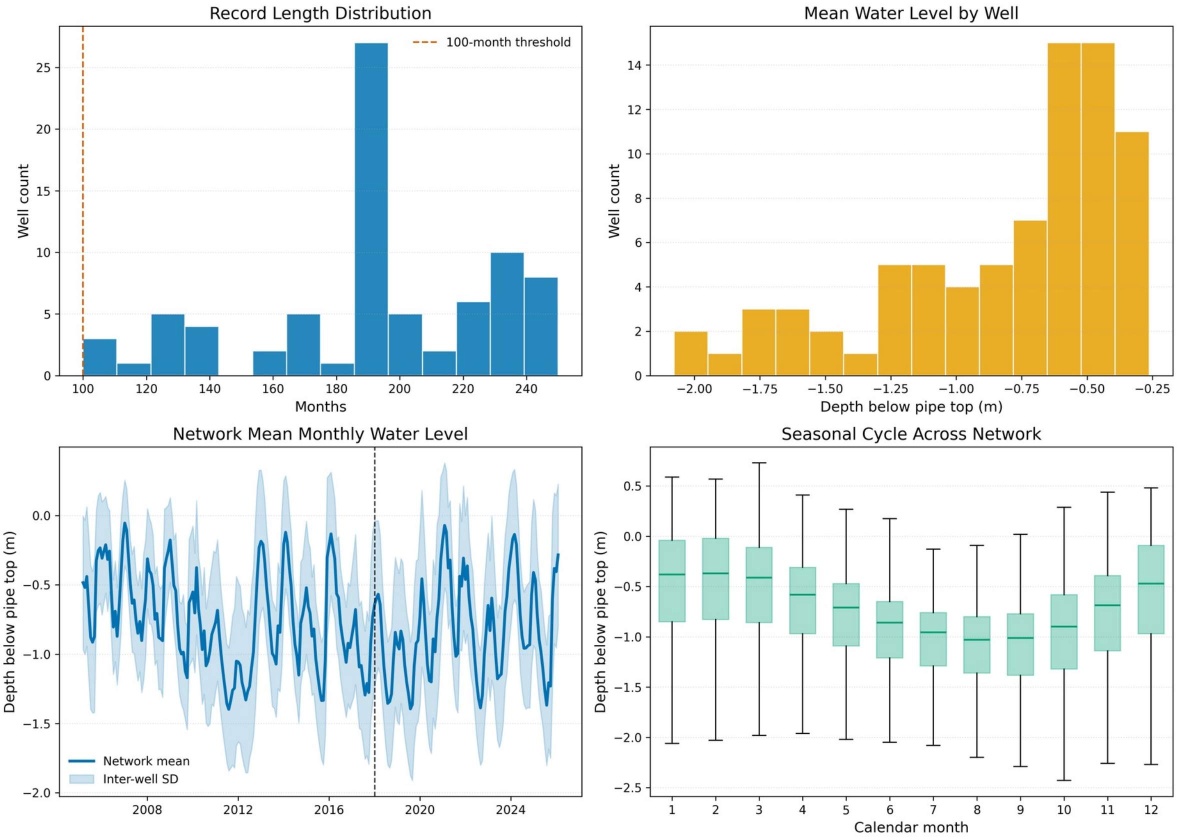

Figure 5. *Well network characterisation for the 66-well reference network. Upper left: frequency distribution of record lengths (months), with the 100-month minimum threshold indicated by dashed line. Upper right: frequency distribution of mean water level by well (m below pipe top; negative values indicate water standing above pipe top). Lower left: network mean monthly water level (m below pipe top) with inter-well standard deviation as shaded envelope; vertical dashed line marks the December 2017 clearfell intervention. Lower right: network-wide seasonal cycle shown as monthly box plots pooling all well-month observations; boxes show interquartile range, horizontal line shows median, whiskers extend to 1.5 × IQR, outliers suppressed. (00_02_well_network_summary_short.png)*

Mapped against the eco-hydrological thresholds of Curreli et al. (2013), the network-wide seasonal extremes reveal a system already under considerable stress. The mean annual summer minimum across the reference network was −1.11 m below pipe top, ranging from −0.33 m to −2.28 m across individual wells. Only one well maintains a mean summer minimum shallower than the wet slack viability threshold of −0.61 m; 35 wells fall within the dry slack zone (−0.61 to −0.98 m); and 37 wells record mean summer minima deeper than the dry slack threshold of −0.98 m, of which 21 lie below the SD16 recovery limit of −1.20 m — placing them beyond the reach of a single scraping event under average conditions. For winter maxima, 22 wells exceed the SD15b wet slack flooding threshold of −0.10 m in an average year and 38 exceed the SD16 dry slack threshold of −0.25 m, though only 11 record mean winter maxima above pipe top. These network-level baselines provide the reference against which cluster-specific dynamics, intervention effects, and climate trajectory are assessed in subsequent sections.

4.2 Network Clustering and Spatial Architecture

The clustering analysis (Section 3.2) identified k = 5 as the preferred cluster solution (silhouette coefficient = 0.39; bootstrap stability analysis confirming four of five clusters with median per-well co-assignment ≥ 0.93). The five clusters are spatially coherent and hydrogeologically interpretable (Table 3; Figures 6–8); their mechanistic characterisation including LCSC values is presented in Section 4.2.2.

The dendrogram confirms a primary binary division separating eastern wells (C1, C2) from western and forest wells (C3, C4, C5) at a Ward linkage distance of approximately 1.0. This top-level split reflects the division between the till-and-estuary-influenced eastern dune aquifer and the deeper, more uniform western sand block. The forest clusters (C4 Main Forest, C5 Coastal Forest) separate from adjacent C3 wells despite geographical proximity, confirming that plantation land cover modifies groundwater behaviour independently of the broader dune aquifer geology. The separation also reflects the position of the C4 wells in the northern part of the study area, closest to the rock ridge, where the water table receives lateral recharge descending from the ridge. C5 (Coastal Forest) occupies the coastal edge of the plantation where canopy and coastal influences combine.

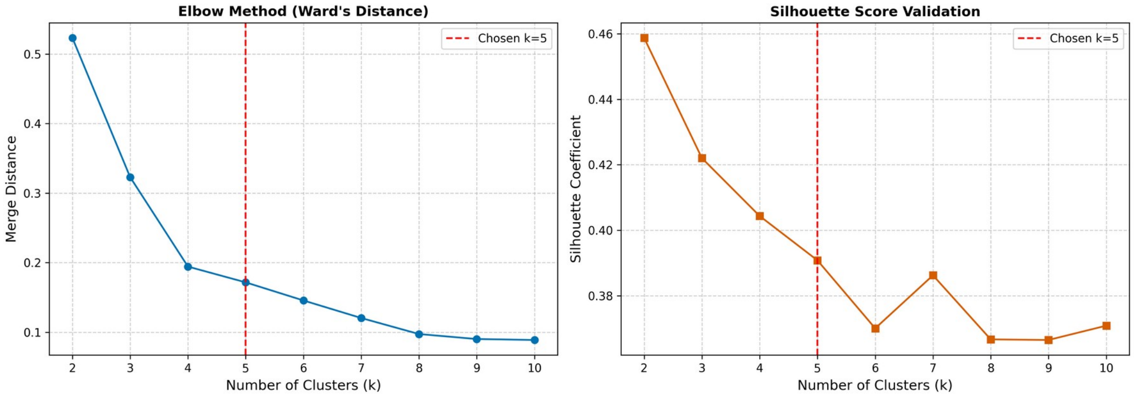

Figure 6. *Cluster validation plots: silhouette coefficient, Calinski–Harabasz index, and Ward's merge distance for k = 2–10, with chosen k = 5 indicated. The silhouette coefficient peaks at k = 2 reflecting the fundamental binary contrast between eastern and western aquifer domains; k = 5 was selected as the finest partition at which bootstrap stability analysis confirms robust, reproducible cluster assignments. (02_02_validation_plots.png).*

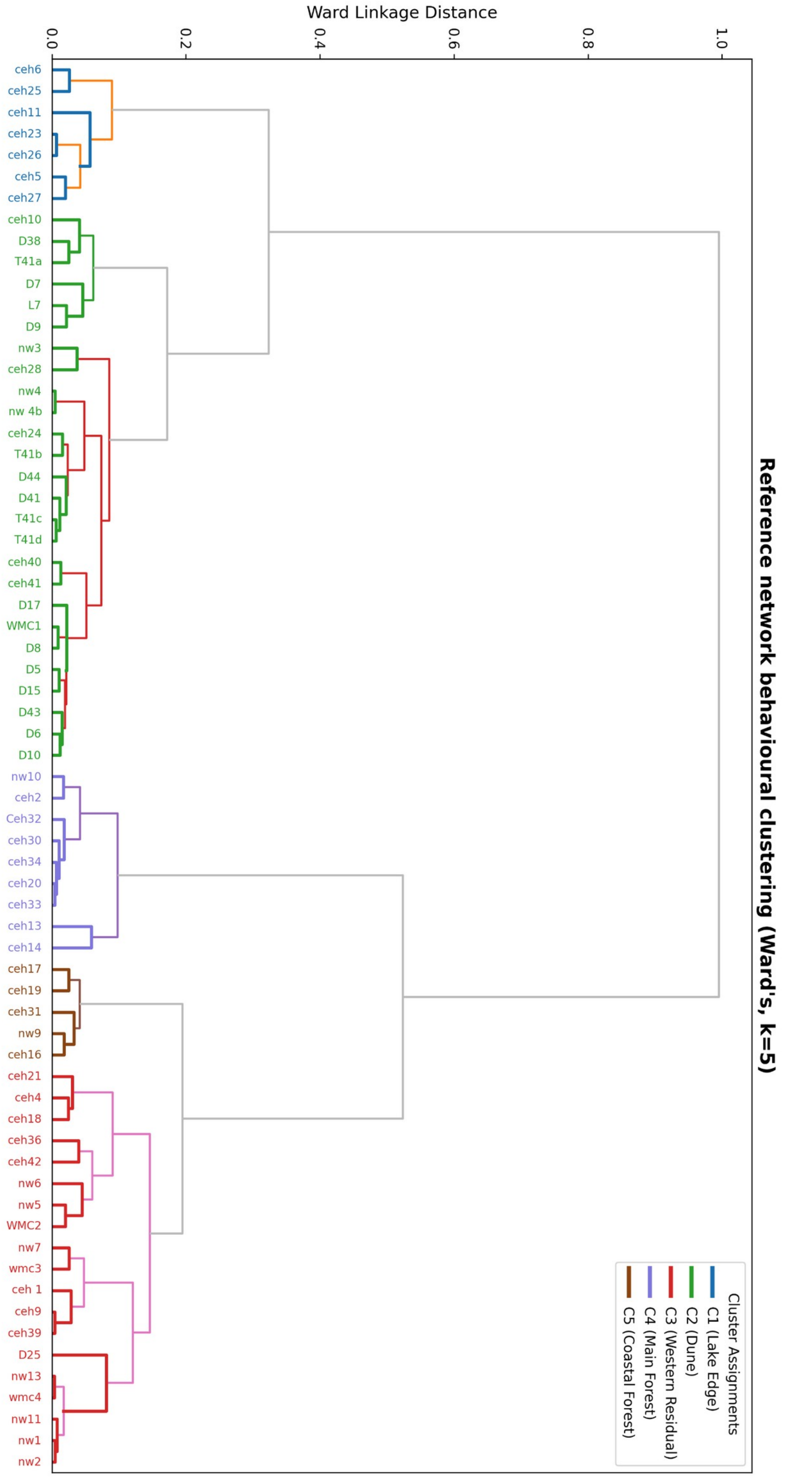

Figure 7. *Clustering results: Ward's hierarchical clustering dendrogram for the 66-well reference network. Wells are labelled by identifier and coloured by assigned cluster. The primary binary split at a Ward linkage distance of approximately 1.0 separates the eastern till-influenced clusters (C1, C2) from the western sand and forest clusters (C3, C4, C5). The Forest cluster (C4) separates from the adjacent Western Residual (C3) at a lower linkage distance, confirming that plantation land cover imposes a distinct hydrogeological signature independent of the shared sandy substrate. (02_01_dendrogram.png).*

4.2.1 Cluster Hydrographs and Seasonal Amplitudes

The five cluster average hydrographs (Figure 8b) reveal a clear stratification of groundwater behaviour that reflects the hydrogeological architecture identified by the clustering analysis. C1 (Lake Edge) shows the shallowest mean depth (−0.45 m), consistent with its low storage capacity and hydraulic connectivity to Llyn Rhos-Ddu. C2 (Dune) tracks C1 closely but sits systematically deeper (mean −0.60 m), reflecting greater distance from the lake boundary. The fundamental geological contrast between the Eastern and Western Blocks is visible as a persistent vertical offset between C1/C2 and C3 (Western Residual), with C3 running approximately 0.3 m deeper on average (mean −0.89 m) despite occupying a broadly similar topographic position — a contrast that reflects the deeper sand accumulation and higher storage capacity of the Western Residual aquifer. C4 (Main Forest) is the most distinctive cluster, sitting 0.6 m deeper than adjacent C3 wells throughout the record (mean −1.47 m). This persistent depression relative to the surrounding dune aquifer is consistent with the canopy interception interpretation: the Corsican pine plantation intercepts a substantial fraction of gross precipitation before it reaches the water table, suppressing recharge and depressing the water table below that of the adjacent unforested dune. The C4 cluster mean hydrograph should be interpreted cautiously prior to 2010, when only three of the nine cluster wells were operational (CEH2, CEH13 and NW10); these wells are disproportionately located at higher elevations on the ridge margin and carry above-average water balance residuals which could be due to boundary subsidies, so the apparent high mean heads in 2006–2008 reflect an unrepresentative spatial sample rather than a genuine network-wide signal. Full nine-well coverage was not achieved until mid-2010. C5 (Coastal Forest) sits at comparable depth to C4 (mean −1.25 m) but with a damped seasonal signal, consistent with its position at the coastal edge of the plantation where both canopy interception and proximity to the coastal discharge boundary attenuate the seasonal response. Per-well seasonal amplitudes and their temporal evolution are analysed below.

The cumulative water balance panel (Figure 8a) confirms that the cluster separation is not an artefact of differential climate forcing — the 12-month rolling mean of the corrected net balance remains close to zero throughout the record, indicating that all clusters are responding to the same climatic regime. The structural differences between clusters therefore reflect hydrogeological and land cover controls rather than spatially variable precipitation or evapotranspiration.

The network mean water table elevation was 7.60 m AOD across the 66-well reference network, ranging from 3.12 m AOD (CEH11, C1 Lake Edge) to 13.30 m AOD at the ridge flank (CEH14, C4 Main Forest). CEH11's assignment to C1 is discussed in Section 3.2.3: its low elevation and coastal position place it near the Menai Strait tidal boundary rather than Llyn Rhos-ddu, and it represents the low-elevation extreme of the reference network rather than a typical lake-edge well. Excluding CEH11, the C1 cluster-mean elevation rises from 8.57 m AOD to 9.44 m AOD, and the network minimum shifts to CEH42 (C2, 4.47 m AOD). The Main Forest cluster (C4) occupies the highest mean water table elevation in the network at 9.52 m AOD, compared with 8.57 m AOD in the Lake Edge cluster (C1), 7.57 m AOD in the Dune cluster (C2), and 7.18 m AOD in the Western Residual (C3) — a contrast that reflects the topographic position of the plantation on the northern ridge flank rather than any difference in recharge behaviour. The Coastal Forest cluster (C5) sits at 4.60 m AOD, the lowest cluster-mean elevation in the network, consistent with its coastal-margin position.

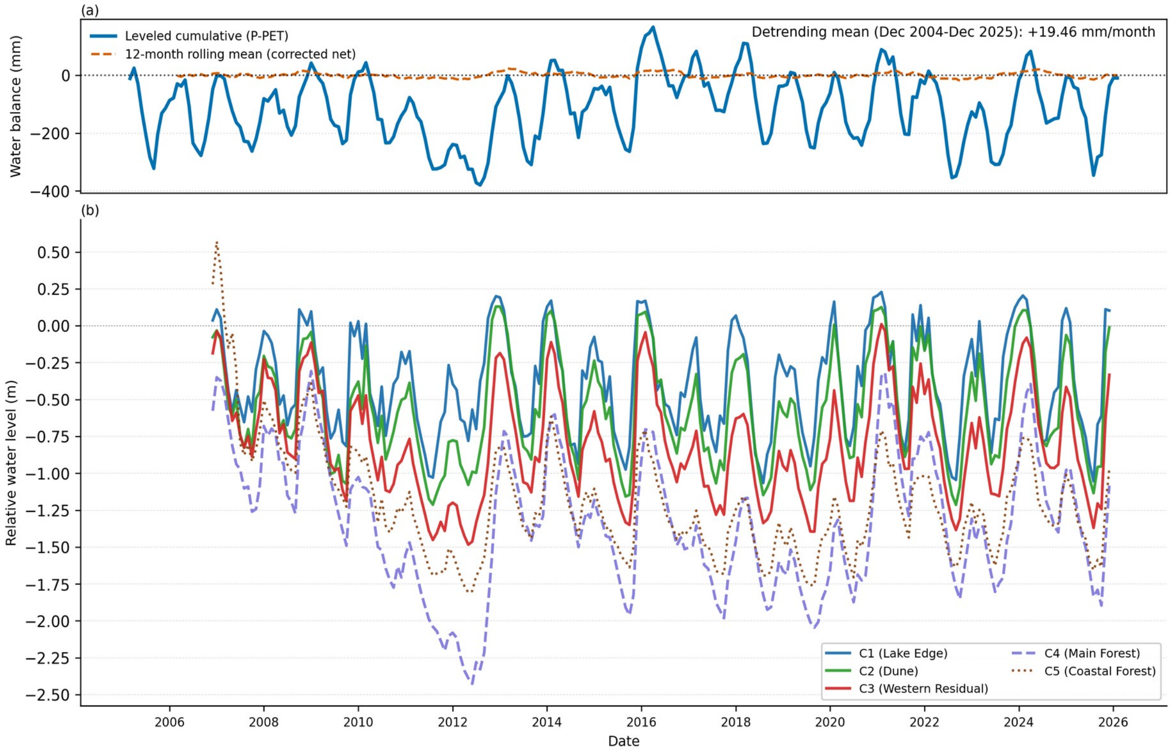

Figure 8. *Cluster-mean relative water level hydrographs for the five hydrogeological clusters, April 2005 to February 2026. (a) Levelled cumulative water balance (P − PET), detrended at the study-period mean of +19.46 mm/month (December 2004 to December 2025); the 12-month rolling mean of the corrected net balance (orange dashed) confirms no long-run directional trend in climate forcing over the monitoring period. (b) Monthly cluster-mean water level relative to pipe-top datum (m), where zero approximates ground surface, computed from all wells assigned to each cluster in the 66-well reference network. Cluster labels and colours follow those used in Figure 7. (02_03_cluster_hydrographs_wb.png)*

Per-Well Seasonal Amplitude. Per-well inter-decile amplitudes (p90 − p10) computed across three windows (full record, pre-2018, post-2018) confirm and refine the cluster-mean amplitude readings (Section 4.2.1; Figure 8). Cluster medians of per-well amplitudes are summarised in Table 2.

Cluster	n	Pre-2018 (m)	Post-2018 (m)	Δ pre→post (%)	Climate-normalised Δ (%)
C1 Lake Edge	7	0.83	1.00	+20%	+16%
C2 Dune	26	1.00	1.02	+2%	−7%
C3 Western Residual	19	1.04	1.06	+1%	−7%
C4 Main Forest	9	1.16	1.09	−6%	−6%
C5 Coastal Forest	5	0.89	0.73	−18%	−21%

Table 2. *Negative values indicate amplitude reduction (damping); positive values indicate amplitude amplification. Climate-normalised values exclude empirically-identified drought summers (2005, 2018, 2022) where Jun–Sep rainfall fell ≥1σ below the 1931–2017 baseline.*

Three patterns emerge. First, the open dune clusters (C2, C3) show stable raw amplitude but modest climate-normalised damping (≈7%), consistent with the post-2018 drought summers temporarily inflating apparent variability while the underlying seasonal range has slightly contracted. Second, **C5 Coastal Forest shows a pronounced amplitude collapse** of 18% (raw) to 21% (climate-normalised), the largest reduction in the network. Combined with C5's anomalously fast summer minimum decline (Section 4.8), this indicates that the C5 water table is becoming both lower in absolute terms and less seasonally responsive — recharge peaks no longer reach their pre-2018 heights even after correcting for drought summers, suggesting either progressive interception increase or a structural change in the recharge pathway at the coastal margin. Third, **C1 Lake Edge shows the opposite pattern**: amplitude has *increased* by 16–20%, driven primarily by deeper summer minima rather than higher winter peaks (median post-2018 summer minimum −1.04 m vs −1.47 m at C3; Table 2). The C3 cluster shows the highest within-cluster spread (post-2018 ratio max:min ≈ 2.0), approaching the amplitude-heterogeneity flagging threshold (Section 3.2.4) and consistent with the geological transition between deep aeolian sand and the western dune-till margin.

The amplitude analysis confirms that the cluster-mean coefficients reported in Section 4.2.2 represent stable averages across wells whose individual amplitude responses differ within most clusters by less than a factor of two, but with sufficient within-cluster variation in C3 and C4 to motivate the per-well coefficient mapping presented in Section 4.9.

4.2.2 Cluster Mechanistic Characterisation

The five clusters are characterized mechanistically in Table 3 using a state-space model fitted to the average hydrograph for each cluster (Section 3.4.1). β_1 is the recharge sensitivity — the water table rise (mm) per mm of rainfall; β_2 is the atmospheric draw — the water table decline (mm) per mm of PET; and $-\beta_3$ is the proportional drainage decay rate — the fraction of antecedent displacement above the drainage datum (Section 3.4) dissipated by lateral drainage each month. All coefficients are statistically significant ($p < 0.002$) at every cluster.

The recharge sensitivity coefficient β_1 is highest in the Lake Edge cluster (C1: 4.58), indicating that each millimetre of rainfall drives a proportionally larger water table rise in the shallow-till aquifer than in the Dune cluster (C2: 3.87) or Western Residual (C3: 3.58). The two forested clusters return near-identical and substantially lower recharge sensitivities (C4: 2.52, C5: 2.44), consistent with canopy interception suppressing the rainfall signal reaching the water table. The atmospheric draw coefficient β_2 is highest at C4 (2.50), followed by C3 (1.81) and C2 (1.74), with C5 (1.37) and C1 (0.96) lowest. The C4 value of β_2 is near-identical to its β_1, meaning that the Main Forest cluster loses as much water table per mm of PET as it gains per mm of rainfall — the plantation suppresses recharge (low β_1) while simultaneously amplifying the atmospheric draw signal relative to unforested clusters, producing a symmetrically sensitive water balance. Because β coefficients fold together vegetation effects, specific yield and soil moisture dynamics, they cannot be decomposed into crop coefficients without independent Sy estimates; this decomposition is developed in Section 5.6.1. That C5 (Coastal Forest) has low β_2 despite also carrying Corsican pine canopy suggests that coastal proximity moderates the atmospheric draw signal, possibly through higher ambient humidity or higher specific yield in the sandier coastal substrate. C1's low β_2 (0.96, the lowest in the network) reflects its shallow water table and rapid lateral drainage to the lake — the head-dependent drainage term captures the dominant loss mechanism, with the depth-dependent PET diagnostic (Section 3.7.1) confirming that C1 has the strongest residual depth-coupling ($\lambda = 1.90\ \mathrm{m}^{-1}$). The drainage decay coefficient $-\beta_3$ is highest at C1 (0.090) where rapid lateral exchange with Llyn Rhos-Ddu drives fast recession, similar at C2 (0.063) and C3 (0.060), lower at C5 (0.045), and lowest at C4 (0.021). The LCSC (= $100/\beta_1$) expresses these contrasts as a directly interpretable management metric: 22–26 mm of rainfall is required to raise Lake Edge and Dune water tables by 10 cm, compared with 28 mm in the Western Residual and 40–41 mm beneath the plantation.

The C1 cluster-mean coefficients in Table 3 include CEH11, which is geographically separated from the other six C1 members (Section 3.2.3). CEH11 sits at 3.6 m AOD on the south-eastern coast, approximately 1.6 km from the nearest lake-edge well, and its damped hydrograph likely reflects proximity to the Menai Strait tidal boundary rather than to Llyn Rhos-ddu. Its β_2 value of 0.77 is consistent with the C1 cluster mean (0.96) but may reflect a coastal constant-head mechanism rather than the lake-buffering process that governs the remaining members. CEH11's inclusion extends the C1 elevation range from 8.6–10.9 m down to 3.6 m; if excluded, the six remaining wells (five excluding CEH6, whose negative β_2 is discussed below) occupy a tight 9.1–10.9 m elevation band. The cluster-mean coefficients are not materially affected by CEH11's inclusion — its β_1 (4.99) and β_3 (0.088) are close to the C1 mean — but per-well spatial analyses within C1 should be interpreted with awareness that CEH11 responds to a different boundary condition.

An extended form of the state-space model incorporating a depth-dependent evapotranspiration coefficient (Section 3.7.1) was evaluated at cluster level; as the results of this analysis bear primarily on the mechanistic interpretation of the Forest cluster β_2 coefficient, they are discussed in Section 5.6.1 rather than reported here. The within-cluster coefficient variation is substantial, particularly in the forest zone where per-well β_2 ranges from 2.181 to 3.759 across the 14 C4 and C5 wells. Whether this variation is driven by substrate properties or canopy differences is examined through per-well spatial regression in Section 4.9.4.

Cluster	Label	n	β_1	β_2	$-\beta_3$	LCSC (%)	R^2
C1	Lake Edge	7	4.581	0.963	0.09	21.80	0.73
C2	Dune	26	3.871	1.745	0.06	25.80	0.75
C3	Western Residual	19	3.576	1.806	0.06	28.00	0.81
C4	Main Forest	9	2.518	2.504	0.02	39.70	0.68
C5	Coastal Forest	5	2.436	1.374	0.05	41.00	0.70

Table 3. Cluster mechanistic characterisation derived from the displacement-formulation SSM (Section 3.4). β_1 = recharge sensitivity (mm water table rise per mm rainfall); β_2 = atmospheric draw (mm water table decline per mm PET); $-\beta_3$ = proportional drainage decay rate (fraction of displacement above the 3.7 m drainage datum dissipated per month). LCSC = $100/\beta_1$; values express the millimetres of rainfall required to raise the cluster average water table by 10 cm. R^2 is one-step diagnostic fit. All coefficients significant at $p < 0.01$ or better. (03_03_cluster_mechanistic_coefficients.csv).

4.2.3 Recharge, Loss and the Boundary Subsidy: A Cluster-Level Water Balance

Water balance decomposition applied to all five clusters (Section 3.7.2) reveals markedly different recharge efficiencies across clusters receiving identical climate forcing. The head-space decomposition is presented in Table 4a and Figure 9; an indicative volumetric conversion is presented in Table 4b and Figure 9b. The derivation and defence of the specific yield values is addressed in Section 4.2.4.

Recharge into the water table ($\beta_1 \cdot \bar{P}$) declines systematically from the Lake Edge cluster (C1: 0.341 m/month) through the Dune (C2: 0.288 m/month) and Western Residual (C3: 0.266 m/month) to the forested clusters (C5: 0.181 m/month; C4: 0.187 m/month). Total losses (atmospheric draw plus drainage) track recharge closely, ranging from 0.186 m/month at C4 to 0.345 m/month at C1. Residuals are less than 2.5% of total losses in all clusters, confirming that the three-term displacement-formulation SSM captures the principal water balance fluxes without requiring an unmeasured boundary contribution. The drainage partition varies markedly across clusters: drainage accounts for 85% of total losses at C1 (reflecting rapid lateral exchange with Llyn Rhos-Ddu) but only 26% at C4, where atmospheric draw dominates. C2 and C3 show drainage-dominated partitions (67%/33% and 63%/37% respectively), while C5 is intermediate (59%/41%).

Cluster	Label	LCSC (%)	Recharge (m/month)	Atm. Draw (m/month)	Drainage (m/month)	Total Loss (m/month)	Residual (m/month)
C1	Lake Edge	21.8	0.341	0.053	0.292	0.345	−0.004
C2	Dune	25.8	0.288	0.095	0.196	0.291	−0.004
C3	Western Residual	28.00	0.266	0.099	0.167	0.266	0.000
C4	Main Forest	39.70	0.187	0.138	0.048	0.186	0.002
C5	Coastal Forest	41.00	0.181	0.075	0.110	0.185	−0.004

Table 4a. *Mean monthly head-space water balance decomposition by hydrogeological cluster, Newborough Warren 2005–2026. Recharge ($\beta_1 \cdot \bar{P}$), atmospheric draw ($\beta_2 \cdot \overline{PET}$) and gravity drainage ($\beta_3 \cdot \bar{h}$_disp) are expressed as metres of water table change per month (m/month), where $\bar{h}$_disp is the mean displacement above the 3.7 m drainage datum (Section 3.4). Coefficients are cluster-centroid fits from the displacement-formulation SSM (Table 3). Residuals are less than 2.2% of total losses in all clusters, confirming that the three-term formulation closes the water balance without requiring an unmeasured boundary contribution. All clusters receive identical climate forcing: $\bar{P}$ = 74.4 mm/month, $\overline{PET}$ = 55.1 mm/month. Drainage and ET percentages shown below. (Source: 16_water_bal_table.csv)*

An indicative volumetric expression of the water balance is presented in Table 4b and Figure 9b. All clusters receive 894 mm/yr of rainfall; for the forested clusters (C4, C5) canopy interception removes 215 mm/yr (24%; Freeman, 2008), leaving net rainfall of 680 mm/yr reaching the ground surface. The ET/drainage partition is bracketed by two independent methods — the SSM headspace β_2/β_3 ratio and a seasonal recession curve analysis — with the mid-point reported in Table 4b. At steady state, total losses equal total inputs (P or P_net for forest clusters) and the hatched band in Figure 9b spans the partition uncertainty range between the two methods.

Cluster	Label	P (mm/yr)	I (mm/yr)	P_net (mm/yr)	ET mid (mm/yr)	Drainage mid (mm/yr)
C1	Lake Edge	893	0	893	246	647
C2	Dune	893	0	893	402	490
C3	Western Residual	893	0	893	308	585
C4	Main Forest	893	214	678	416	262
C5	Coastal Forest	893	214	678	246	433

Table 4b. *Indicative annual volumetric water balance by cluster, 2005–2026. At steady state, total losses (ET + drainage) equal total inputs (P for open dune clusters; P − I for forested clusters). I = canopy interception (24% of P; Freeman, 2008), applied to C4 and C5 only. ET and drainage mid-point values represent the mean of two independent partitioning methods: the SSM headspace* β_2/β_3 *ratio and seasonal recession curve analysis. The partition uncertainty range is shown graphically in Figure 9b as a hatched band. Interception appears identically on both input and loss bars of Figure 9b and cancels in the net surplus. (Source: 16_water_bal_vol_table.csv)*

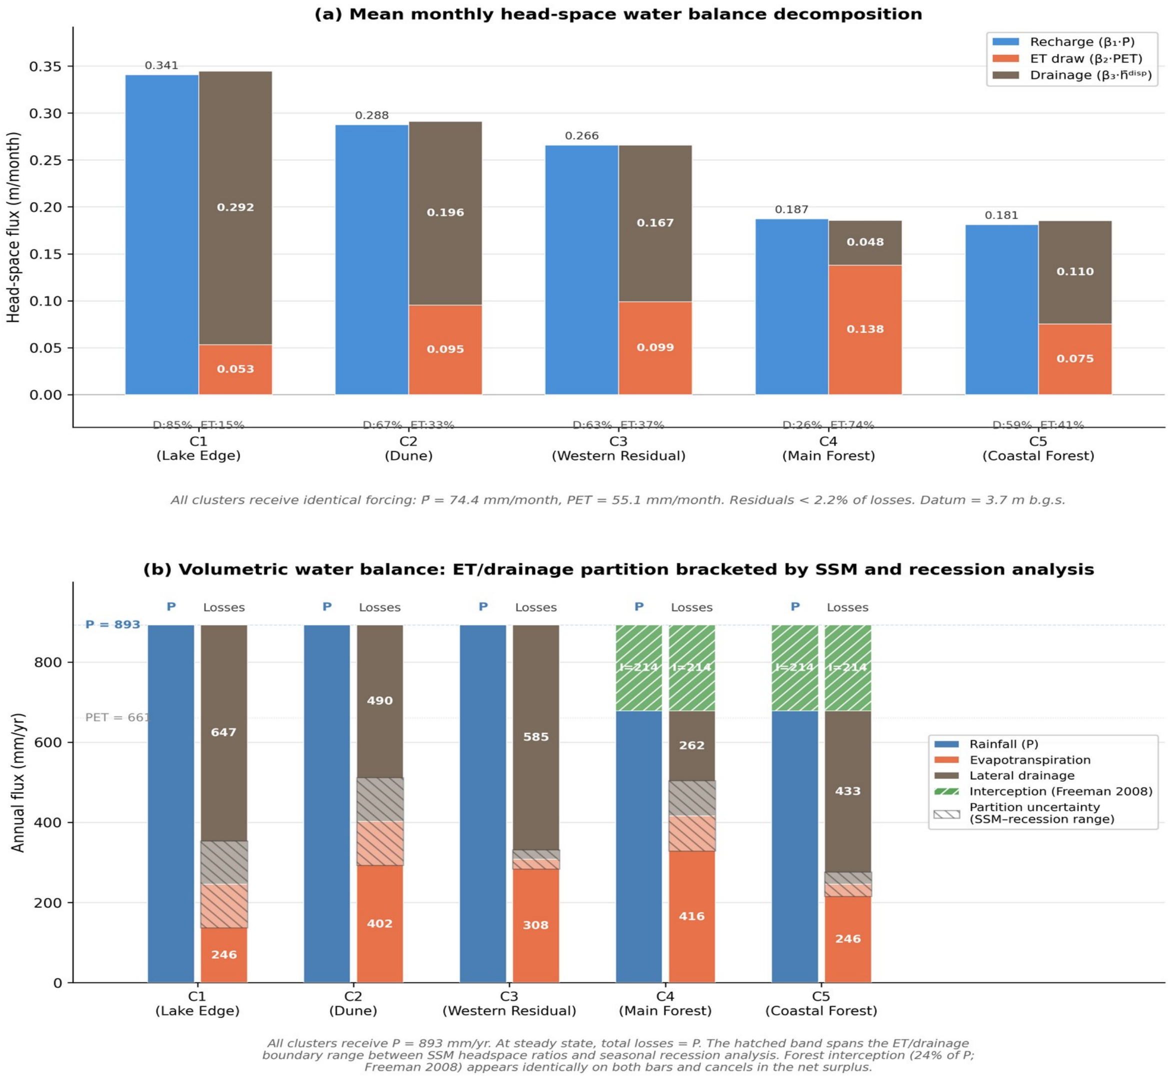

(a) Mean monthly head-space water balance decomposition

All clusters receive identical forcing: $\bar{P}$ = 74.4 mm/month, PET = 55.1 mm/month. Residuals < 2.2% of losses. Datum = 3.7 m b.g.s.

(b) Volumetric water balance: ET/drainage partition bracketed by SSM and recession analysis

All clusters receive P = 893 mm/yr. At steady state, total losses = P. The hatched band spans the ET/drainage boundary range between SSM headspace ratios and seasonal recession analysis. Forest interception (24% of P; Freeman 2008) appears identically on both bars and cancels in the net surplus.

Figure 9*. Water balance decomposition by hydrogeological cluster, Newborough Warren 2005–2026. (a) Mean monthly head-space decomposition: for each cluster, the left bar shows recharge ($\beta_1 \cdot \bar{P}$) and the right bar shows losses decomposed into gravity drainage ($\beta_3 \cdot \bar{h}$_disp) and atmospheric draw ($\beta_2 \cdot PE\bar{T}$). Drainage and ET percentages are shown below each cluster label. All clusters receive identical climate forcing: $\bar{P}$ = 74.4 mm/month, $PE\bar{T}$ = 55.1 mm/month. Drainage datum = 3.7 m b.g.s. Residuals < 2.2% of losses. (b) Indicative annual volumetric water balance. All clusters receive P = 893 mm/yr; for forested clusters (C4, C5) canopy interception (I = 214 mm/yr, 24% of P; Freeman, 2008) appears identically on both input and loss bars and cancels in the net surplus. The hatched band spans the ET/drainage partition uncertainty between the SSM headspace ratio and seasonal recession analysis. (16_water_bal_bar_ms.png, generated by 16_water_bal.py).*

4.2.4 Specific Yield Estimation

WTF event-based median Sy estimates for all five clusters, derived from 35–53 qualifying monthly rising-limb events per cluster, are presented in Table 4c. Uncorrected median values range from 0.210 (C1 Lake Edge) to 0.360 (C5 Coastal Forest), with broad interquartile ranges reflecting the inherent noise in monthly cluster-mean data. The two forested clusters return uncorrected medians (C4: 0.317, C5: 0.360) that exceed the open dune range (0.210–0.341), because gross rainfall overstates the recharge flux reaching the water table beneath the canopy — the WTF method attributes the full rainfall to recharge when in reality 24% is intercepted, inflating the apparent Sy. After applying the interception correction of 24% of P (Freeman, 2008), the forested cluster medians fall to 0.243 (C4) and 0.320 (C5), within the range of the open dune clusters (0.210–0.341). This convergence confirms that the attenuated water table response beneath the plantation reflects the surface boundary condition (canopy interception) rather than a fundamentally different substrate. Within this corrected dataset, a secondary pattern is visible: the two clusters with the most constrained substrate — C1 (Lake Edge, shallow till over flat clay; corrected median 0.210) and C4 (Main Forest, irregular bedrock on the ridge flank; 0.243) — return lower Sy than the two clusters on deeper, more uniform sand — C3 (Western Residual; 0.341) and C5 (Coastal Forest; 0.320), with C2 (Dune; 0.292) intermediate. At C1 the constraint is a flat clay/till base with a short drainage path to the lake; at C4 it is irregular bedrock topography where buried ridges impede lateral drainage and create within-cluster heterogeneity in β_3. The pattern recurs in the per-well β_2 regression (Section 4.9.4) where elevation — reflecting both vadose zone thickness and proximity to buried bedrock — explains 95% of within-forest β_2 variance.

Individual well Sy values underlying the spatial mapping of Figure 10 are considered separately below. All five WTF medians exceed the assumed Sy values used in the water balance (C1: 0.08; C2–C5: 0.12) by a factor of two to three, consistent with the known tendency of monthly WTF to conflate gravity drainage with capillary fringe release (Healy and Cook, 2002; Scanlon et al., 2002). The WTF estimates should therefore be interpreted as upper bounds on effective storage. Direct field measurement by slug test or pumping test at representative wells per cluster is identified as a priority for future work.

Cluster	Label	n events	Sy assumed	Sy WTF median	Q25	Q75	Interception corrected?
C1	Lake Edge	61	0.08	0.21	0.13	0.272	No
C2	Dune	63	0.12	0.292	0.216	0.38	No
C3	Western Residual	61	0.12	0.341	0.289	0.439	No
C4	Main Forest	50	0.12	0.317	0.26	0.399	No
C5	Coastal Forest	34	0.12	0.36	0.335	0.433	No
C4	Main Forest (corrected)	61	0.12	0.243	0.187	0.336	Yes — Freeman (2008)
C5	Coastal Forest (corrected)	50	0.12	0.32	0.243	0.381	Yes — Freeman (2008)

Table 4c. *Specific yield (Sy) estimates by hydrogeological cluster derived from the water table fluctuation (WTF) method (Healy and Cook, 2002). Event-based median Sy computed from monthly rising-limb events satisfying Δh > 5 mm and net recharge P − PET > 10 mm over the monitoring period April 2005 to February 2026. For the forested clusters (C4, C5), two estimates are reported: uncorrected (using gross rainfall) and corrected (net recharge adjusted for 24% canopy interception following Freeman, 2008). The corrected C4 median (0.243) and C5 median (0.320) fall within the range of the open dune clusters (0.210–0.341), confirming broadly uniform aquifer storage properties across the site. IQR = interquartile range; wide IQR reflects noise inherent in monthly cluster-mean data rather than genuine spatial heterogeneity in aquifer storage. Monthly WTF estimates are known to overestimate effective Sy through conflation of gravity drainage with capillary fringe release and should be interpreted as upper bounds rather than point estimates of long-term effective storage (Healy and Cook, 2002; Scanlon et al., 2002). Generated by 17_wtf_specific_yield.py (output: 17_wtf_01_sy_estimates.csv).*

Individual well WTF Sy estimates were derived from 27–72 qualifying rising-limb events per well (mean 52 across the 66-well reference network). Per-well cluster-mean values — C1: 0.208, C2: 0.260, C3: 0.316, C4: 0.246, C5: 0.309 — are consistent with the cluster-aggregate estimates of Table 4c, confirming that the spatial pattern is robust. These per-well means are by-products of the spatial interpolation: they describe the inputs to Figure 10 and preserve within-cluster variability, but each per-well Sy is noisier than the cluster-aggregate estimate of Table 4c, which fits a single median to the smoothed cluster-mean hydrograph and is therefore more robust as an estimate of cluster storage. The well-level analysis is presented here for its spatial content rather than as a competing storage estimate. The spatial distribution of Sy, shown in Figure 10, is broadly uniform across the site, with no abrupt discontinuities coinciding with the cluster margins, providing further confirmation that the cluster behavioural differences reflect surface hydrology and land cover rather than heterogeneity in aquifer storage architecture.

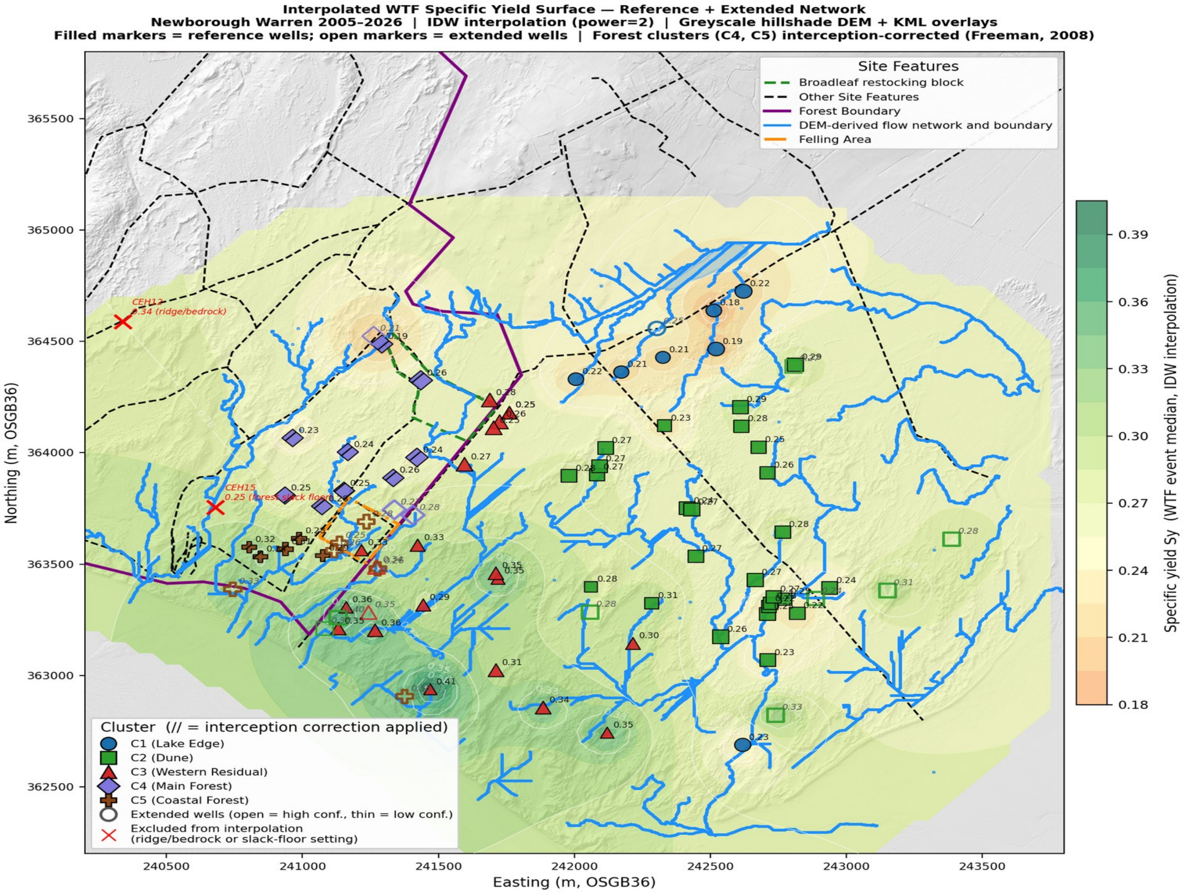

Figure 10. *Interpolated WTF specific yield surface for the combined reference and extended networks, Newborough Warren 2005–2026. Per-well event-median Sy interpolated using IDW (power = 2) over 1 m LiDAR DEM hillshade. Filled markers denote reference-network wells; open markers denote extended-network wells classified by Pearson affinity (Section 4.3, best-match cluster). Forest cluster (C4, C5) wells are interception-corrected following Freeman (2008). Two wells are excluded from interpolation: CEH12 (ridge/bedrock setting) and CEH15 (forest slack floor), both of which produce anomalous Sy values inconsistent with the surrounding aquifer. Per-well cluster-mean Sy values — C1: 0.208, C2: 0.260, C3: 0.316, C4: 0.246, C5: 0.309 — are consistent with the cluster-aggregate estimates of Table 4c. The spatial distribution is broadly uniform across the open dune (Sy 0.22–0.31), with no abrupt discontinuities coinciding with cluster margins, providing further confirmation that the cluster behavioural differences reflect surface hydrology and land cover rather than heterogeneity in aquifer storage architecture. (18_wtf_04_sy_contour_extended.png, generated by 18_wtf_spatial.py) LiDAR data: © NRW & OS; see Section 10 for full attribution. Individual well Sy estimates for all 86 wells are provided in Table S4.*

4.3 Checking the Clusters Pearson Affinity Analysis

Of the 66 wells in the reference network, 63 confirmed their primary cluster membership through the Pearson Affinity Analysis (Section 3.3). Membership is characterized by three distinct tiers based on the correlation delta (Δr) between assigned and neighbouring cluster centroids: Core members (17 wells) with an unambiguous dominant affinity to their assigned cluster (Δr > 0.05), representing stable hydrogeological anchors within their designated geomorphological zone; Fuzzy members (46 wells) exhibiting closely competing affinities between their primary and secondary clusters (Δr < 0.05), identifying areas where the groundwater regime transitions smoothly between adjacent structural blocks; and three Spy classifications where the well's highest temporal affinity belongs to a different cluster than Ward's assignment. CEH21 and CEH4 (both assigned C3 Western Residual) show strongest affinity with C5 Coastal Forest, consistent with their coastal-margin position where the Western Residual and Coastal Forest signatures overlap; NW13 (assigned C3) shows strongest affinity with C2 Dune. All three Spy deltas are small (|Δr| < 0.01), confirming that these are borderline assignments at gradational boundaries rather than gross misclassifications.

Classification confidence was further characterized by the Multi-Cluster Affinity flag, raised where three or more cluster centroids achieved a correlation of r > 0.90. This criterion was met at 46 of 66 wells, reflecting the high mutual correlation among the five cluster centroids in a contiguous, highly transmissive sand aquifer. Of these, 21 are concentrated within the gradational C1/C2/C3 transitional zone of the mid-eastern dune sector, 14 lie along the C2/C3/C5 boundary where the Dune, Western Residual and Coastal Forest signatures converge, and 9 lie along the C3/C4/C5 forest–dune boundary. The prevalence of MCA flags confirms the continuous nature of the sand matrix, where boundary wells respond to system-wide climatic forcing rather than localized topographic drivers, and identifies the precise locations of greatest structural ambiguity within the network (Figure 12a).

The extended network analysis (Section 3.3.3) classified all 23 extended wells against the reference centroids, yielding 4 Ext_Core and 19 Ext_Fuzzy assignments (Figure 12b). The high proportion of Fuzzy classifications reflects the shorter and earlier records of these wells rather than genuine hydrogeological ambiguity — with less temporal overlap against the centroid time series, affinity margins are naturally compressed.

Within the extended network, the clearfell treatment wells showed mixed affinities across the C3/C4/C5 boundary, with all wells carrying Multi-Cluster Affinity flags (r > 0.90 with three or more centroids). FE1 and LIS1 returned strongest affinity with C4 Main Forest (r = 0.989 and 0.972 respectively), while FE2 and FE4 showed marginally strongest affinity with C5 Coastal Forest (r = 0.974 and 0.970) with C4 within 0.006. FE3 showed strongest affinity with C5 (r = 0.984) followed by C3 (r = 0.975). WMC3, in the reference network, showed strongest affinity with C3 (r = 0.965) with C5 a close second (r = 0.958). The small deltas (Δr = 0.002–0.037) confirm that the clearfell zone sits at the gradational boundary between the Western Residual, Main Forest and Coastal Forest clusters rather than within any single hydrogeological unit. The persistence of forest-range affinities (C4 or C5) at these wells despite the December 2017 clearfell is consistent with either residual canopy influence from the surrounding intact plantation, incomplete hydrological recovery, or the deep sandy substrate at high elevation generating the forest signature independently of the canopy.

WMC3, which is in the reference network, similarly showed strongest affinity with C3 (r = 0.973) with C4 a close second (r = 0.949) (Figure 11).

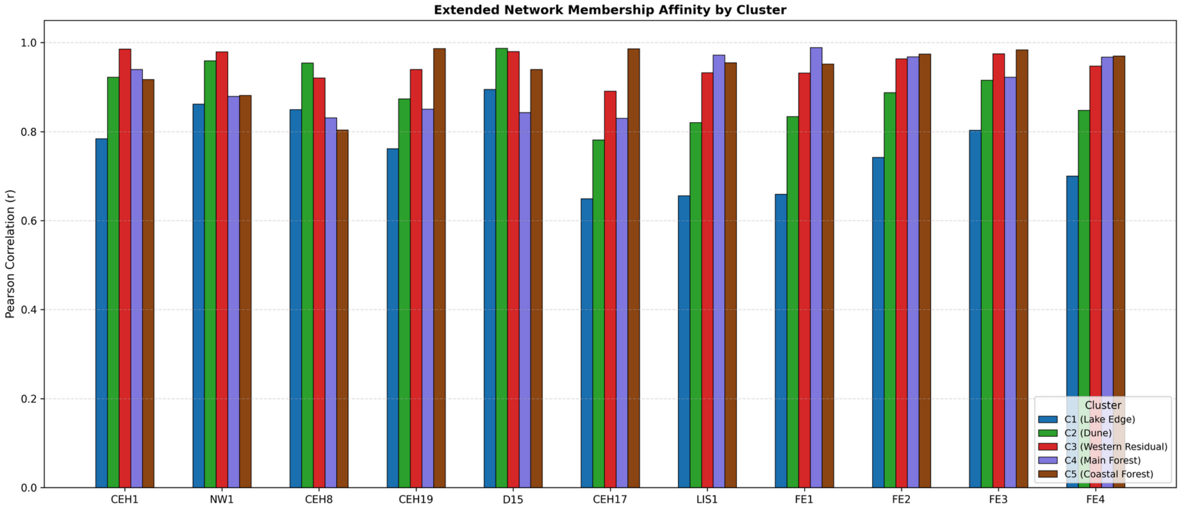

Figure 11. *Pearson affinity analysis of selected reference network wells, showing the correlation coefficient (r) between each well's z-scored time series and each of the five cluster centroids. Wells are selected to illustrate the range of affinity confidence from strong single-cluster dominance (Core members) to near-equal competing affinities (Fuzzy members), along with wells in or close to the clearfell experiment. (06_pear_01_affinity_chart_extended.png)*

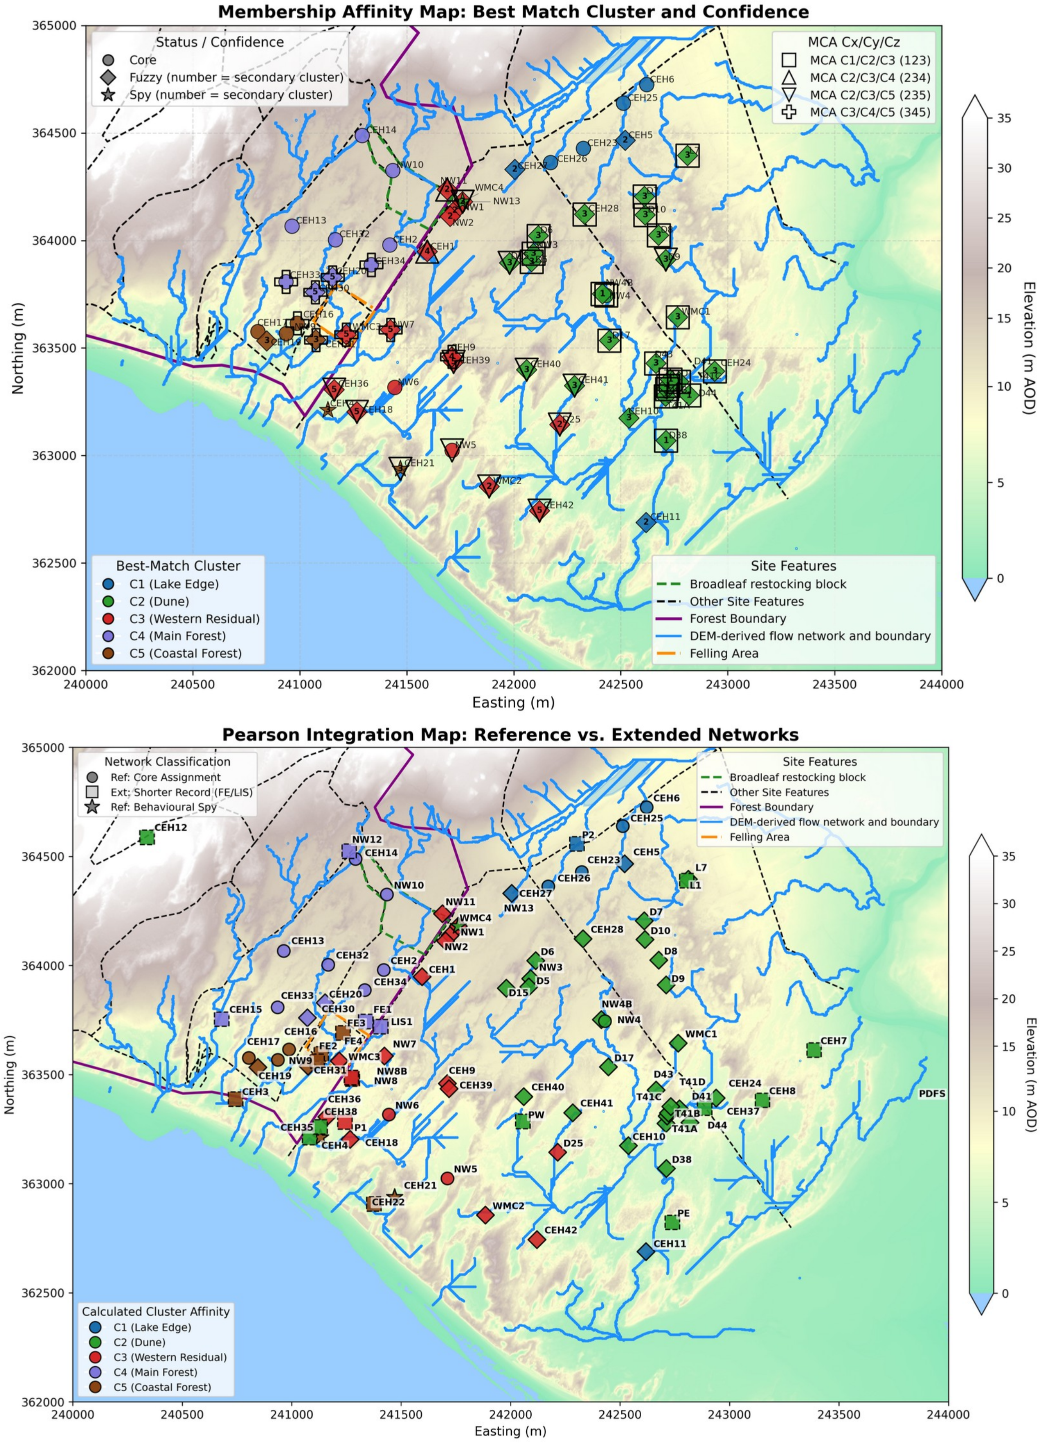

Figure 12. *Spatial distribution of the five primary hydrogeological clusters validated via Pearson affinity analysis. Points are coloured by best-match cluster following z-score standardisation. Symbology denotes classification confidence: circles represent Core members (Δr > 0.05), diamonds represent Fuzzy members (Δr < 0.05), and triangles represent Spy classifications (best-match ≠ assigned cluster). Panel (a): Reference network (66 wells) with records of at least 100 months extending to February 2026. Large hollow overlays identify wells with Multi-Cluster Affinity (MCA), exhibiting high convergence (r > 0.90) with three or more cluster average hydrographs. (05_pear_01_spatial_confidence_map.png) Panel (b): Pearson affinity classifications for all 89 wells across the combined reference and extended networks. (06_pear_02_integration_map.png) (LiDAR data: © NRW & OS; see Section 10 for full attribution).*

4.4 Model Benchmarking: State-Space Model Comparison with Traditional Linear Models

The SSM was compared to a TLM as outlined in Section 3.4.2. In one-step diagnostic mode — where each prediction is corrected with the observed water level from the previous month — both models track the observed record almost identically (median R^2 0.91 and 0.92 respectively; Table 5). The difference emerges under autonomous simulation. Over 100 months without observational correction, the TLM diverges rapidly at wells with strong head-dependent drainage (e.g. CEH6, C2 Dune: TLM NSE = −1.15, SSM NSE = +0.63; Figure 13), while the SSM maintains physical consistency through its storage decay term.

Across the full 66-well reference network, the SSM achieves a median iterative NSE of 0.76 compared with 0.14 for the TLM — a median gain of +0.62 NSE units — and returns a positive iterative NSE at 65 of 66 wells compared with 42 of 66 for the TLM (Table 5). The improvement varies systematically across the five clusters (Figure 14a, 13b): C1 Lake Edge shows the largest median gain (+1.292), where strong head-dependent drainage to Llyn Rhos-Ddu makes the storage decay term essential; C5 (+0.726), C3 (+0.624) and C2 (+0.538) show substantial gains; and C4 Main Forest shows the smallest (+0.229), reflecting its characteristically low β_3 (0.021) which produces a near-linear climate-to-water-table relationship that the TLM handles more adequately. The sole well where the SSM fails is CEH14 (C4, iterative NSE = −4.31), a ridge-flank well where the displacement formulation does not adequately capture the boundary condition. Spatial visualisation (Figure 14) confirms this pattern geographically, with the largest gains concentrated along the forest–Warren boundary and the C1 lake-buffer sector.

The depth-dependent PET grid search (Section 3.7.1) provides an independent diagnostic of how completely the standard fixed-β_2 formulation captures the atmospheric draw signal. The grid search produced statistically significant NSE improvement at all five clusters, where λ (m^{-1}) is the exponential decay rate controlling how rapidly the PET signal attenuates with water table depth: C1 ($\Delta NSE = +0.182$, $\lambda = 1.90\ m^{-1}$), C4 (+0.104, $\lambda = 0.20$), C2 (+0.073, $\lambda = 0.80$), C5 (+0.032, $\lambda = 0.45$) and C3 (+0.025, $\lambda = 0.35$). AIC improvement is decisive at every cluster (ΔAIC +17 to +98). The fitted λ values are physically interpretable as inverse 1/e decay depths: C1's strong decay (1/e = 0.53 m) is consistent with direct evaporation from a frequently shallow water table near the lake margin, while C4's weak decay (1/e = 5.0 m) confirms that C4's high β_2 reflects low specific yield rather than a depth-dependent ET pathway. The mechanistic interpretation is developed in Section 5.6.1.

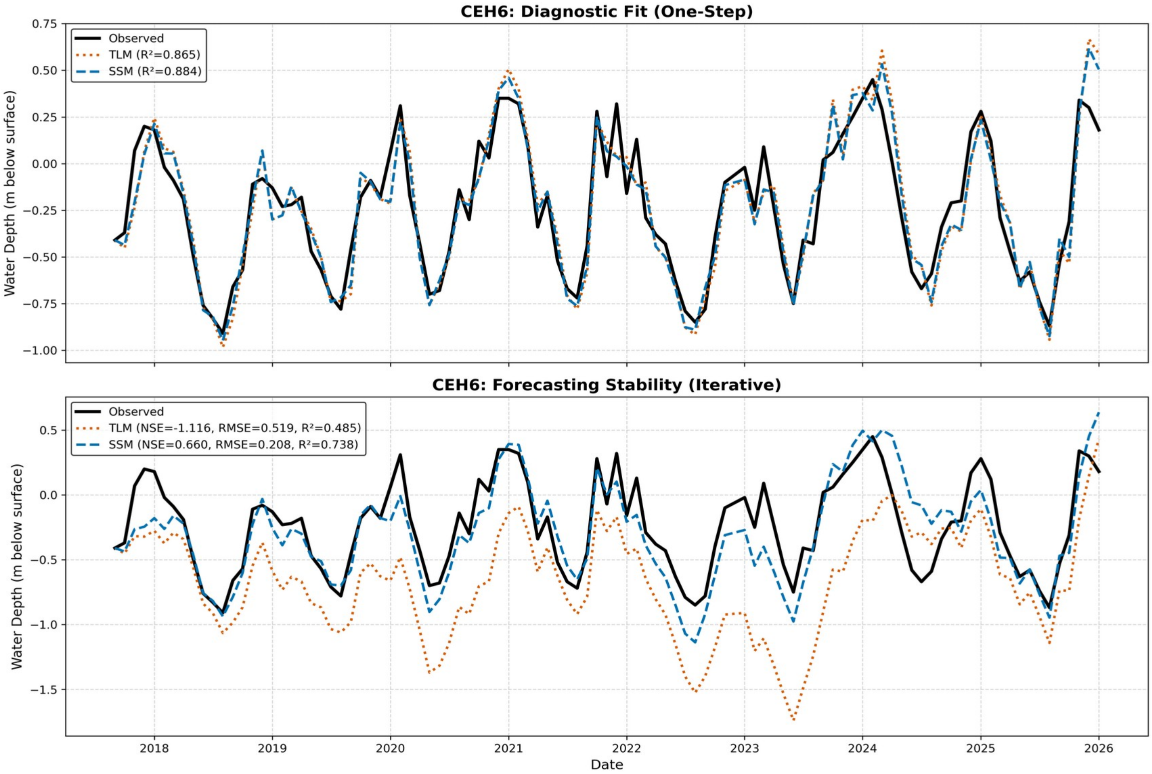

Figure 13*. Dual-model performance comparison at dune well CEH6. Upper panel: one-step diagnostic fit — both TLM and SSM track observed water levels closely when corrected with real observations at each step. Lower panel: iterative 100-month autonomous simulation without observational correction — the TLM (NSE = −1.15) diverges rapidly, while the SSM (NSE = +0.63) maintains physical consistency through its storage decay term. CEH6 records the largest single NSE improvement (+1.78) in the 66-well reference network. (08_lcsc_01_ceh6_showdown.png)*

Metric	**TLM**	**SSM**	**Improvement Δ**
Median one-step R²	0.91	0.92	0.01
Median iterative NSE	0.14	0.76	0.62
Wells with iterative NSE > 0	42 / 66	65 / 66	—
Max NSE improvement (CEH25)	−1.47	0.68	2.15
CEH6 iterative NSE	−1.12	0.66	1.78
CEH14 iterative NSE	0.53	−4.31	−4.84

Table 5. *Model benchmarking summary comparing the Traditional Linear Model (TLM) and State-Space Model (SSM) across the 66-well reference network. Performance evaluated over the most recent 100 months at each well. One-step R² uses observed h(t−1) at each step; iterative NSE uses simulated h(t−1) over the full 100-month evaluation window without observational correction. CEH25 is the well recording the largest single NSE improvement; CEH6 is the illustrative dune well shown in Figure 13; CEH14 is the sole well where the SSM underperforms the TLM. (08_lcsc_model_stats.csv, generated by 08_model_benchmarking.py).*

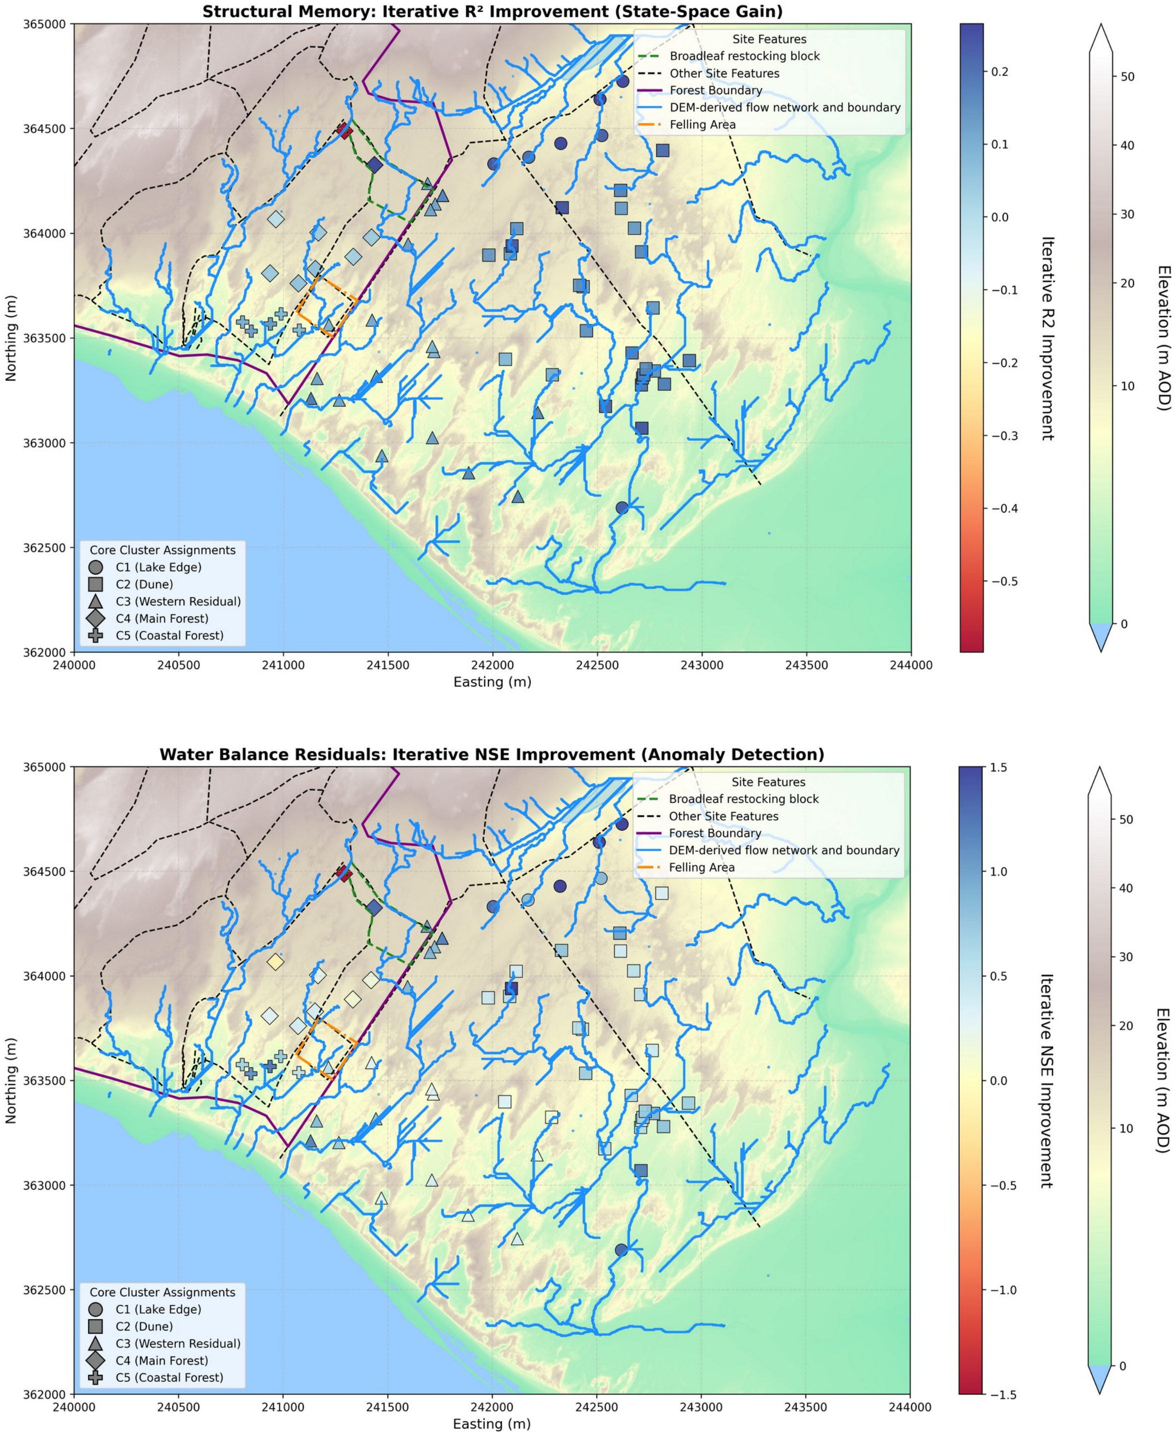

Figure 14. *SSM gain over TLM across all 66 reference wells. Panel (a): iterative R² improvement — largest gains in the Lake Edge sector and eastern dune, smallest in the Main Forest cluster where the TLM already performs adequately. (08_lcsc_02_r2_improvement_map.png) Panel (b): NSE improvement — corresponding spatial pattern with the largest gains concentrated in the C1 Lake Edge cluster. The spatial distribution of improvement physically delineates subsurface structural boundaries without invasive coring. (08_lcsc_03_nse_improvement_map.png) LiDAR data: © NRW & OS; see Section 10 for full attribution.*

4.5 Dune Scraping Intervention Analysis

4.5.1 Background Environmental Drift

Three dipwells within the dune slack system — CEH36, CEH18 and CEH21 — were subject to ground surface lowering by mechanical scraping, with CEH36 scraped in April 2015 and CEH18 and CEH21 scraped in October 2023. The analysis uses a two-tier CUSUM-BACI framework (Section 3.5) to isolate the scraping signal from background environmental drift, with CEH4 and CEH22 serving as paired controls. Tier 1 analysis at CEH4 and CEH22 both show accelerating negative CUSUM trajectories across the monitoring period, reaching −11.51 m and −17.46 m respectively by early 2026 (Section 3.5.4; Figure 15). The convergent deterioration of two spatially separated control wells — CEH22 and CEH4 near the coast — independently of any management action confirms that progressive coastal erosion of the dune edge is driving a structural decline in water table depth across this part of the network. CEH22's trajectory is the more severe, dominated by a coastal drawdown signal that accelerated markedly from 2021 onwards.

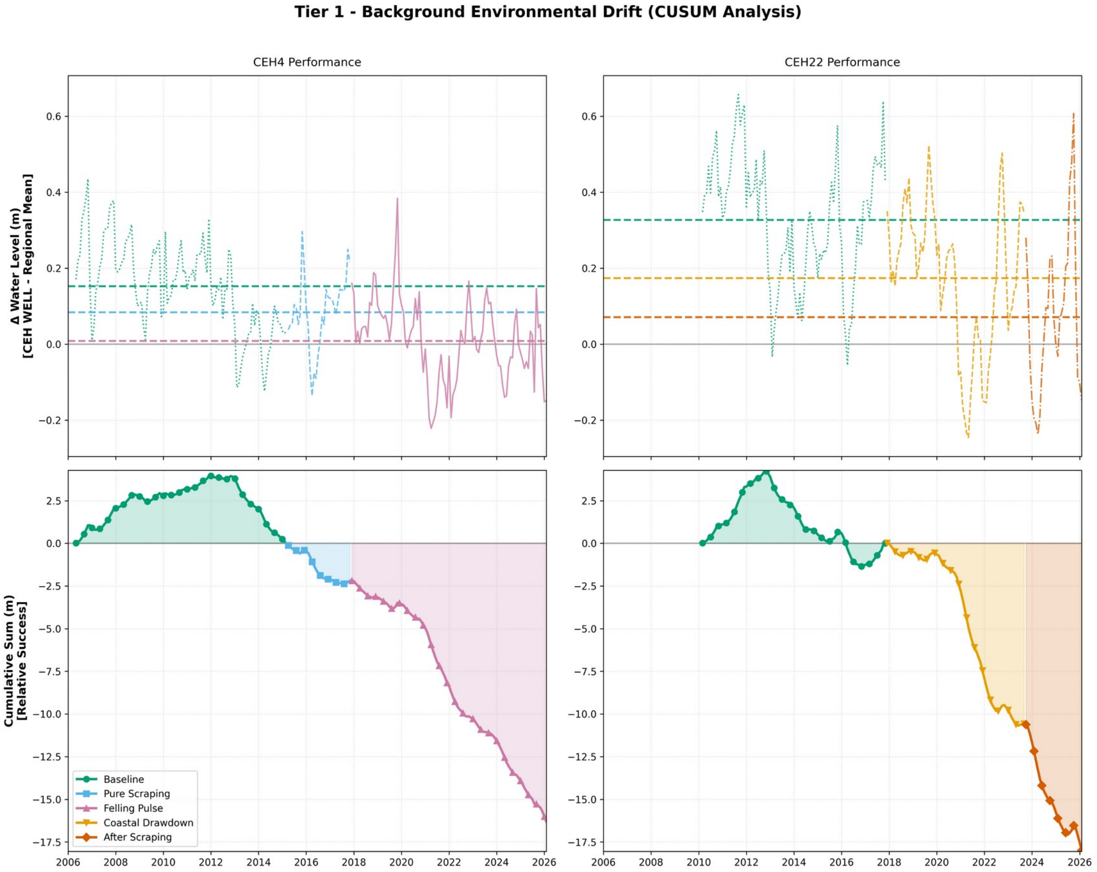

Figure 15. *Tier 1 background environmental drift analysis. CUSUM time series for control wells CEH4 (left) and CEH22 (right) decomposed by management era relative to the regional mean. Upper panels show monthly Δ water level (treatment well minus regional mean) with era mean levels indicated by dashed horizontal lines. Lower panels show cumulative CUSUM trajectories. A rising CUSUM indicates the well is performing above the regional mean (shallower water table); a falling CUSUM indicates below-average performance (deeper water table). The accelerating negative CUSUM at both control wells confirms a system-wide deterioration trend independent of local management interventions. (09_scrape_05_tier1_background_drift.png)*

Against this background of coastal erosion, CEH4 shows an additional secondary deterioration signal from December 2017 onwards consistent with the clearfell felling pulse, though this is superimposed on and inseparable from the longer-run erosion trend. These Tier 1 trajectories establish that both local controls are experiencing substantial background decline that must be accounted for before any scraping benefit can be assessed.

4.5.2 Paired BACI and Robustness

In this context, Tier 2 paired analysis shows a modest positive signal at CEH36 during the Pure Scraping era (April 2015 to December 2018), the only well and era combination to show a net benefit after Tier 1 correction. Three independent methods converge on a consistent estimate: raw BACI vs CEH4 (+0.131 m), synthetic control from 11 donor wells (+0.135 m), and SSM forward residual (+0.083 m) — reflecting permanent ground surface lowering rather than a hydrodynamic deviation from the climate-driven trajectory (Figure 16). The cumulative CUSUM for CEH36 rises strongly from 2018 onwards, though this reflects the accelerating deterioration of CEH4 as much as any genuine improvement at CEH36, which is why the synthetic control method is critical to isolating the true scraping signal from CEH4's own trajectory (Figure 17). CEH18 sits consistently negative relative to CEH4 throughout the full record including the pre-scraping baseline, indicating that this well has always sat deeper than its control rather than having been driven down by the clearfell; only a modest recovery is visible after the October 2023 scraping. CEH21's rising CUSUM from 2022 onwards should be interpreted cautiously as its paired control CEH22 is itself deteriorating through coastal erosion, which inflates the apparent relative improvement at CEH21 independently of any scraping benefit; these results are noted for completeness only.

After Tier 1 correction for background drift, the net benefit at CEH36 reversed to −0.046 m during the Felling Pulse era, as the coastal erosion signal in CEH4 — compounded from 2018 by the secondary clearfell pulse — progressively eroded the scraping gain. The raw Tier 2 CUSUM tells a different story, with CEH36 rising strongly relative to CEH4 through this period — but as established above, this reflects CEH4's own deterioration rather than genuine improvement at CEH36. At CEH18 and CEH21, scraped in October 2023, no scraping-specific response survived Tier 1 correction for background drift. CEH18's corrected net benefit was −0.029 m after scraping; CEH21's paired control (CEH22) is itself on a coastal erosion trajectory, so a net benefit cannot be reliably estimated. The monitoring window for both wells is short (two post-scraping summers) and the analysis is likely underpowered to detect benefits of the magnitude observed at CEH36.

The annual summer minimum records for all three treatment wells, partitioned by management era and plotted against the Curreli et al. (2013) ecological thresholds, are presented in Figure 19. The progressive deepening of the CEH4 control well across all eras — from 1.16 m at baseline to 1.18 m during the Pure Scraping era and 1.40 m during the Felling Pulse — is consistent with the system-wide climate drying trend documented in Section 4.8.1 and confirms that management benefit must be assessed against this background trajectory rather than in absolute terms.

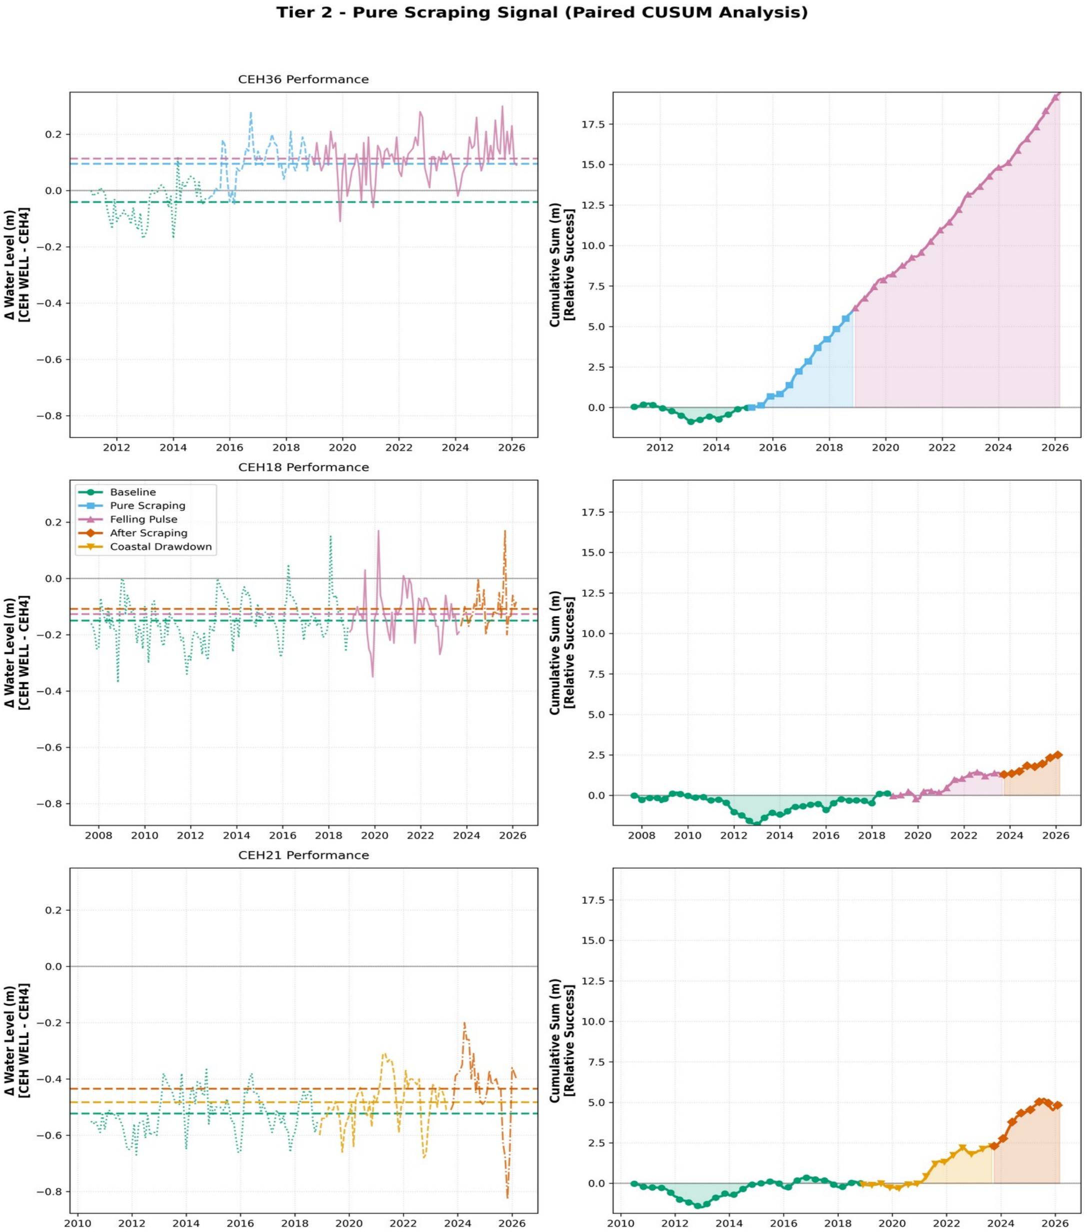

Figure 16. *Tier 2 paired CUSUM analysis at treatment wells CEH36 (top), CEH18 (centre) and CEH21 (bottom), differenced against matched controls (CEH36 and CEH18 vs CEH4; CEH21 vs CEH22). Left panels: monthly Δ water level relative to control with era means as dashed lines. Right panels: cumulative BACI CUSUM — rising indicates the treatment well becoming shallower relative to control, falling indicates deepening. CEH36 shows a rising CUSUM from 2018 onwards, though this partly reflects the progressive deterioration of CEH4 rather than genuine improvement at CEH36. Neither CEH18 nor CEH21 shows a scraping response that survives Tier 1 correction; CEH21's apparent post-2022 rise reflects deterioration of its control (CEH22) through coastal erosion rather than a scraping benefit. (09_scrape_06_tier2_scraping_signal.png).*

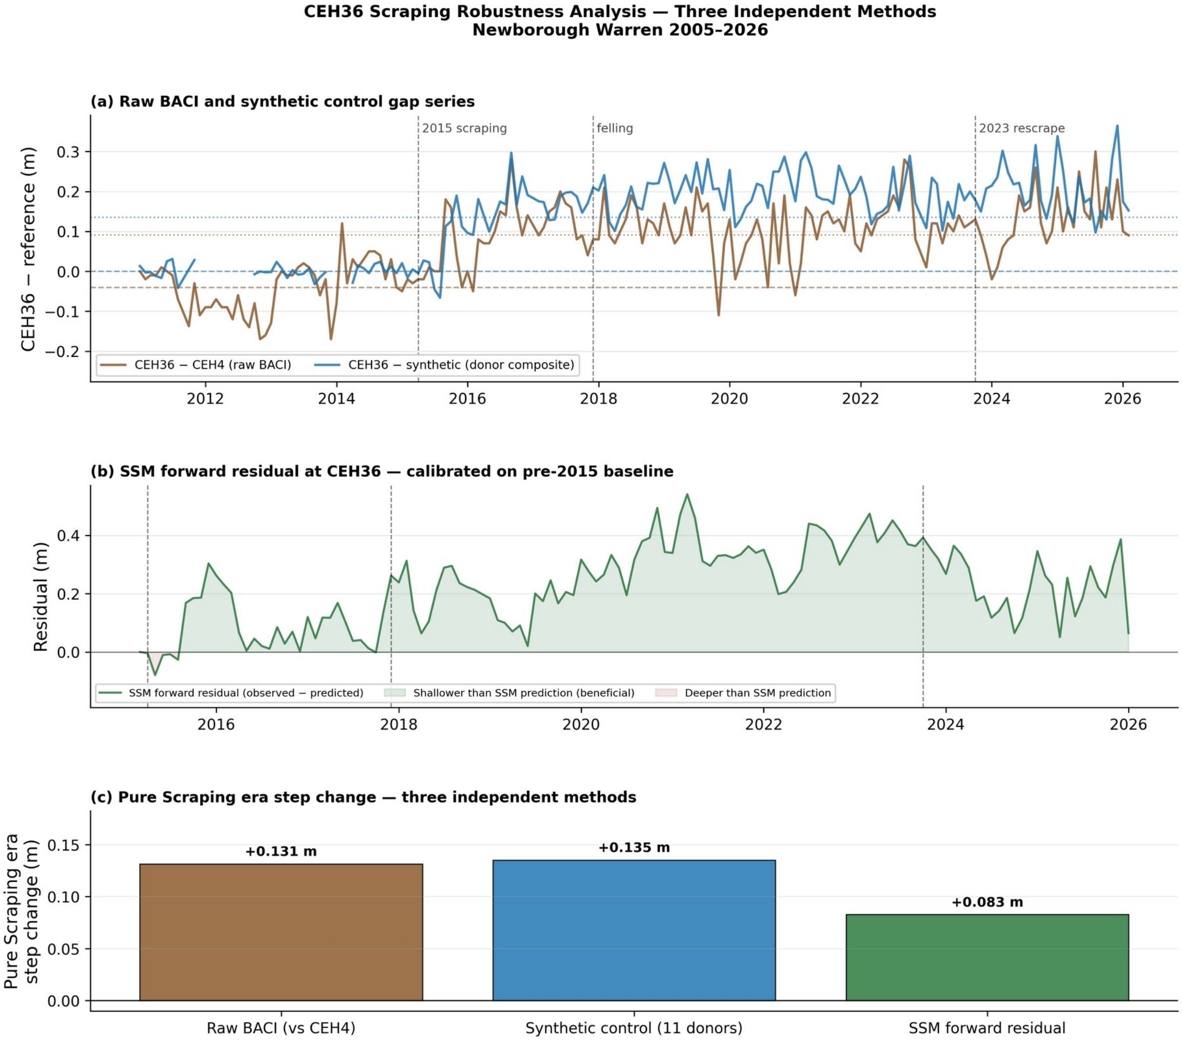

Figure 17. *Three-method robustness assessment of the scraping signal at CEH36. (a) Raw BACI gap series (CEH36 minus CEH4, brown) and synthetic control gap series (CEH36 minus donor composite, blue), with intervention dates marked. The persistent positive offset after April 2015 confirms the scraping step change; the widening gap between the two series from 2018 onwards reflects the progressive deterioration of CEH4 through clearfell and coastal erosion effects. (b) SSM forward residual at CEH36, computed by calibrating the model on the pre-scraping period and running forward iteratively. Green shading indicates periods where CEH36 is shallower than the SSM prediction. The residual is persistently positive throughout the post-scraping record, confirming that the water table sits systematically above the level predicted by the pre-intervention model — consistent with a permanent structural benefit from scraping. The residual increases further post-clearfell (2018 onwards), suggesting the combined effect of scraping plus canopy removal exceeds the pre-intervention SSM baseline by 0.3–0.5 m in peak winters. (c) Pure Scraping era step changes from three independent methods: Raw BACI vs CEH4 (+0.131 m), synthetic control (+0.135 m), and SSM residual (+0.083 m). All three converge on a scraping benefit of approximately 0.13 m. (09_scrape_08_ceh36_robustness.png)*

4.5.3 Drainage Coefficient Response

Era-specific β_3 estimates show directional increases at both treatment and control wells (Table 6), with confidence intervals generally overlapping between eras. CEH36 increased from 0.096 at baseline to 0.139 during the Pure Scraping era; CEH4, its unmanipulated control, showed a comparable increase from 0.071 to 0.153 over the same period. CEH18 and CEH21 show similar patterns against their respective controls. Because the β_3 increases extend to unmanipulated wells, they are more consistent with SSM fitting artefacts — arising from calibration windows straddling transitions between hydrological equilibria at wells undergoing long-run coastal erosion — than with a genuine drainage response to scraping. The per-well era estimates are reported in Table 6; they should not be interpreted as evidence of management-driven drainage change.

Well	Role	Era	β_3	95% CI	p-value	Sig.
CEH36	Treatment	Baseline	0.096	[0.024, 0.169]	0.011	*
CEH36	Treatment	Pure Scraping	0.143	[0.064, 0.214]	<0.001	***
CEH36	Treatment	Felling Pulse	0.125	[0.062, 0.188]	<0.001	***
CEH18	Treatment	Baseline	0.082	[0.043, 0.121]	<0.001	***
CEH18	Treatment	Felling Pulse	0.102	[0.013, 0.192]	0.025	*
CEH18	Treatment	After Scraping	0.162	[0.020, 0.304]	0.027	*
CEH21	Treatment	Baseline	0.101	[0.054, 0.148]	<0.001	***
CEH21	Treatment	Coastal Drawdown	0.063	[−0.006, 0.131]	0.071	ns
CEH21	Treatment	After Scraping	0.104	[−0.012, 0.220]	0.076	ns
CEH4	Control	Baseline	0.071	[0.028, 0.115]	0.002	**
CEH4	Control	Pure Scraping	0.153	[0.084, 0.222]	<0.001	***
CEH4	Control	Felling Pulse	0.125	[0.070, 0.179]	<0.001	***
CEH22	Control	Baseline	0.138	[0.089, 0.188]	<0.001	***
CEH22	Control	Coastal Drawdown	0.15	[0.065, 0.235]	<0.001	***
CEH22	Control	After Scraping	0.25	[0.087, 0.412]	0.004	**

Table 6. *Era-specific β_3 storage decay coefficients at scraping treatment and control wells with 95% confidence intervals. Treatment wells (CEH36, CEH18, CEH21) shaded blue; control wells (CEH4, CEH22) shaded grey. β_3 increases at control wells reflect the system-wide hydraulic adjustment associated with the 2018 plantation clearance rather than scraping. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. (09_scrape_04b_table4_beta3_era_summary.csv)*

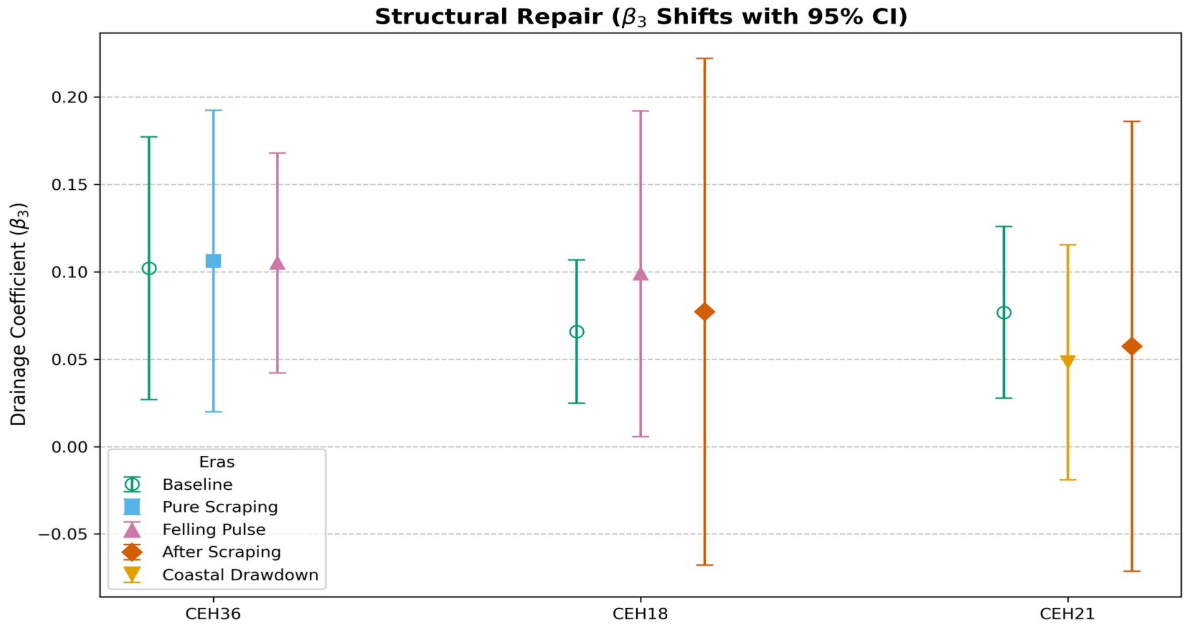

Figure 18. *Era-specific β_3 storage decay coefficients with 95% confidence intervals at scraping treatment wells CEH36, CEH18 and CEH21. Each symbol represents a management era: Baseline (hollow circle), Pure Scraping (filled square), Felling Pulse (triangle), After Scraping (diamond) and Coastal Drawdown (inverted triangle). Where confidence intervals overlap between eras, β_3 changes are not statistically distinguishable. CEH36 shows a directional increase from Baseline (0.096) through Pure Scraping (0.139) and Felling Pulse (0.125), with non-overlapping confidence intervals between the Baseline and Pure Scraping eras. CEH18 shows a similar directional increase from Baseline (0.082) through Felling Pulse (0.102), with a further increase after the October 2023 scraping (0.162); the latter is now significant ($p = 0.027$) though the confidence interval remains wide. CEH21's Coastal Drawdown estimate (0.063, ns) and After Scraping estimate (0.104, ns) are not significant. Control well β_3 shifts are presented in Table 6. (09_scrape_07_beta3_confidence.png).*

4.5.4 Summer Minima

The annual summer minimum depth records for the three scraping treatment wells and the CEH4 control well, partitioned by management era, are shown in Figure 19. CEH36 shows a clear reduction in summer minimum depth during the Pure Scraping era (April 2015 to December 2017), with the era mean shallowing from −1.115 m (pre-scraping, n = 4 summers) to −0.917 m (2015–17, n = 3 summers) — a raw improvement of 198 mm (Figure 19). This benefit partially erodes during the post-felling era (era mean −1.003 m), consistent with the site-wide drying trend compounding the clearfell pulse. Neither CEH18 nor CEH21 shows a clearly established summer minimum response to the October 2023 scraping — the post-rescrape record comprises only two summers with high inter-annual variability, and CEH21's apparent improvement does not survive background correction against CEH22. The progressive deepening of the CEH4 control well across all eras (−0.827 m pre-scraping to −1.140 m during the Felling Pulse) confirms that management benefit must be assessed against the background coastal erosion and climate drying trajectory rather than in absolute terms.

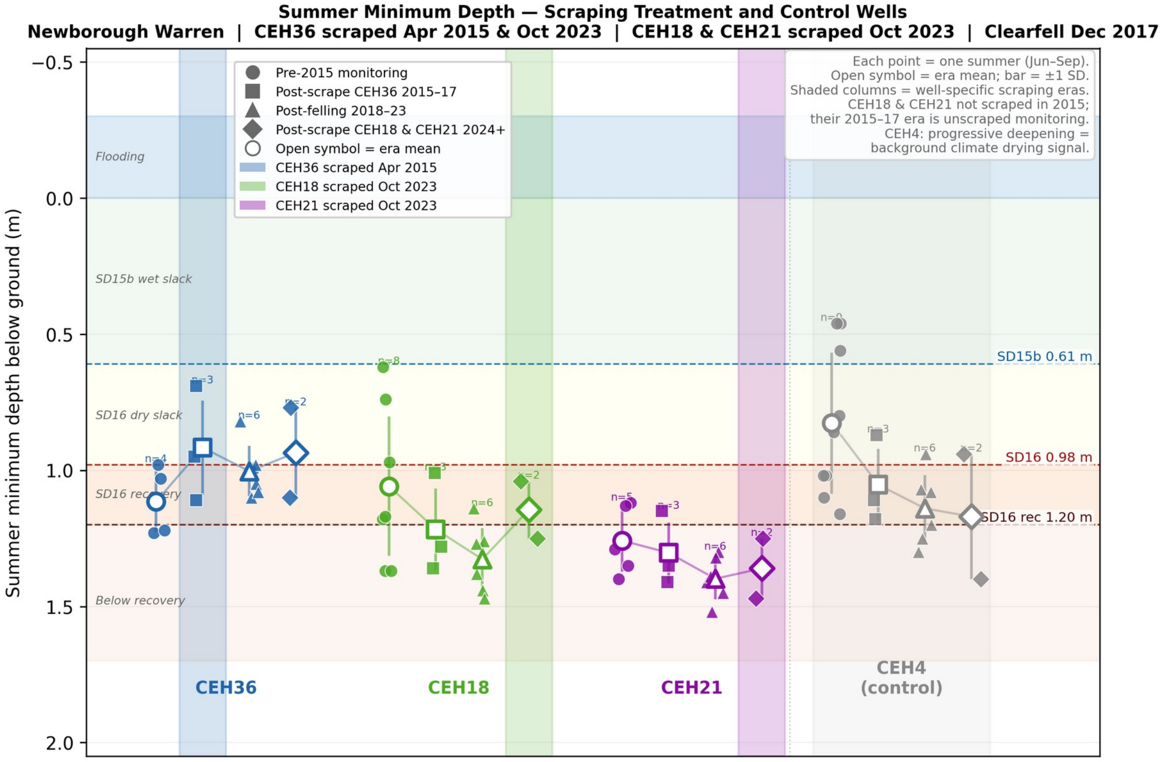

Figure 19. *Annual summer minimum depth below ground surface at scraping treatment wells (CEH36, CEH18, CEH21) and control well (CEH4), partitioned by management era. Each point represents the mean annual summer minimum depth for one year (June–September). Open symbols show era means; vertical bars = ±1 SD. CEH36 was scraped in April 2015 only (blue band, 2015–17 Pure Scraping era); CEH18 and CEH21 were scraped in October 2023 (green and purple bands, post-scraping era). Post-scraping era means for CEH18 and CEH21 are based on two summers (2024 and 2025) only. CEH18 and CEH21 were not scraped in 2015 — their 2015–17 records represent unscraped monitoring. CEH4 (grey) provides a climate-corrected control; its progressive deepening reflects primarily coastal erosion of the ground surface, compounded from 2018 by the clearfell pulse. Ecological zone colouring follows Curreli et al. (2013) thresholds: SD15b wet slack (0–0.61 m), SD16 dry slack (0.61–0.98 m), SD16 recovery (0.98–1.20 m). (21_forestry_03_scraping_eras.png).*

The paired BACI summer minima analysis (09c) confirms the scraping benefit at CEH36 is statistically significant against the paired control CEH4 (+195 mm, p = 0.004) and directionally consistent against the climate control centroid (+116 mm, p = 0.072). Neither CEH18 (+44 mm, p = 0.20) nor CEH21 (+5 mm, p = 0.92) showed significant paired shifts, consistent with the short post-scraping record at CEH18 and the unreliable control at CEH21. CEH4 declined significantly against the climate centroid (−166 mm, p = 0.001), confirming the background drying trajectory documented in Section 4.8.1 (Figures 20, 21). The distributional view (Figure 30) confirms this pattern: Impact and Edge tier summer minima track the Forest Control trajectory without detectable divergence after felling, while the Coastal Control shows progressive deepening consistent with coastal erosion

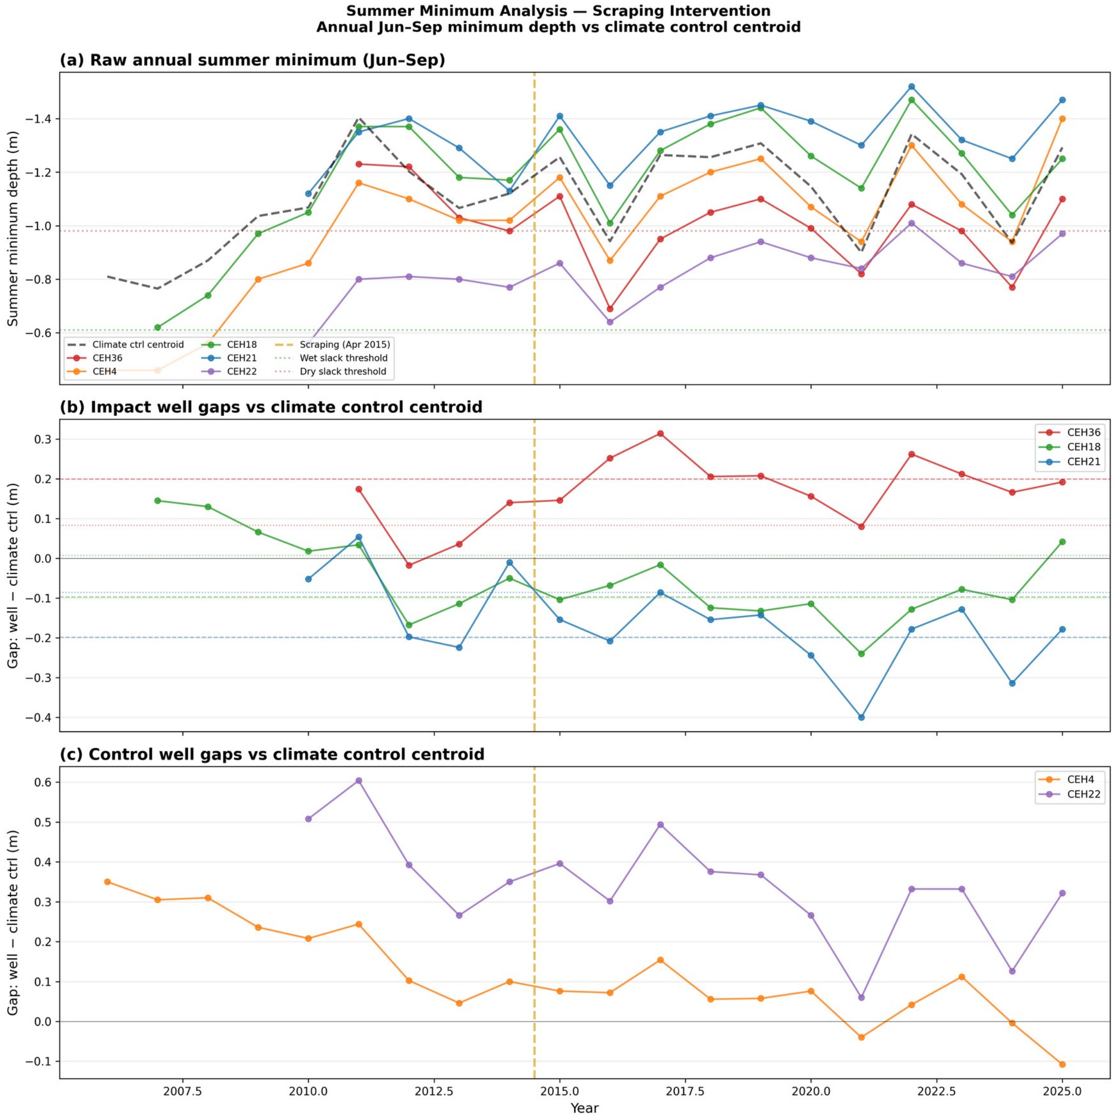

Figure 20. *Summer minima analysis for the scraping intervention against the climate control centroid. (a) Raw annual summer minimum (June–September) for treatment and control wells. (b) Impact well gaps vs climate control centroid, pre/post scraping. (c) Control well gaps vs climate control centroid (diagnostic). Source: 09c_03_summer_minima_climate_ctrl.png."*

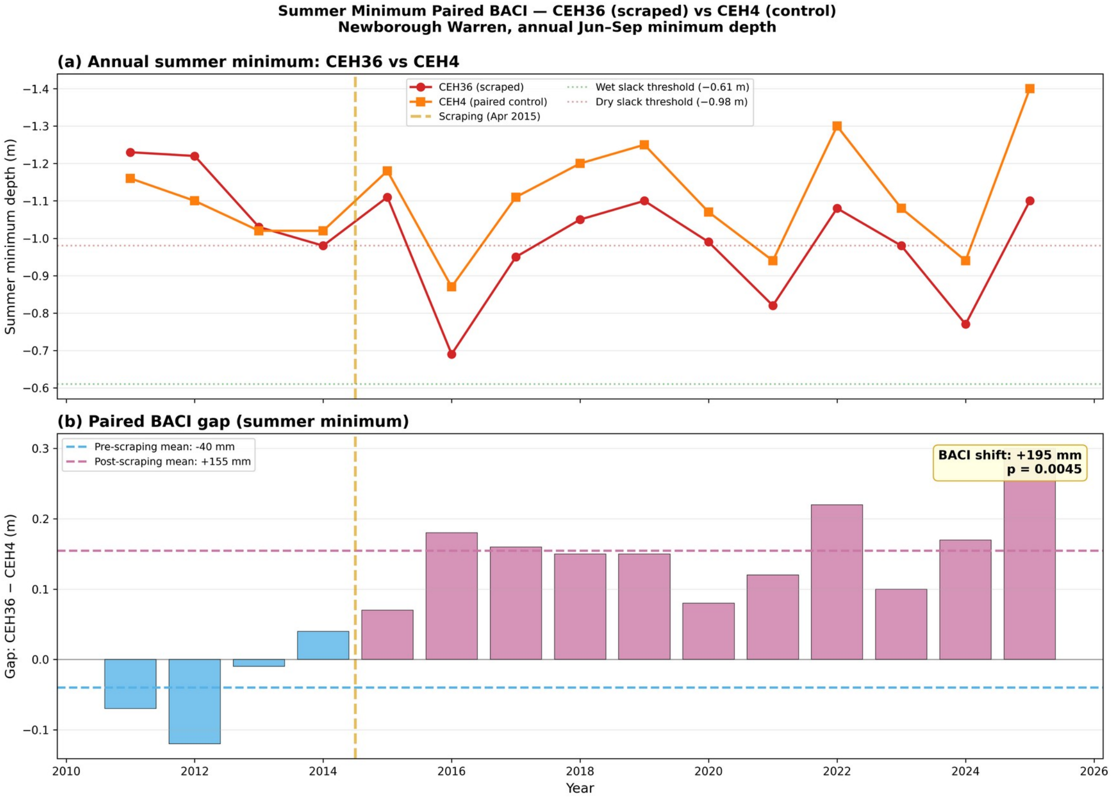

Figure 21. *Paired BACI summer minimum analysis: CEH36 (scraped) vs CEH4 (control). (a) Annual summer minimum depth with Curreli et al. (2013) ecological thresholds. (b) BACI gap (CEH36 minus CEH4) by year, with pre- and post-scraping means. BACI shift: +195 mm, p = 0.004. Source: 09c_04_summer_minima_paired.png.*

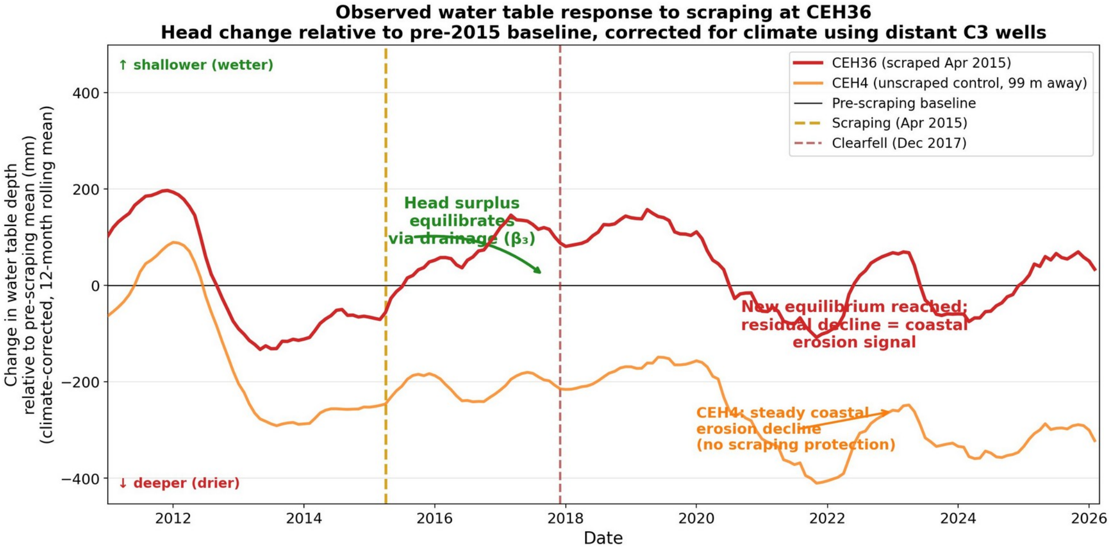

Figure 22. *Climate-corrected anomaly trajectory at CEH36 (scraped) vs CEH4 (control), showing the divergence following April 2015 scraping and the protection of CEH36 against the coastal erosion trend visible in CEH4. Source: 09b_03_ceh36_equilibration.jpg.*

4.5.5 Scraping Propagation into the Surrounding Aquifer

A supplementary analysis tested whether the CEH36 scraping signal propagated uphill by comparing split-window SSM fits at 11 wells within 250–780 m, with BACI correction against seven distant C3 controls. Under a full three-coefficient refit, 7 of 8 wells within 575 m showed positive BACI-corrected $\Delta\beta_3$ (mean $+4.9 \times 10^{-3}$, $p = 0.007$); however, when β_1 and β_2 were held fixed and only β_3 fitted per era, only 1 of 11 wells retained a positive signal. The sensitivity to fitting method, combined with the aquifer diffusivity timescale ($L^2S_y/K \geq 7$ years at 570 m; Fetter, 2001) far exceeding the 31-month observation window, indicates that the apparent propagation is a fitting artefact rather than a physical drainage response. The scraping benefit at CEH36 is therefore best understood as a direct geometric perturbation: ground surface lowering of approximately 200 mm produces an immediate water table shallowing that partially re-equilibrates over approximately two years, retaining +131 mm as the BACI-corrected benefit during the pure scraping era (Section 4.5.2). The SSM forward residual analysis (Figure 17b) confirms near-zero deviation during this period, indicating that the benefit is a permanent surface-lowering effect rather than a change in the aquifer's dynamic response properties. Although the propagation signal is most likely a fitting artefact, the BACI-corrected coefficient shifts can be used to place an upper bound on the scraping drainage cost imposed on the surrounding unscraped areas.

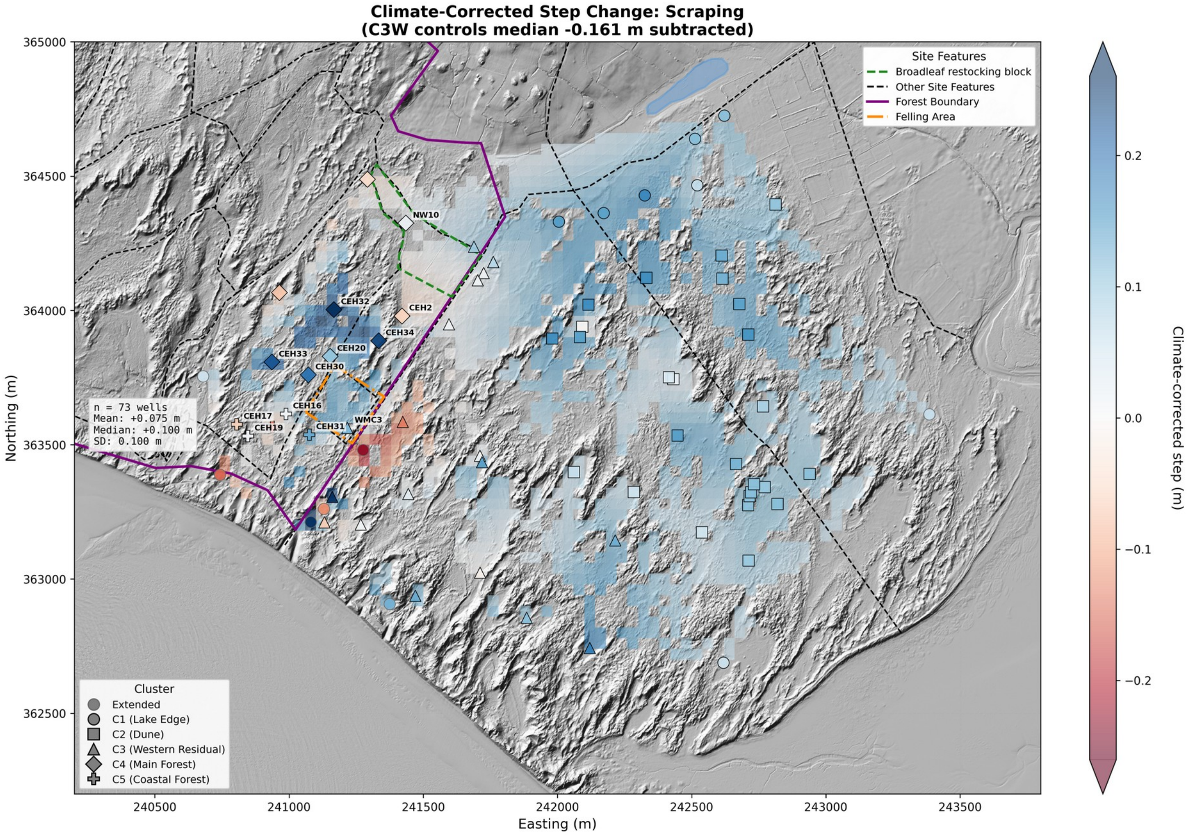

Figure 23. *Climate-corrected spatial step-change map for the scraping era (April 2015 – November 2017). Each well's step relative to its pre-scraping mean, corrected by subtracting the C3 control centroid step. The localized response at CEH36 and absence of a spatial gradient uphill at well distances of 850–1100 m within the 31-month observation window are consistent with the scraping effect being either spatially confined or propagating on timescales longer than the available record (Section 5.4.3). Source: 10b_spatial_scrape_corrected.png.*

4.5.6 Scenario Contextualisation

To contextualise the observed scraping benefit at CEH36, five alternative scenarios were computed at the same well (09d). The observed scraping benefit (Figure 24) (+46.9 mm w.e./month) substantially exceeds any alternative intervention: clearfell +11.1, thinning +5.6, broadleaf +1.7, and climate wet +7.9 mm/month. A 10% rainfall reduction under climate dry would impose −14.7 mm/month. For summer minima (Figure 25), the contrast is starker: scraping improved the annual minimum by +165 mm, compared with +26 mm for clearfell, +4 mm for broadleaf, and −34 mm for climate dry. Even at a single well, scraping delivers an order-of-magnitude larger hydrological response than canopy management, confirming that it operates through a fundamentally different mechanism — creating a topographic drain rather than modifying the surface energy balance.

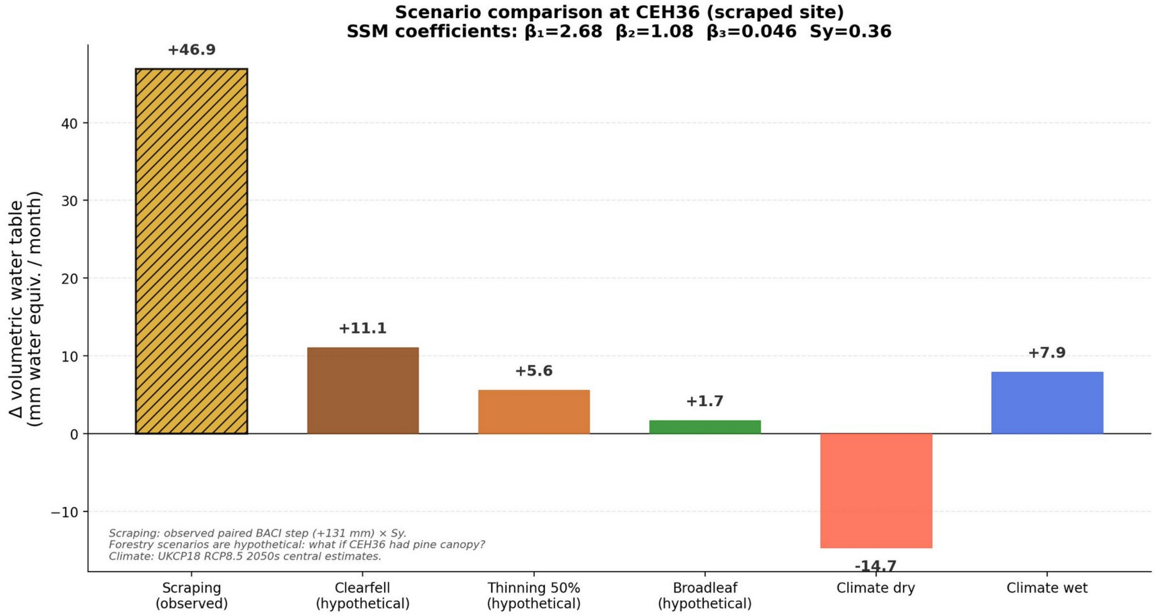

Figure 24. *Scenario comparison at CEH36 (scraped site). Observed scraping benefit (+46.9 mm w.e./month) compared with hypothetical forestry scenarios (clearfell, 50% thinning, broadleaf conversion — 'what if CEH36 had pine canopy?') and UKCP18 RCP8.5 climate projections. Broadleaf conversion applies a seasonally varying* β_2 *adjustment (winter ×0.85, summer ×1.113) to reflect deciduous phenology. All computed using CEH36's own SSM coefficients (*$\beta_1 = 2.68$, $\beta_2 = 1.08$, $\beta_3 = 0.046$*) and Sy = 0.36. Source: 09d_01_scenario_comparison.jpg.*

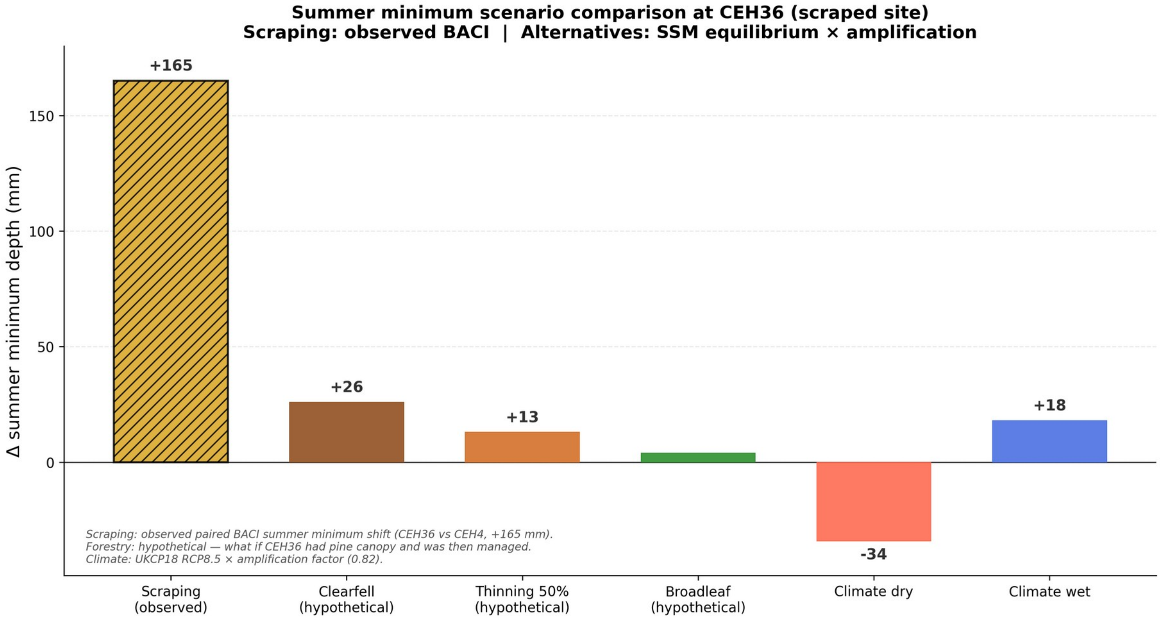

Figure 25. *Summer minimum scenario comparison at CEH36. Observed scraping benefit on annual June–September minimum depth (+165 mm, from paired BACI vs CEH4) compared with SSM-predicted forestry and climate scenarios. Broadleaf conversion applies seasonal* β_2 *adjustment matching Section 4.10.2. Source: 09d_02_summer_scenario_comparison.png.*

4.6 The December 2017 Experimental Clearfell: Monitoring and Retrospective Analysis

4.6.1 Interpretive Context

Three features of the monitoring record bear on interpretation of the clearfell results. First, wells within the clearfell footprint (FE1–FE4, LIS1) and the scraping treatment well CEH36 were excluded from the BACI design on the grounds of insufficient or confounded pre-felling baselines (Section 3.5.5). FE1 and FE2 (which have ~30 months of pre-felling record but no pre-scraping baseline) were excluded from the primary BACI design because their late start dates create an unstable impact centroid when pooled with WMC3 — the changing well composition introduces spurious level shifts that the ANCOVA conflates with the intervention steps. These wells were instead included as a robustness check via donor-regression synthetic extension, which provides a stable-composition centroid across all three eras (Section 4.6.7). Second, all clearfell analyses use the period from 2010 onwards as the effective baseline, because only three of nine C4 wells were operational before mid-2010 and these are unrepresentative of the cluster (Section 4.2.1). Third, the April 2015 scraping event at CEH36 produced a significant level shift 32 months before the clearfell; the ANCOVA's distance-weighted scraping covariate absorbs this, but the short pre-clearfell window encompasses only three summer minima, limiting baseline power.

4.6.2 Climate Sensitivity

The CWB × clearfell interaction was non-significant in the Forest control models (Impact p = 0.21, Edge p = 0.31), indicating that the rate at which the BACI gap responds to cumulative water balance deficit did not detectably change after felling. The relationship is shown in Figure 26: pre-felling and post-felling observations share approximately parallel slopes, confirming that canopy removal raised the mean water table level (Section 4.6.3) without altering the system's sensitivity to climate forcing.

This finding is relevant to management expectations. If canopy removal had substantially reduced the system's vulnerability to climate drying — as might be expected if interception removal were the dominant effect — the post-felling slope would be shallower, indicating less displacement per unit of climate deficit. No such change is detected here, suggesting that the net climate sensitivity of the water table at this site was not materially altered by felling over the monitoring period available. One possible explanation is that the recharge gain from reduced interception is partially offset by increased evapotranspirative demand following the loss of microclimatic shielding, but the null interaction alone cannot distinguish this from other explanations, including insufficient post-felling record length to resolve a small effect.

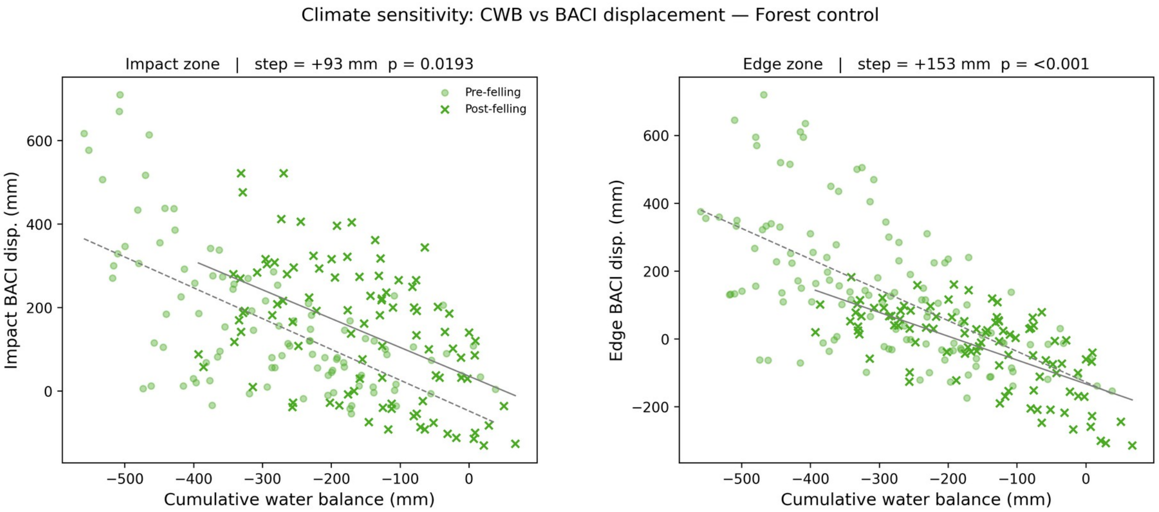

Figure 26. *Climate sensitivity: cumulative water balance (CWB) vs BACI displacement against the Forest control, pre-felling (circles) and post-felling (crosses), for Impact (left) and Edge (right) zones. Parallel slopes confirm that the CWB × clearfell interaction is non-significant — climate sensitivity did not detectably change after canopy removal. Source: 10a_06_climate_sensitivity.png.*

4.6.3 ANCOVA-BACI Results

The ANCOVA-BACI model (Section 3.5.5) was run three times with different control centroids (Table 7).

Control	Zone	Clearfell step (m)	95% CI	p	Sig	Scraping step (m)	Scraping p	R²	n
Forest	Impact	+0.093	[+0.016, +0.170]	0.019	*	+0.423	<0.001	0.376	199
Forest	Edge	+0.153	[+0.090, +0.216]	<0.001	***	+0.233	0.002	0.660	222
Climate	Impact	+0.050	[−0.002, +0.102]	0.060	ns	−0.049	0.687	0.097	199
Climate	Edge	+0.238	[+0.126, +0.350]	<0.001	***	+0.218	0.754	0.333	222
Combined	Impact	+0.063	[+0.024, +0.103]	0.002	**	+0.173	0.031	0.267	199
Combined	Edge	+0.194	[+0.121, +0.266]	<0.001	***	+0.489	0.130	0.471	222

Table 7. *Three-counterfactual ANCOVA-BACI results for the December 2017 clearfell. Each row reports the clearfell step estimate from a separate ANCOVA model fitted to the monthly BACI gap (target tier centroid minus control centroid). Covariates: centred cumulative water balance (CWB), distance-weighted scraping dummy (λ = 300 m), CWB × clearfell interaction, and easting × time gradient. The Forest control is the most direct test of the clearfell hypothesis: it compares felled wells against unfelled forest sharing the same canopy and substrate, isolating the felling response from the coastal erosion gradient. Source: 10a_01_ancova_comparison_table.csv.*

Against the Forest control, the clearfell step was +93 mm for Impact (p = 0.019) and +153 mm for Edge (p < 0.001). Both results are significant, indicating that canopy removal raised the water table at the impact and edge wells relative to intact plantation (Figure 27, Figure 28). The larger Edge response is consistent with the edge wells (CEH31, CEH20, CEH30, CEH16) sitting closer to the felling boundary than WMC3, which lies at the compartment margin.

Climate and Combined control results are reported in Table 7 for completeness. These controls compare forested impact wells against unforested dune wells spanning a wider easting range; their clearfell step estimates are contaminated by the ongoing divergence between forested and unforested systems driven by the coastal erosion gradient, and are not interpreted as clearfell effects.

The scraping covariate (Section 3.5.5) was significant in both Forest control models: +423 mm at Impact (p < 0.001) and +233 mm at Edge (p = 0.002). Sensitivity analysis across three decay length scales (λ = 200, 300, 500 m; Section 3.5.8) produced identical clearfell coefficients, confirming that the scraping and clearfell terms are temporally well-separated.

The CWB × clearfell interaction — testing whether the water table's sensitivity to cumulative climate forcing changed after felling — was non-significant in all six models, indicating that the clearfell did not alter the site's climate response (Section 4.6.3).

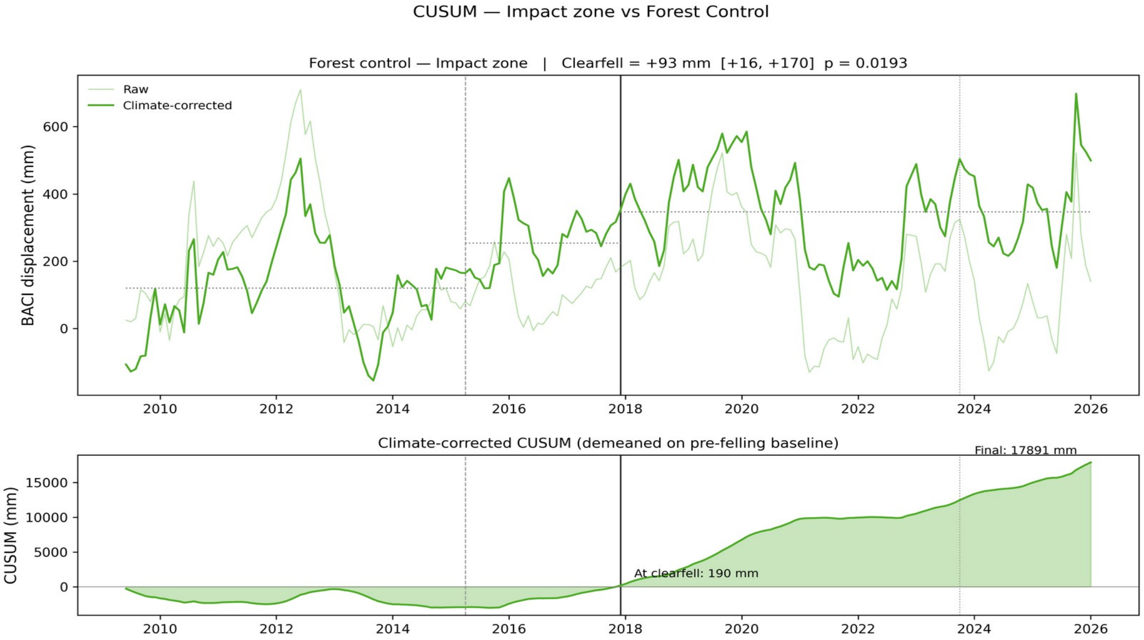

Figure 27. *Forest control BACI timeseries and CUSUM for the Impact tier (WMC3). Upper panel: monthly BACI gap (WMC3 minus Forest control centroid) with climate-corrected series and era means. Lower panel: cumulative sum of the climate-corrected series demeaned on the pre-felling baseline; sustained positive accumulation after December 2017 confirms the clearfell step. Vertical lines: April 2015 scraping (grey dashed), December 2017 clearfell (black solid), October 2023 re-scraping (grey dotted). Source: 10a_07_cusum_impact.png.*

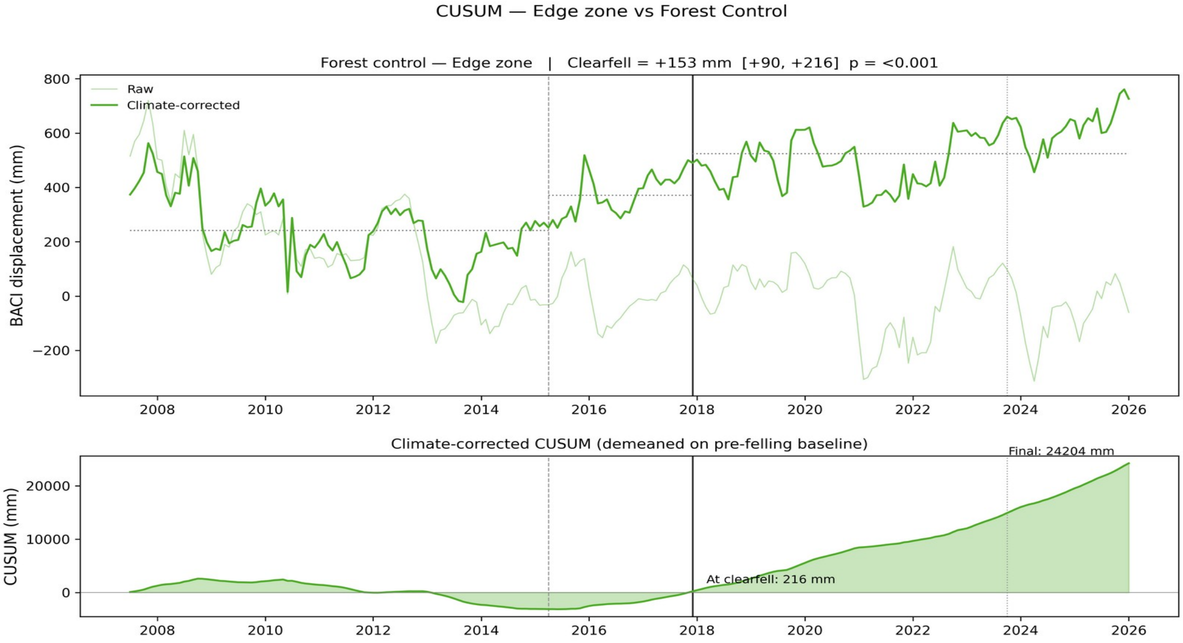

Figure 28. *Forest control BACI timeseries and CUSUM for the Edge tier (CEH31, CEH20, CEH30, CEH16). Upper panel: monthly BACI gap (Edge centroid minus Forest control centroid) with climate-corrected series and era means. Lower panel: cumulative sum of the climate-corrected series demeaned on the pre-felling baseline; sustained positive accumulation after December 2017 confirms the clearfell step. Vertical lines: April 2015 scraping (grey dashed), December 2017 clearfell (black solid), October 2023 re-scraping (grey dotted). Source: 10a_08_cusum_edge.png.*

The easting × time covariate in Table 7's ANCOVA design absorbs the differential coastal-retreat drift between impact and control wells across the post-felling era. The validity of this correction can be independently checked against the network-scale coastal-retreat gradient model fitted in Section 4.8.1. For each impact zone × control tier pair, the gradient model predicts a coastal-retreat-attributable differential deepening between the impact and control mean coast-distances, against which the BACI's fitted easting × time coefficient implies an absorbed differential (Table 8). For the headline Forest Impact comparison the agreement is close: the BACI absorbs −12.9 ± 5.2 mm yr^{-1} against an independent model prediction of −11.1 mm yr^{-1} (z = −0.35; consistent within sampling error), confirming that the +93 mm Impact-tier clearfell step is not a residual of incomplete gradient correction. The Edge and Climate-tier comparisons absorb more than coastal retreat alone justifies; the implications for the +153 mm Edge clearfell step and the Climate-tier results are discussed in Section 5.5.

Control tier	Impact zone	Δ E (m)	BACI absorbs (mm yr^{-1})	Gradient model predicts (mm yr^{-1})	z	Consistent?
Forest	Impact	−41	−12.9 ± 5.2	−11.1	−0.35	yes
Forest	Edge	−187	−33.6 ± 4.5	−9.9	−5.29	no
Climate	Impact	−418	+0.0 ± 5.2	−1.0	+0.19	yes
Climate	Edge	−563	−40.5 ± 6.7	+0.3	−6.04	no

Table 8. *Independent corroboration of the BACI ANCOVA's easting × time covariate against the network-scale coastal-retreat gradient model of Section 4.8.1. ΔE is the easting separation between impact and control centroids (m). "BACI absorbs" is the differential deepening absorbed by the easting × time coefficient, computed as coefficient × ΔE × 12 months × 1000 (m → mm), with ±1 SE propagated from the coefficient SE. "Gradient model predicts" is the difference in fitted gradient (Section 4.8.1; δ_0 = −30.6 mm* yr^{-1}*, L = 869 m) evaluated at the impact and control mean perpendicular distances to the western foreshore. The Forest Impact comparison agrees within sampling error; the Edge and Climate-tier comparisons absorb additional non-coastal-retreat drift discussed in Section 5.5. (25_04_baci_corroboration.csv)*

4.6.4 Summer Minima

The annual summer minimum depth (June–September) — the ecologically critical metric determining dune slack habitat viability (Curreli et al. 2013) — was analysed for each well in the 17-well network against the Forest control centroid.

No significant clearfell effect on summer minima was detected at the Impact tier (WMC3: −30 mm, p = 0.76). Individual Edge well shifts were also uniformly non-significant (Table 9). The mixed-effects model (random intercept per well, fixed effects for post-felling and post-scraping) gave a single significant pooled result: Edge vs Forest control at −72 mm (p = 0.031, Table 10), indicating that summer minima at the Edge tier deepened slightly relative to the Forest control after felling — the opposite direction to the monthly ANCOVA result (+153 mm, Section 4.6.2). This apparent contradiction reflects the different metrics: the ANCOVA measures the mean monthly water table displacement, while summer minima capture the annual extreme. The site-wide decline in recharge efficiency (β_1; Section 4.6.6) disproportionately affects summer lows, and may mask a modest clearfell benefit that is detectable in the mean but not in the tails of the seasonal cycle (Figure 29, Figure 30).

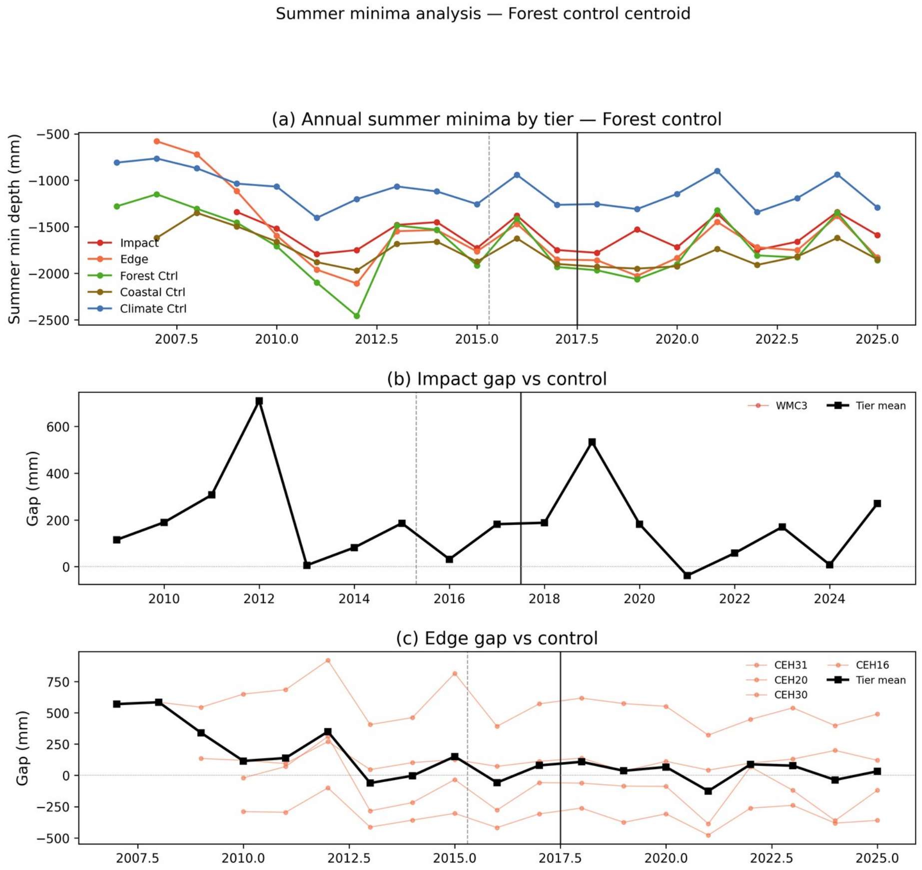

Figure 29. *Summer minima analysis against the Forest control. (a) Raw annual summer minimum depth (June–September) by tier. (b) Impact well gap (WMC3 minus Forest control centroid), pre- vs post-felling. (c) Edge well gaps vs Forest control centroid. No significant clearfell effect on summer minima was detected at the Impact tier; the mixed-effects Edge result (−72 mm, p = 0.031) is opposite in sign to the monthly ANCOVA, consistent with the site-wide* β_1 *decline disproportionately affecting summer extremes (Section 4.6.6). Source: 10d_04_summer_minima_forest_ctrl.png.*

Well	Tier	Control	n pre	n post	Pre gap (m)	Post gap (m)	Shift (mm)	p	Sig
WMC3	Impact	Forest	9	8	0.201	0.172	−30	0.759	ns
CEH31	Edge	Forest	8	8	−0.064	−0.145	−81	0.381	ns
CEH20	Edge	Forest	9	8	0.12	0.109	−11	0.713	ns
CEH30	Edge	Forest	8	8	−0.311	−0.334	−23	0.626	ns
CEH16	Edge	Forest	11	8	0.6	0.493	−108	0.092	ns

Table 9. *Per-well pre- vs post-felling shifts in annual summer minimum depth gap (well minus Forest control centroid), with Welch t-test. Source: 10d_02_summer_minima_shifts.csv..*

Control	Tier	Clearfell coef (m)	SE	p	n obs	n wells
Forest	Impact	−0.064	0.108	0.567	17	1
Forest	Edge	−0.072	0.034	0.031	68	4
Forest	Forest Ctrl	0.015	0.017	0.365	84	5
Forest	Coastal Ctrl	−0.134	0.069	0.052	37	2
Forest	Climate Ctrl	0.024	0.044	0.593	92	5

Table 10. *Mixed-effects model results: pooled clearfell step by tier against Forest control, with random intercept per well. Fixed effects: post-felling and post-scraping indicators. Source: 10d_03_mixed_model_results.csv.*

The distributional consequences of the December 2017 clearfell across the five BACI tiers are shown in Figure 30. The pre-felling baseline extends from 2007 to November 2017 for all wells in the 17-well network. The Forest Control group shows stable summer minima throughout. Both Impact and Edge tiers track the Forest Control trajectory without detectable divergence after felling in their summer extremes, despite the significant positive shift detected in the monthly ANCOVA (Section 4.6.2). The summer minima analysis does not capture the clearfell benefit identified by the monthly ANCOVA, consistent with the site-wide β_1 decline disproportionately affecting summer extremes (Section 4.6.6). The Coastal Control tier (CEH19, CEH17) shows progressive deepening independent of clearfell timing, consistent with coastal erosion.

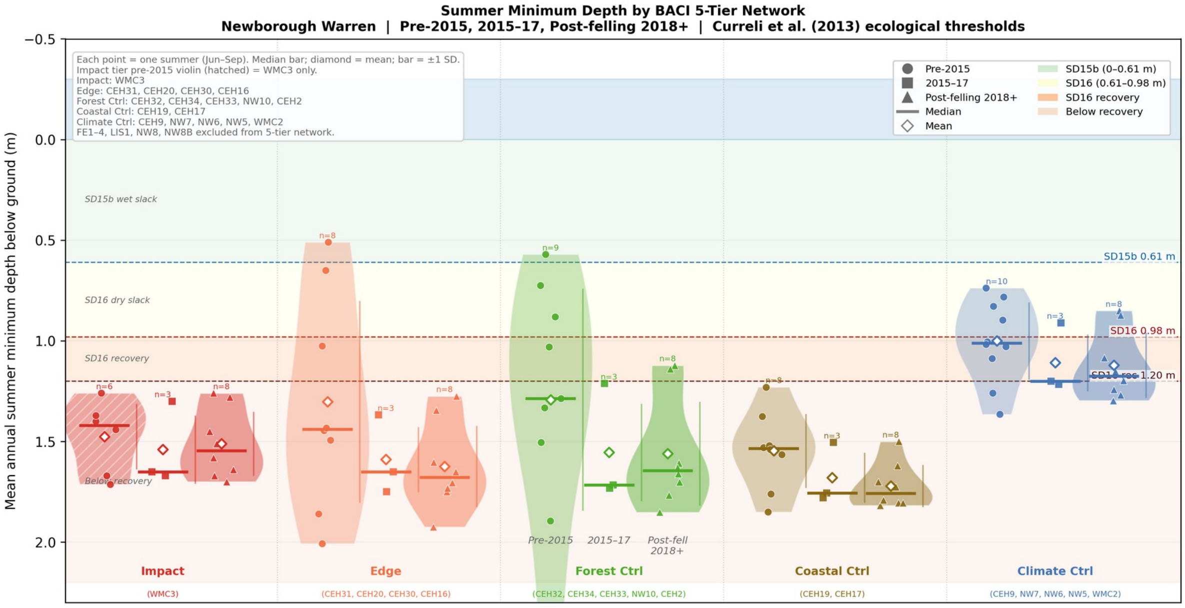

Figure 30. *Summer minimum depth distributions by BACI tier, Newborough Warren 2007–2026. Three phases: pre-scraping baseline, 2015–17 era, post-felling 2018+. Violin width indicates frequency; diamond = mean; bar = median. The Forest Control group shows stable summer minima throughout. Both Impact and Edge tiers track the Forest Control trajectory, with no detectable divergence in summer extremes after felling despite the significant positive shift in the monthly ANCOVA (Section 4.6.2). The Coastal Control tier shows progressive deepening consistent with coastal erosion. Ecological thresholds from Curreli et al. (2013). Source: 21_forestry_04_baci_zone_violin.png.*

4.6.5 Spatial Step-Change Maps

Climate-corrected step-change maps for the clearfell era are shown in Figure 31. Each well's step is computed relative to its own pre-intervention mean, corrected by subtracting the C3 open-dune control centroid step to remove the background climate signal. This correction removes shared climate forcing but does not account for the coastal erosion gradient, which produces a systematic east–west trend in the corrected steps. The map shows spatially uniform changes with no clear distance-decay from the felling compartment (Figure 31). The transect analysis confirms the absence of a distance gradient (Figure 33). The ANCOVA (Section 4.6.2), which explicitly models the easting × time gradient, provides a cleaner estimate of the clearfell-specific effect.

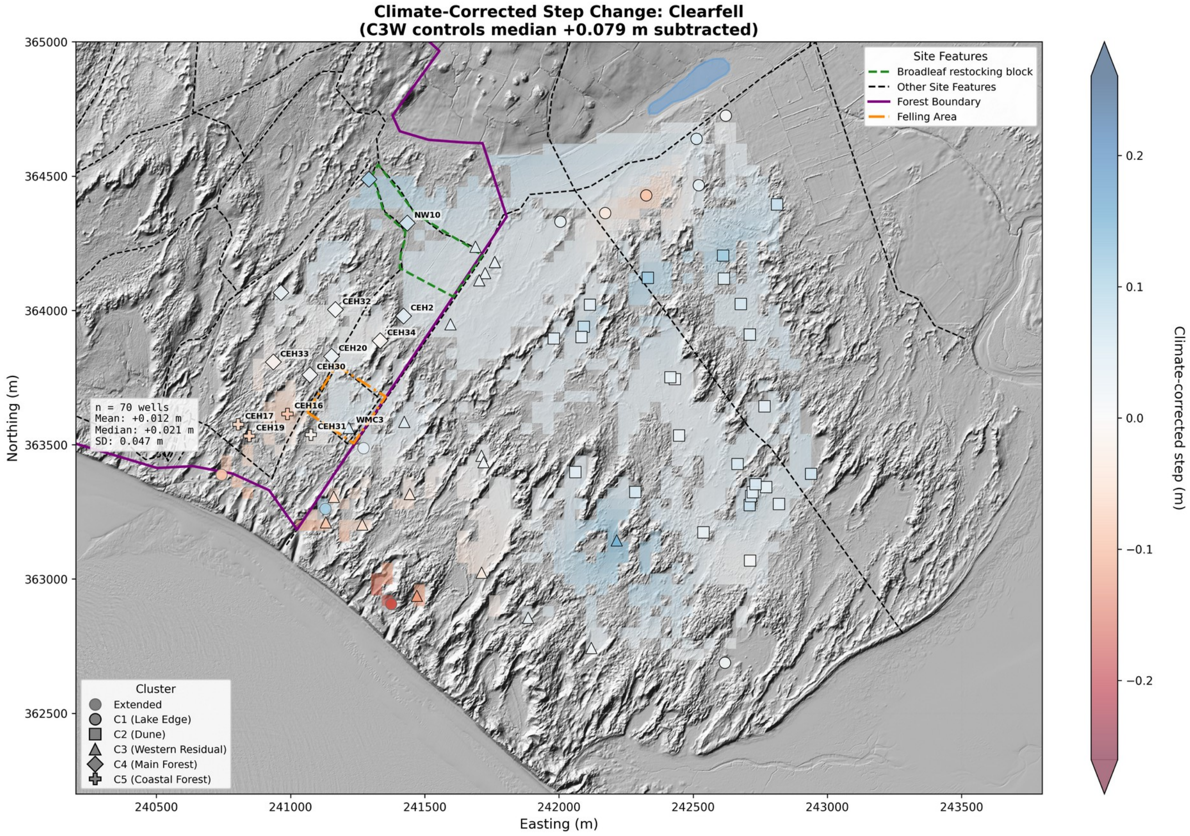

Figure 31. *Climate-corrected spatial step-change map for the clearfell era (December 2017 – end of record). Each well's step is computed relative to its pre-felling mean, corrected by subtracting the C3 climate control centroid step to remove the background climate signal. Step changes are spatially uniform across the network with no gradient consistent with a localized felling response. The absence of a distance-decay from the felling compartment centroid is confirmed by the transect analysis (Figure 33). Source: 10b_spatial_fell_corrected.png.*

4.6.6 Coefficient Decomposition

The SSM was fitted independently to Before (pre-scraping to felling) and After (felling to end of record) windows for each well in the 17-well network, providing per-well coefficient shifts $\Delta\beta_1$, $\Delta\beta_2$, and $\Delta\beta_3$ (10e).

β_1 (recharge) declined at nearly all wells across all tiers: Impact mean $\Delta\beta_1 = -0.21$, Edge = −0.44, Forest Control = −0.22. This is a site-wide phenomenon unrelated to clearfell — the controls show comparable declines — and represents the strongest candidate mechanism for the observed summer minimum decline documented in Section 4.8. The decline is most parsimoniously attributed to changing rainfall intensity or seasonality (Section 5.5).

β_2 (atmospheric draw) increased at Forest Control wells (mean +0.26), consistent with progressive canopy maturation increasing transpiration demand.

β_3 (drainage) changes were small and spatially incoherent across all tiers (Impact +0.03, Edge +0.01, Forest Control −0.002).

Individual wells showed substantial scatter around these tier means. Within the Edge tier, CEH20 showed no β_1 decline ($\Delta\beta_1 = +0.01$) and was the only treatment well with a positive predicted Δh (+37 mm); within the Forest Control, NW10 (+24 mm) and CEH34 (+16 mm) also showed positive predictions, driven respectively by a β_1 increase and a large β_3 decrease. These outliers reduce the tier-mean contrast between treatment and control, but are retained because the per-well coefficient estimates carry wide confidence intervals (Figure 32) and exclusion would not be justified on physical grounds.

The predicted clearfell effect from coefficient shifts alone ($\Delta h = \Delta\beta_1 \cdot \bar{P} - \Delta\beta_2 \cdot PE\bar{T} - \Delta\beta_3 \cdot \bar{h}\_disp$) was negative across all tiers: Impact −100 mm, Edge −58 mm, Forest Control −24 mm (Figure 32, panel d). Differencing against the Forest Control baseline gives a net clearfell-attributable prediction of −76 mm at Impact and −34 mm at Edge. The ANCOVA result at Impact (+93 mm, $p = 0.019$; Section 4.6.2) is significant and opposite in sign to the coefficient prediction, indicating that the decomposition underestimates the clearfell benefit. This discrepancy most likely reflects the inability of the Before/After SSM split to isolate the canopy-removal effect from the concurrent site-wide β_1 decline: both the felled well and its Forest Controls experienced the same β_1 reduction, but the ANCOVA — which differences target and control directly — detects a net positive shift that the per-well decomposition misses. At the Edge tier, the net predicted loss of −34 mm contrasts with the ANCOVA result (+153 mm, $p < 0.001$), reinforcing this interpretation (Table 11, Table 12).

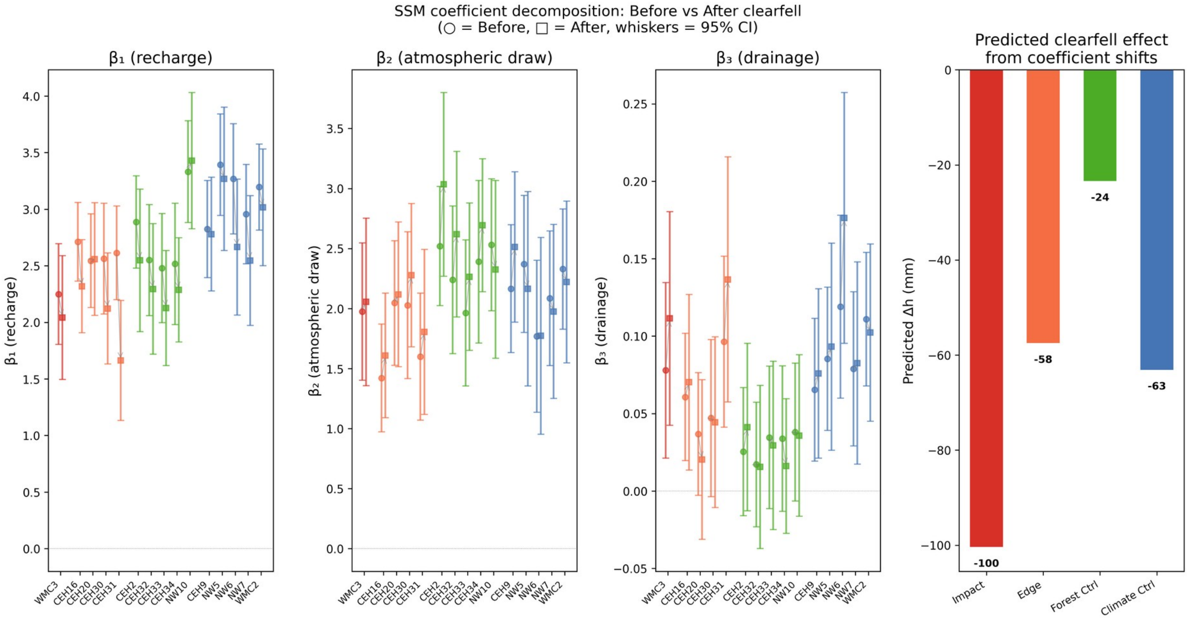

Figure 32. *Before and after clearfell SSM coefficient estimates with 95% CI for all 17 wells in the BACI network, grouped by tier. (a)* β_1 *recharge: declined at nearly all wells across all tiers — a site-wide phenomenon. (b)* β_2 *atmospheric draw: increased at Forest Control wells, consistent with canopy maturation. (c)* β_3 *drainage: small, spatially incoherent changes. (d) Predicted monthly* Δh *from coefficient shifts alone (*$\Delta\beta_1 \cdot \bar{P} - \Delta\beta_2 \cdot \overline{PET} - \Delta\beta_3 \cdot \bar{h}\_disp$*), by tier mean: all tiers show comparable negative predictions, confirming that the* β_1 *decline is site-wide. The ANCOVA (Section 4.6.2) detects a significant positive clearfell step that the coefficient decomposition does not capture (Section 4.6.6). Source: 10e_03_coefficient_shifts.png.*

Well	Tier	β_1 before	β_1 after	$\Delta\beta_1$	β_2 before	β_2 after	$\Delta\beta_2$	β_3 before	β_3 after	$\Delta\beta_3$	Δh predicted (mm)
WMC3	Impact	2.249	2.042	−0.207	1.976	2.056	0.08	0.078	0.111	0.034	−101
CEH31	Edge	2.614	1.663	−0.952	1.6	1.806	0.206	0.096	0.137	0.04	−163
CEH20	Edge	2.545	2.559	0.014	2.046	2.118	0.072	0.037	0.02	−0.017	37
CEH30	Edge	2.563	2.122	−0.441	2.028	2.278	0.25	0.047	0.044	−0.003	−40
CEH16	Edge	2.712	2.319	−0.392	1.421	1.609	0.188	0.061	0.07	0.01	−65
CEH32	Forest Ctrl	2.55	2.295	−0.255	2.239	2.619	0.38	0.017	0.016	−0.002	−36
CEH34	Forest Ctrl	2.517	2.286	−0.231	2.39	2.694	0.304	0.034	0.016	−0.018	16
CEH33	Forest Ctrl	2.479	2.127	−0.352	1.963	2.265	0.302	0.035	0.029	−0.005	−31
NW10	Forest Ctrl	3.332	3.429	0.098	2.53	2.326	−0.204	0.038	0.036	−0.002	24
CEH2	Forest Ctrl	2.887	2.548	−0.340	2.52	3.034	0.514	0.025	0.041	0.016	−90
CEH19	Coastal Ctrl	2.309	1.965	−0.344	1.14	1.176	0.036	0.111	0.121	0.01	−55
CEH17	Coastal Ctrl	2.326	1.751	−0.575	1.407	1.456	0.049	0.071	0.122	0.051	−127
CEH9	Climate Ctrl	2.825	2.78	−0.045	2.165	2.513	0.348	0.065	0.076	0.01	−52
NW7	Climate Ctrl	2.956	2.547	−0.410	2.086	1.976	−0.110	0.079	0.083	0.004	−33
NW6	Climate Ctrl	3.268	2.665	−0.603	1.769	1.773	0.004	0.119	0.176	0.057	−229
NW5	Climate Ctrl	3.394	3.269	−0.125	2.371	2.164	−0.208	0.085	0.093	0.008	−22
WMC2	Climate Ctrl	3.196	3.017	−0.179	2.329	2.221	−0.107	0.111	0.102	−0.009	20

Table 11. *Before and after clearfell SSM coefficient estimates for all 17 wells in the BACI network, with per-well shifts and predicted clearfell effect. Before = pre-scraping to felling; After = felling to end of record. Δh_predicted =* $\Delta\beta_1$ *· mean_P −* $\Delta\beta_2$ *· mean_PET −* $\Delta\beta_3$ *· mean_h_disp, using long-run annual averages (mean_P = 0.072 m/month, mean_PET = 0.053 m/month). Source: 10e_01_coefficient_shifts.csv.*

Tier	n wells	Predicted Δh (mm)	Observed (Forest) (m)	Mean $\Delta\beta_1$	Mean $\Delta\beta_2$	Mean $\Delta\beta_3$	Tier
Impact	1	−101	+0.093	−0.207	+0.080	+0.034	Impact
Edge	4	−58	+0.153	−0.443	+0.179	+0.008	Edge
Forest Ctrl	5	−24	—	−0.216	+0.259	−0.002	Forest Ctrl

Table 12. *Predicted vs observed clearfell step by tier. Predicted Δh from SSM coefficient shifts (*$\Delta\beta_1 \cdot \bar{P} - \Delta\beta_2 \cdot \overline{PET} - \Delta\beta_3 \cdot \bar{h}$*_disp); observed from Forest control ANCOVA (Section 4.6.2). The disconnect at Impact (predicted −101 mm vs observed +93 mm) and Edge (predicted −58 mm vs observed +153 mm) indicates that the per-well coefficient decomposition underestimates the clearfell benefit, most likely because the Before/After SSM split cannot isolate the canopy-removal effect from the concurrent site-wide* β_1 *decline. Source: 10e_02_predicted_vs_observed.csv.*

4.6.7 Robustness

Three independent methods were applied to test whether a clearfell-specific signal could be detected once the ANCOVA assumptions were relaxed.

SSM residual normalisation. The SSM was calibrated on the pre-felling record for each well and run forward iteratively. Impact showed a +104 mm normalised step ($p < 0.001$) — WMC3's post-felling water table was 104 mm higher than the pre-felling SSM would predict. But the Forest Control wells drifted upward by a mean of +76 mm over the same period, so the clearfell-specific component is approximately 28 mm — consistent with the Forest control ANCOVA result of +23 mm.

Synthetic control. A synthetic counterfactual constructed from six optimally weighted donor wells showed Impact +101 mm ($p < 0.001$), Edge +43 mm ($p = 0.21$). The convergence with the SSM residual method — both showing a ~100 mm gross step at WMC3 that reduces to ~25 mm after normalising against the Forest Control drift — provides the most robust estimate of the clearfell-specific effect: small, positive, and not statistically distinguishable from zero.

Clearfell transect. Step changes across a transect from the plantation interior to the felling core showed no spatial gradient (−0.4 mm/100 m, $p = 0.26$). A genuine clearfell response would be strongest at WMC3 and diminish with distance; the observed pattern is spatially uniform, consistent with a common climate signal (Figure 33). However, the post-felling spread in transect hydrographs is visibly wider than pre-felling (Figure 33, panel a), with WMC3 consistently above the pack and CEH16 below — suggesting a distance-dependent response that the transect regression lacks the power to resolve with seven wells.

Synthetic FE well extension. FE1 and FE2 — which sit inside or at the immediate edge of the clearfell compartment but lack pre-scraping baselines — were extended backwards to August 2010 using OLS donor regression calibrated on Forest Control wells CEH34, CEH2, and CEH33 (calibration $R^2 > 0.99$, RMSE = 18–25 mm; Section 4.6.1). Against the Forest control, the centroid of WMC3 + FE2 produced a clearfell step of +103 mm [+25, +182] ($p = 0.011$, $R^2 = 0.44$), consistent with the WMC3-only ANCOVA estimate of +93 mm. FE2's post-felling divergence from its donor counterfactual was +28 mm ($p = 0.001$), confirming that the within-compartment response is significant. Including FE1 — which sits in standing forest ~20 m from the boundary — diluted the signal to non-significance (+45 mm, $p = 0.16$), consistent with its position outside the clearfell footprint.

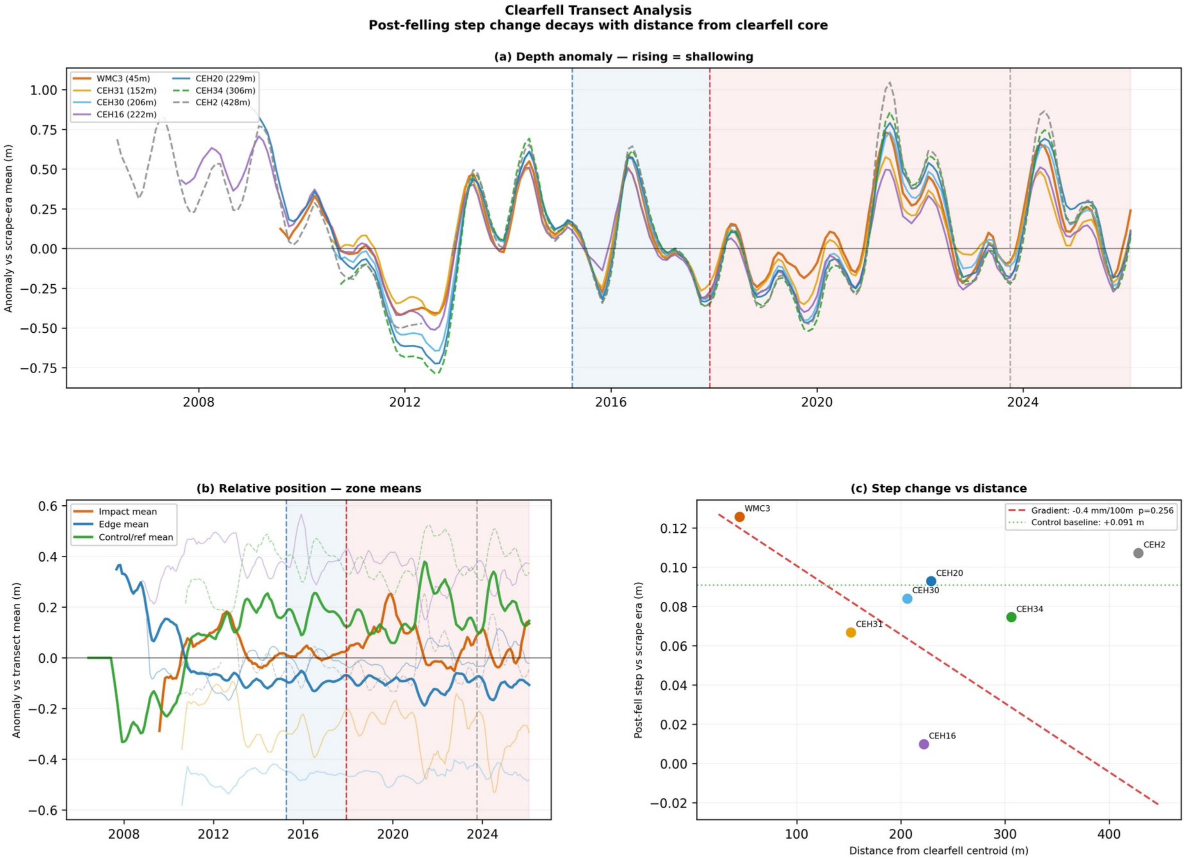

Figure 33. *Clearfell transect: post-felling step change vs distance from felling compartment centroid. Gradient = −0.4 mm/100 m (p = 0.26, ns). No spatial decay consistent with a localized clearfell response; the uniform pattern indicates a common climate signal. Source: 10g_02_clearfell_transect.png.*

4.7 Critical Rainfall Threshold Forecasting - Predicting Winter flooding and Summer droughts

4.7.1 Mechanistic State-Space Equations

The cluster-level state-space equations (Section 3.6.1) fitted without intercept produced statistically significant coefficients at all five clusters (all $p < 0.001$ except C4 β_3 where $p = 0.002$), confirming that the physical mass-balance formulation is well-supported by the data. Model fit ranged from $R^2 = 0.810$ (C3) to 0.682 (C4), reflecting the progressively greater storage buffering captured by the increasing LCSC values across clusters (Table 3). The full coefficient values are reported in Section 4.2.2; note that β coefficients used in the forecasting tool are expressed in m/mm rainfall, one thousand times smaller than the m/m values in Table 3.

4.7.2 Seasonal Prediction Equations

The prediction equations were fitted at hydrological-year resolution, one equation per cluster under the $k = 5$ partition (Table 13). The two predictor terms — winter rainfall total (P_winter) and antecedent summer minimum (h_min) — show a clear spatial partitioning of dominance across the site.

In the Eastern Block, winter rainfall is the dominant predictor. P_winter is significant at Lake Edge ($p = 0.002$) and Dune ($p = 0.005$), while h_min is non-significant in both ($p = 0.548$ and 0.538 respectively). This is consistent with the faster-draining, lower-memory character of the eastern aquifer, where same-season rainfall dominates the winter flooding response.

In the Western Residual and forest clusters, the pattern reverses. The antecedent summer minimum is the dominant predictor: Western Residual ($p = 0.010$, $R^2 = 0.661$), Forest ($p < 0.001$, $R^2 = 0.871$) and Coastal Forest ($p < 0.001$, $R^2 = 0.881$). Winter rainfall is non-significant at all three ($p = 0.057$, 0.710 and 0.361 respectively). The depth of the preceding summer drought is a stronger predictor of winter flooding potential than the actual winter rainfall amount in these more buffered domains. This reflects the aquifer's multi-month memory — a deep summer minimum requires progressively more winter rainfall simply to restore the water table to its pre-drought position before any flooding can occur. The significant intercepts at Forest ($p = 0.016$) and Coastal Forest ($p = 0.003$) are consistent with the strong structural offset imposed by the plantation canopy on water table depth.

Two individual coefficients require comment. The negative h_min coefficient at Lake Edge (−0.134) is not mechanistically meaningful given its non-significance; it is an artefact of the small sample ($n = 19$) and the dominance of P_winter in that cluster. The negative P_winter coefficient in the Forest cluster (−0.00022) reflects multicollinearity with h_min and should not be interpreted mechanistically.

Block	Equation	R^2	n	P (P_winter)	P (h_min)
Lake Edge (C1)	h_peak = 0.00109·P_winter − 0.134·h_min − 0.689	0.47	19	0.002	0.548
Dune	h_peak = 0.00159·P_winter + 0.199·h_min − 0.798	0.52	20	0.005	0.538
Western Residual (C3)	h_peak = 0.00124·P_winter + 0.867·h_min − 0.033	0.66	20	0.057	0.01
Main Forest (C4)	h_peak = −0.00022·P_winter + 1.362·h_min + 1.490	0.87	19	0.71	<0.001
Coastal Forest (C5)	h_peak = 0.00042·P_winter + 1.616·h_min + 1.386	0.88	19	0.361	<0.001

Table 13. *Winter peak prediction equations fitted at hydrological-year resolution (Section 3.6.2), one equation per cluster under the k = 5 partition. P_winter = total winter rainfall (mm) accumulated over the cluster-specific recharge horizon (October–January for C1 and C2; October–February for C3, C4 and C5). h_min = preceding summer water table minimum (m below surface). (11_forecast_winter_transfer_functions.csv)*

4.7.3 Summer Drought Prediction Equations

The summer drought prediction equations performed substantially better than the winter peak functions (Section 3.6.2). Both predictors — antecedent winter peak (h_max_winter) and cumulative summer precipitation (P_summer) — were significant in all five blocks (all $p < 0.05$). Model fit was highest for the Forest ($R^2 = 0.912$) and Coastal Forest ($R^2 = 0.907$) blocks, high for the Western Residual ($R^2 = 0.711$), moderate-to-high for the Dune ($R^2 = 0.626$), and moderate for the Lake Edge Cluster($R^2 = 0.404$). These functions provide a reliable basis for forecasting summer drought severity from the preceding winter state and are the more operationally useful of the two prediction equation pairs (Table 14).

Block	Equation	R^2	P (P_summer)	P (h_max)
Lake Edge (C1)	h_min = 0.00139·P_summer + 0.415·h_max_winter − 1.430	0.404	0.005	0.049
Dune (C2)	h_min = 0.00164·P_summer + 0.508·h_max_winter − 1.573	0.626	<0.001	<0.001
Western Residual (C3)	h_min = 0.00116·P_summer + 0.537·h_max_winter − 1.437	0.711	0.012	<0.001
Forest (C4)	h_min = 0.00106·P_summer + 0.712·h_max_winter − 1.436	0.912	0.013	<0.001
Coastal Forest (C5)	h_min = 0.00072·P_summer + 0.526·h_max_winter − 1.346	0.907	0.032	<0.001

Table 14. *Summer drought prediction equations (Section 3.6.2). P_summer = total April–September rainfall (mm); h_max_winter = preceding winter water table peak (m below surface). All coefficients significant at $p < 0.05$. (11_forecast_summer_transfer_functions.csv).*

4.7.4 Spatial Distribution of P_flood and Flooding Status Across the Well Network

Application of the cluster-level P_flood equations derived in Section 3.6.3 to each well in the 87-well classified network (66 reference wells with per-well SSM coefficients; 21 extended wells assigned cluster-level coefficients via the Pearson affinity audit of Section 3.3), using that well's observed mean summer minimum as h_0, produces a per-well critical rainfall threshold expressed as either millimetres of cumulative winter rainfall or as a dimensionless multiplier $\lambda = P\_flood / \Sigma\ P\_clim$, where $\Sigma\ P\_clim$ is the long-term mean rainfall over the cluster-specific recharge horizon (464 mm for C1/C2 over Oct–Feb; 524 mm for C3/C4/C5 over Oct–Mar). A well with $\lambda = 1.0$ would flood under an average winter; $\lambda > 1.0$ indicates that above-average rainfall is required.. The network-wide picture is summarised in Figures 35-39.

The summer minimum depth map (Figure 34) classifies each well against the five eco-hydrological bands derived from Curreli et al. (2013) and introduced in Section 2: SD15b wet slack viability (<0.61 m), SD15b-recoverable (0.61–0.75 m), SD16 dry slack (0.75–0.98 m), SD16-recoverable (0.98–1.20 m), and beyond-recoverable (>1.20 m). The two forest clusters (C4 and C5) show the deepest mean summer minima across the network, consistently below 0.98 m. All C4 and most C5 wells fall within the SD16 or beyond-recoverable zone, reflecting the suppression of net recharge by canopy interception. By contrast, the majority of C1 Lake Edge wells maintain mean summer minima within or close to the SD15b wet slack zone, reflecting the high-frequency recharge dynamics of the shallow till substrate. Of the 87 wells in the classified network, 41 have mean summer minima deeper than 0.98 m — the SD16 dry slack threshold — of which 25 are below the SD16-recovery threshold of 1.20 m.

The P_flood map (Figure 35) quantifies the rainfall demand implied by these summer minima and reveals a progressive, cluster-ordered closure of the operational window (Table 15). All wells in the classified network require rainfall above their respective climatological mean to reach the slack floor (all $\lambda > 1.0$). Under mildly wet winters ($1.0 < \lambda < 1.5$), 62 wells (71%) become reachable; a further 16 (18%) require the wettest winters in the instrumental record ($1.5 \leq \lambda < 2.0$); the remaining 9 (10%) are structurally unreachable under rainfall totals unprecedented in the RAF Valley record.

Cluster	n	P_flood range (mm)	Median P_flood (mm)	Median λ	Character
C1 Lake Edge	8	499–561	513	1.11	Uniformly recoverable; median h_0 = 0.76 m
C2 Dune	35	469–806	539	1.16	Engineered-recoverability zone; BACI scraping benefit (+0.131 m) can shift marginal wells into compliance
C3 Western Residual	20	616–990	670	1.28	Wet-winter reachable, deepest wells (λ ~1.9) approaching unreachable
C4 Main Forest	13	677–1,237	860	1.64	Substantial structural impairment; deepest wells unreachable
C5 Coastal Forest	11	769–1,651	1,038	1.98	Wide range reflecting heterogeneity of coastal margin

Table 15. *Per-cluster P_flood summary across the 87-well classified network. λ = P_flood / Σ P_clim; values above 1.0 indicate rainfall above the climatological mean is required. (11b_03_pflood_per_well.csv)*

The winter maximum depth map (Figure 36) shows that Eastern Block wells achieve the shallowest winter maxima, consistent with the flashy, fill-and-spill regime of the shallow till aquifer. Western Residual and Forest wells rarely approach the 0 m threshold required by SD15b communities. The flooding frequency map (Figure 37) gives the most ecologically direct summary: C1 wells show the highest flooding frequency across the monitoring period; C3 wells show very low frequency; C4 and C5 Forest wells show essentially zero, with the water table failing to reach the ground surface in any monitored year.

The Eastern Block (C1 and C2) therefore remains within the operational domain of current scraping practice. C3 Western Residual is approaching the margin of operational recoverability. The two forest clusters (C4 and C5), as a whole, cannot be returned to ecological compliance by the existing excavation strategy alone under the present hydrological regime.

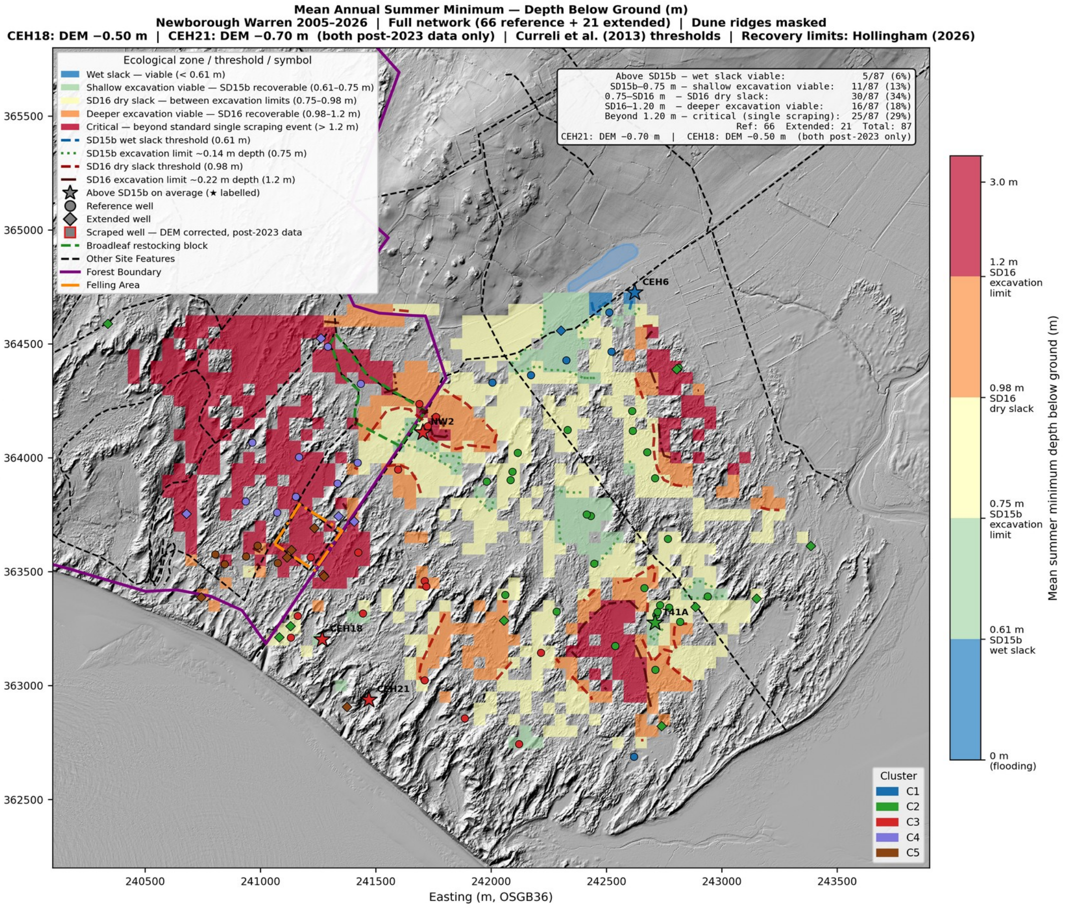

Figure 34. *Spatial distribution of mean annual summer minimum water table depth below ground across the full dipwell network (66 reference wells and 21 extended wells), Newborough Warren 2005–2026. Depth is expressed as metres below the LiDAR ground surface (DEM_Ground_Elev − mean maOD), with positive values indicating the water table lies below the surface. The continuous IDW-interpolated surface is masked across dune ridges (cells where the DEM raster exceeds the interpolated well surface by more than 1 m). Ecological zone boundaries follow the eco-hydrological thresholds of Curreli et al. (2013): SD15b wet slack viability (0.61 m); SD15b excavation limit following one scraping event (0.75 m); SD16 dry slack threshold (0.98 m); SD16 excavation limit (1.20 m). Recovery limits are indicative, derived from the observed BACI scraping benefit of +0.131 m at CEH36 (Section 4.5); deeper scraping or repeat interventions may achieve greater uplift. Wells CEH18 and CEH21 use post-October 2023 data only, with DEM corrections of −0.50 m and −0.70 m respectively to account for material removed during scraping. Well symbols are coloured by cluster; circles indicate reference wells, diamonds indicate extended network wells. (LiDAR data: © NRW & OS; see Section 10 for full attribution.) A ridge-masked variant of this surface is available in the Newborough Warren Hydrogeological Modelling Suite scenario viewer (Hollingham, 2026b). (11b_01_summer_minima_depth.png).*

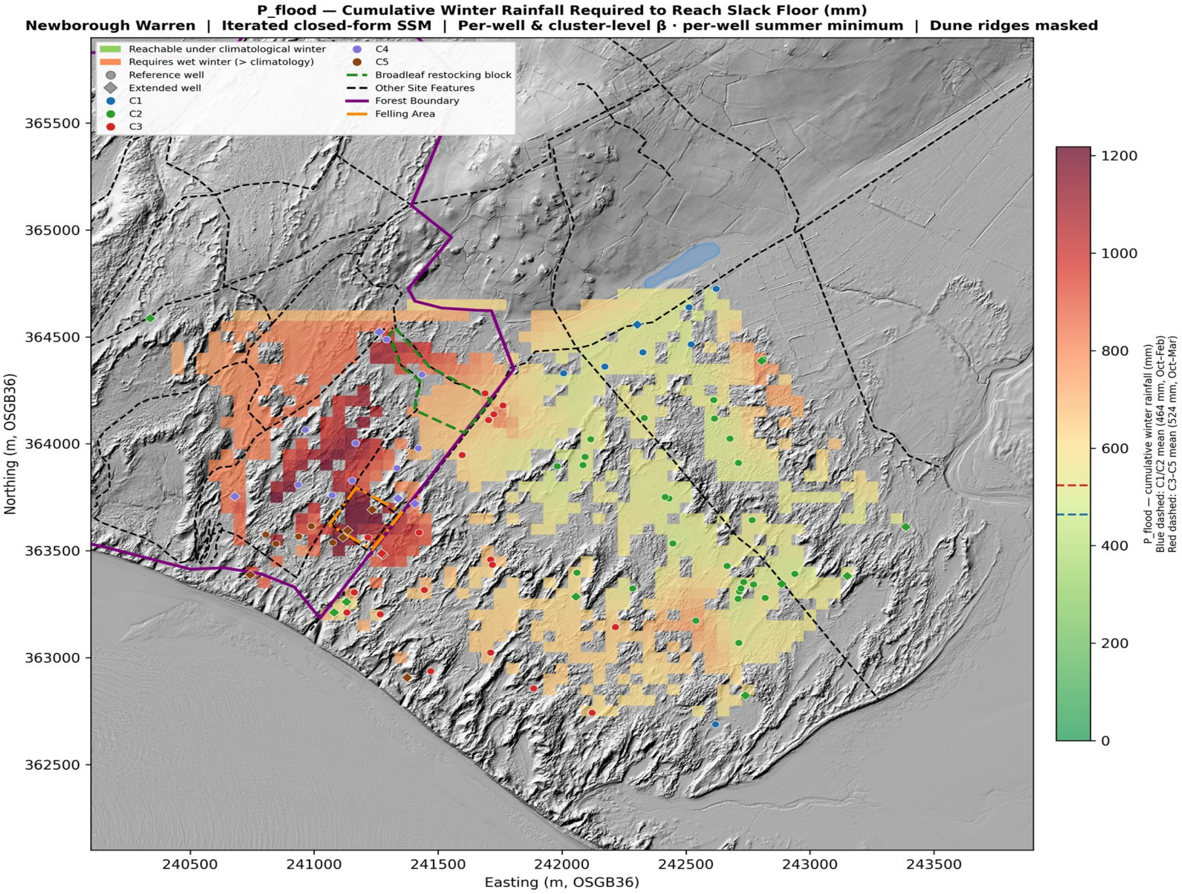

Figure 35. *Spatial distribution of P_flood — the minimum cumulative winter rainfall (mm) required to raise the water table from its observed mean summer minimum to the local slack floor (h_target = 0 m) — across the 87-well classified network, Newborough Warren. P_flood is computed at each well using per-well SSM coefficients where available (66 reference wells) or cluster-level centroid coefficients assigned via Pearson affinity (21 extended wells), applying the iterated closed-form threshold equation of Section 3.6.3 to that well's observed mean summer minimum. The IDW-interpolated surface is masked across dune ridges. Dashed reference lines on the colourbar indicate the cluster-specific climatological rainfall totals: 464 mm for C1/C2 (Oct–Feb, blue) and 524 mm for C3–C5 (Oct–Mar, red). Wells whose P_flood falls below their cluster's reference line are reachable under average winter rainfall. (LiDAR data: © NRW & OS; see Section 10 for full attribution.) (11b_03_pflood.png).*

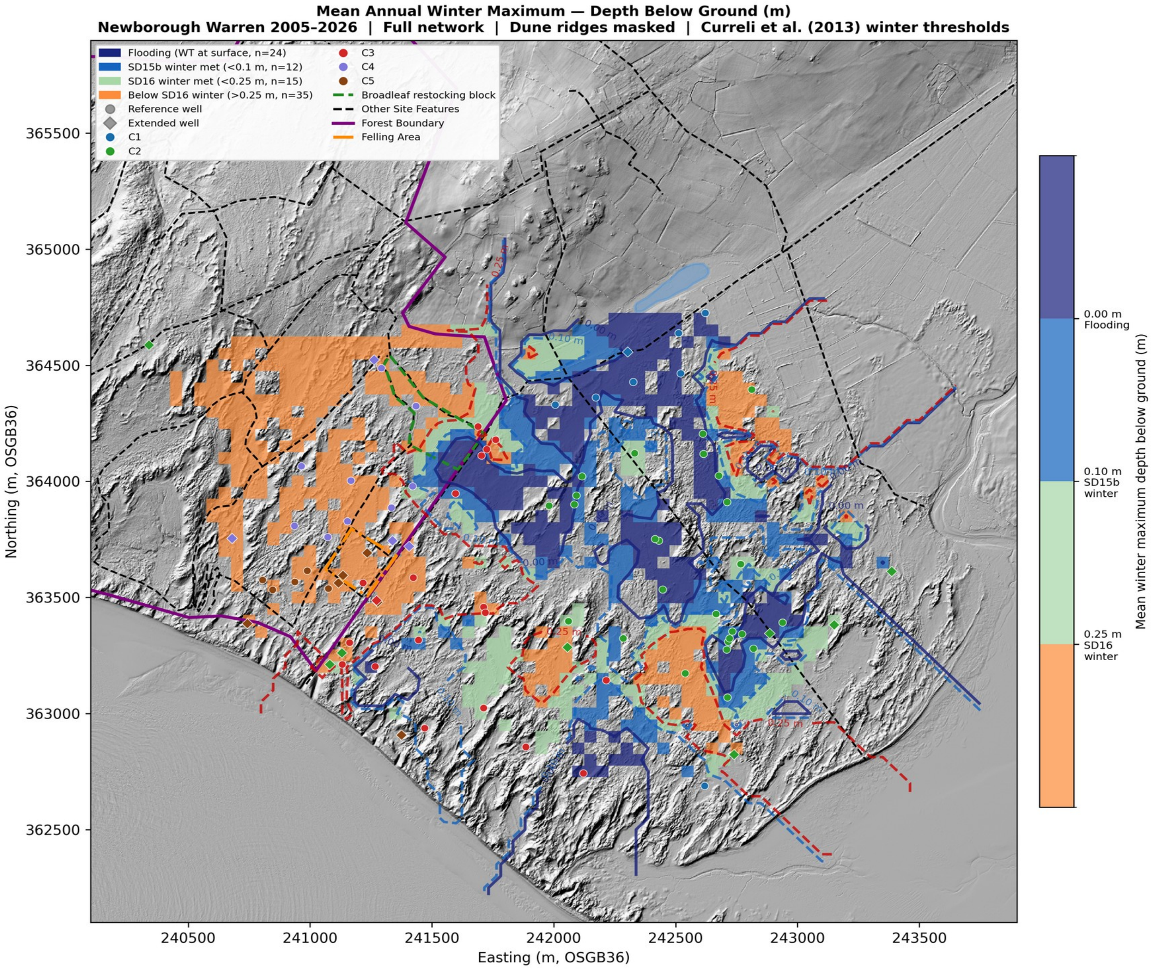

Figure 36. *Spatial distribution of mean annual winter maximum water table depth below ground across the full dipwell network (66 reference wells and 21 extended wells), Newborough Warren 2005–2026. Winter maximum is defined as the highest monthly maOD value in October–March of each hydrological year (October start); the mean is taken across all years with at least three October–March readings. Depth is expressed using the same DEM-referenced convention as Figure 34. The IDW-interpolated surface is masked across dune ridges. Contours mark the eco-hydrological thresholds of Curreli et al. (2013) for winter water table requirements: 0 m (ground surface — required for SD15b wet slack communities); 0.10 m below ground (minimum winter maximum for SD15b viability); 0.25 m below ground (SD16 dry slack winter threshold). The spatial contrast between the C1 Lake Edge — where winter maxima frequently reach or approach the surface — and the C3 Western Residual, C4 Main Forest and C5 Coastal Forest, where winter maxima rarely approach 0 m, reflects the contrasting aquifer storage architectures identified in Section 4.2. (LiDAR data: © NRW & OS; see Section 10 for full attribution.) (11b_02_winter_maxima_depth.png).*

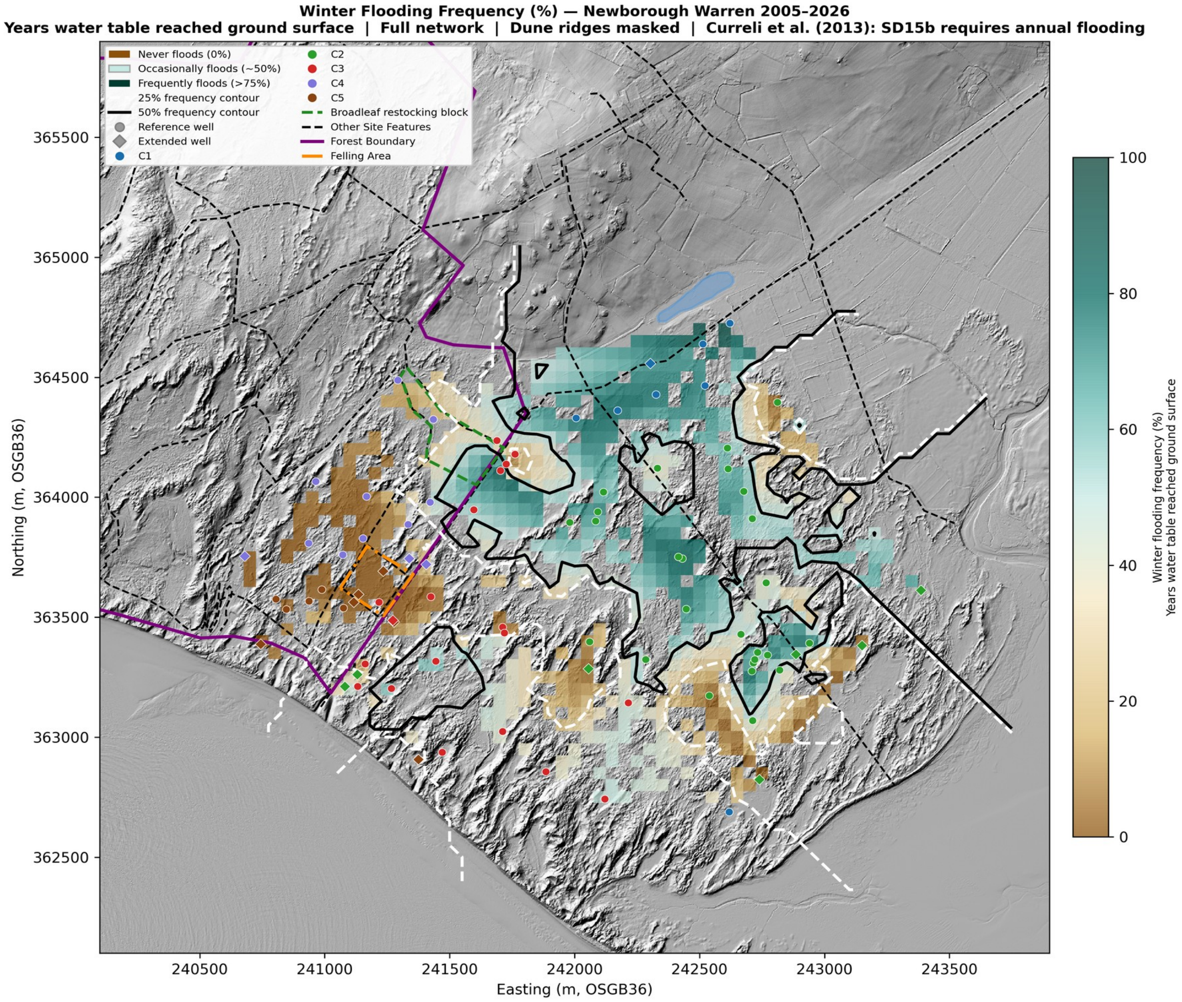

Figure 37. *Winter flooding frequency across the full dipwell network (66 reference wells and 21 extended wells), Newborough Warren 2005–2026, expressed as the percentage of monitored hydrological years in which the winter maximum water table reached or exceeded the local ground surface (DEM ground elevation used as ground surface proxy at each dipwell location). Only wells with at least five years of winter data are included. The IDW-interpolated surface is masked across dune ridges. Contours mark the 25% and 50% frequency thresholds. The C1 Lake Edge shows the highest flooding frequency, consistent with the flashy, fill-and-spill dynamics of the shallow till substrate. C3 Western Residual wells show very low flooding frequency, and C4 and C5 Forest wells record essentially zero — the water table fails to reach the ground surface in any monitored year. This map provides the most ecologically direct assessment of habitat viability: the SD15b wet slack community identified by Curreli et al. (2013) requires the winter water table to rise within 0.10 m of the ground surface, a condition approached regularly only by the C1 Lake Edge under current conditions. (LiDAR data: © NRW & OS; see Section 10 for full attribution.) (11b_04_flood_frequency.png).*

4.7.5 Critical Rainfall Threshold Equations (P_flood)

The structural difference between clusters is immediately apparent from the threshold equations (Table 16; Section 3.6.3). The coefficients from Table 3, together with the cluster-specific recharge horizons from October to each cluster's historical peak month (five months for C1 and C2, six for C3, C4 and C5) and the long-term monthly RAF Valley rainfall and PET climatology, propagate through the iterated closed form of Section 3.6.3 to produce a linear relationship between observed summer-minimum depth and required winter rainfall.

The iterated derivation reduces to the form P_flood = A · d + B for each cluster, where d is the observed summer-minimum depth below ground (m, positive) and P_flood is the cumulative winter rainfall required over the cluster's recharge horizon (mm). The slope A represents the marginal rainfall demand per additional metre of summer drawdown, and the intercept B represents the rainfall demand required even for a well already at the surface — non-zero because accumulated winter PET continues to draw down the water table through the recharge season. Both terms express the cluster gradient directly: A rises from 167 mm per metre at C1 through 215 (C2) and 230 (C3) to 355 (C5) and 371 (C4), and B ranges from 357 mm (C4) to 510 mm (C5).

As a worked example, at a typical C3 well with a mean summer minimum of 1.0 m below ground, P_flood = 229.67 × 1.0 + 463.00 = 692.7 mm of cumulative winter rainfall over October–March, corresponding to a rainfall multiplier λ = 693/524 = 1.32 times the long-term climatological total. The same depth at a C2 well produces P_flood = 215.00 × 1.0 + 367.24 = 582.2 mm (λ = 582/464 = 1.25), a 110 mm difference reflecting the stronger antecedent-depth dependence and longer recharge horizon of C3. These differences translate directly into operational recoverability status: the C2 well lies within a marginally-wet-winter-reachable range, while the C3 well requires slightly wetter conditions over a longer accumulation months endeason.

Cluster	**Label**	**Horizon**	**P_flood equation (P_Flood**(mm)**=)**	**Σ P_clim (mm)**
C1	Lake Edge	Oct–Feb (5 mo)	167.18·d + 398.64	464
C2	Dune	Oct–Feb (5 mo)	215.00·d + 367.24	464
C3	Western Residual	Oct–Mar (6 mo)	229.67·d + 463.00	524
C4	Main Forest	Oct–Mar (6 mo)	370.91·d + 356.35	524
C5	Coastal Forest	Oct–Mar (6 mo)	354.65·d + 509.90	524

Table 16. *Cluster-specific linear forms derived from the iterated closed-form P_flood solution of Section 3.6.3, reduced by substituting fitted SSM coefficients (β_1, β_2, β_3), cluster recharge horizons (n = 5 for C1/C2, 6 for C3/C4/C5), and RAF Valley long-term monthly rainfall and PET climatology (2005–2026). Input variable d is the observed mean summer minimum water table depth below ground in metres, entered as a positive number (e.g. d = 1.26 for a water table 1.26 m below the surface). Output P_flood is the cumulative rainfall (mm) required over the cluster's recharge horizon (October to the historical mean peak month) to raise the water table to the slack floor (h_target = 0 m), assuming monthly rainfall scales uniformly by a multiplier λ and monthly PET follows climatological means. The rainfall multiplier is λ = P_flood / $\Sigma \bar{P}_i$; values above 1.0 indicate rainfall above the climatological mean is required. (11_forecast_pflood_summary.csv)*

4.8 Climate Trajectory and Threshold Exceedance

The spatial distribution of mean annual summer minima and winter maxima across the full dipwell network is presented in Figures 35 and 37 respectively, classified against the Curreli et al. (2013) eco-hydrological thresholds introduced in Section 2. The cluster-level trajectories and threshold exceedance analysis presented in this section quantify the rate and direction of change over the 21-year monitoring period.

4.8.1 Summer Minimum Trajectory

Summer minimum trends were fitted for all five clusters over the monitoring period (Section 3.6.4; Figure 38, panel a). All clusters show declining summer minima, although the rates and significance levels vary substantially. C1 Lake Edge declined at −0.0097 m yr^{-1} ($R^2 = 0.200$, $p = 0.048$, $n = 20$) and C5 Coastal Forest at −0.0381 m yr^{-1} ($R^2 = 0.400$, $p = 0.003$, $n = 20$) — the only two clusters reaching statistical significance at the 5% level. C2 Dune was marginal (−0.0105 m yr^{-1}, $p = 0.084$, $n = 21$). C3 Western Residual (−0.0077 m yr^{-1}, $p = 0.273$) and C4 Main Forest (−0.0115 m yr^{-1}, $p = 0.371$) were non-significant. The moderate R^2 values across all clusters reflect substantial interannual variability rather than an absence of signal.

These trends translate into progressively worsening threshold compliance. C1 has been below the wet slack summer viability threshold (−0.61 m) in all 20 monitored years, with the shallowest summer minimum recorded in 2016 (−0.705 m). C2 has similarly remained below the threshold throughout the record, with its shallowest minimum in 2005 (−0.682 m). C3 has never approached the wet slack threshold, ranging between −0.955 m and −1.586 m, and falling below the dry slack threshold (−0.98 m) in 20 of 21 years.

Winter Maximum Trends were fitted for all five clusters. C1 showed a positive but non-significant trend (+0.0088 m yr^{-1}, $p = 0.119$); C2 (+0.0056, $p = 0.516$), C3 (+0.0014, $p = 0.904$) and C4 (−0.0036, $p = 0.842$) were all non-significant. C5 showed a significant declining winter maximum trend (−0.0368 m yr^{-1}, $p = 0.035$), indicating that this cluster is deteriorating in both seasons.

The C5 Coastal Forest trajectory anomaly warrants specific comment. At −0.038 m yr^{-1} in summer and −0.037 m yr^{-1} in winter, C5's decline rate is 3.9 times the summer rate of C1 — the only other cluster with a statistically significant summer trend — and C5 is the only cluster with significant trends in both seasons. The magnitude of this divergence argues against a purely climate-driven explanation, since all clusters share the same climate forcing. Three candidate explanations, which are not mutually exclusive, are identified: amplified head response per unit climate forcing due to thin saturated thickness at the coastal margin, where a given recharge deficit produces a proportionally larger head change at lower Sy; progressive coastal boundary retreat, reducing aquifer volume at the seaward edge and producing an ongoing decline with no equilibrium; and progressive canopy maturation, increasing interception and transpiration demand as the Corsican pine stand ages.

The post-2018 collapse in C5 seasonal amplitude (−18% raw, −21% climate-normalised; Section 4.2.1) — the largest amplitude reduction in the network — reinforces the structural interpretation: recharge peaks are no longer reaching their pre-2018 heights even after correcting for drought summers, suggesting that at least one site-specific process is operating beyond the common climate signal. The non-stabilisation of the C5 CUSUM trajectory through to the end of the monitoring record favours an ongoing process (coastal erosion or continued canopy maturation) over a one-off step change.

Network-scale partition of the summer-minimum decline. A non-linear panel regression was fitted to the full 21-year monthly record (12,647 observations across 72 wells; FE wells, the scraped well CEH36, and the BACI clearfell-zone wells excluded) to test whether the C5 anomaly is quantitatively consistent with coastal boundary retreat operating site-wide. Each well's perpendicular distance to the eroding Caernarfon Bay High Water Mark was computed from the OS Open Map Local TidalBoundary (Menai Strait and Llanddwyn Island excluded as non-eroding boundaries). The fitted Dupuit–Forchheimer steady-state strip-aquifer form is $\delta(d) = \max(-30.6 \cdot (1 - d/869), 0) - 5.9$ mm yr^{-1}, with a coast-edge anomaly δ_0 of −30.6 ± 1.8 mm yr^{-1} above the climate background, an inland reach L of 869 ± 39 m beyond which the gradient effect vanishes, and a climate background c of −5.9 ± 0.6 mm yr^{-1} consistent within one standard error with the network-mean per-well summer-minimum slope. The linear-with-cutoff form is preferred over an exponential decay by ΔAIC = 29, and is robust to forest-cover exclusion: refitting after dropping C4 and C5 entirely shifts no parameter by more than one standard error (Script 25; Figure 38b).

Applied to C5's mean coast-proximity of 419 m, the model predicts a summer-minimum slope of −21.8 mm yr^{-1} (climate −5.9 + gradient −15.9), accounting for approximately 42% of C5's observed −38.1 mm yr^{-1} decline. C1, C2 and C4 all lie beyond the fitted inland reach L; their observed slopes are accounted for by the climate background plus residual interannual noise. Progressive coastal boundary retreat therefore moves from a qualitative candidate to a quantified, substantial, but partial contributor to the C5 anomaly. The remaining −16 mm yr^{-1} residual at C5 is consistent with the other two candidate mechanisms named above — amplified head response under thin saturated thickness, and progressive canopy maturation — operating in parallel. Independent corroboration of the BACI ANCOVA's easting × time covariate against this gradient model is provided in Section 4.6.3 (Table 8).

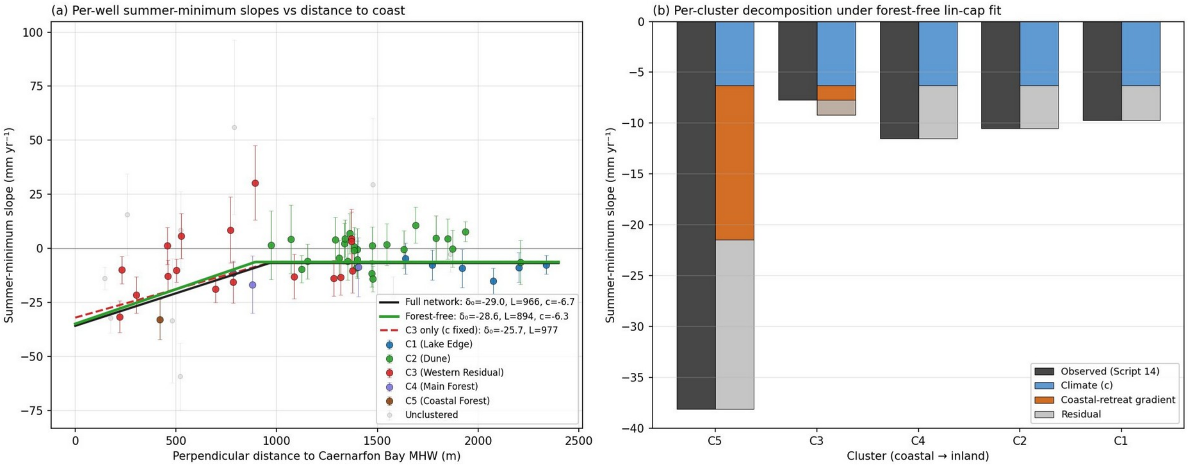

Figure 38b. *Network-scale coastal-retreat gradient. Per-well summer-minimum slope (mm yr*$^{-1}$*, error bars = ±1 SE) plotted against perpendicular distance to the Caernarfon Bay High Water Mark, with three linear-with-cutoff fits overlaid: full-network panel (black), forest-free panel after dropping C4 and C5 (green, headline), and C3-only with climate background held to the forest-free value (red, dashed). The three fits are statistically indistinguishable on δ_0 and L. C5 wells (brown) sit closest to the coast; C1, C2 and C4 lie beyond the fitted reach L = 869 m. (25_05_fit_diagnostic.jpg*

4.8.2 Observed Winter Maximum Flooding Record

The observed winter flooding record reveals a stark spatial contrast between structural domains (Figure 38, panel b). The C1 Lake Edge cluster exceeded the wet slack winter flooding threshold (−0.10 m) in 14 of 20 monitored years (70%), confirming that the eastern lake-influenced slacks retain substantial winter flooding capacity. C2 Dune exceeded the threshold in 11 of 21 years (52%), indicating episodic but reliable flooding in the eastern mature dune zone. C3 Western Residual exceeded the wet slack threshold in only 1 of 21 years (5%) — occurring in 2007 — and the dry slack threshold (−0.25 m) in 6 of 21 years, confirming that the Western Residualfails to achieve the minimum winter flooding conditions required to maintain even dry slack communities in all but exceptional years. C4 Main Forest did not exceed either threshold in any monitored year. C5 Coastal Forest exceeded the wet slack threshold in 1 of 20 years (5%).

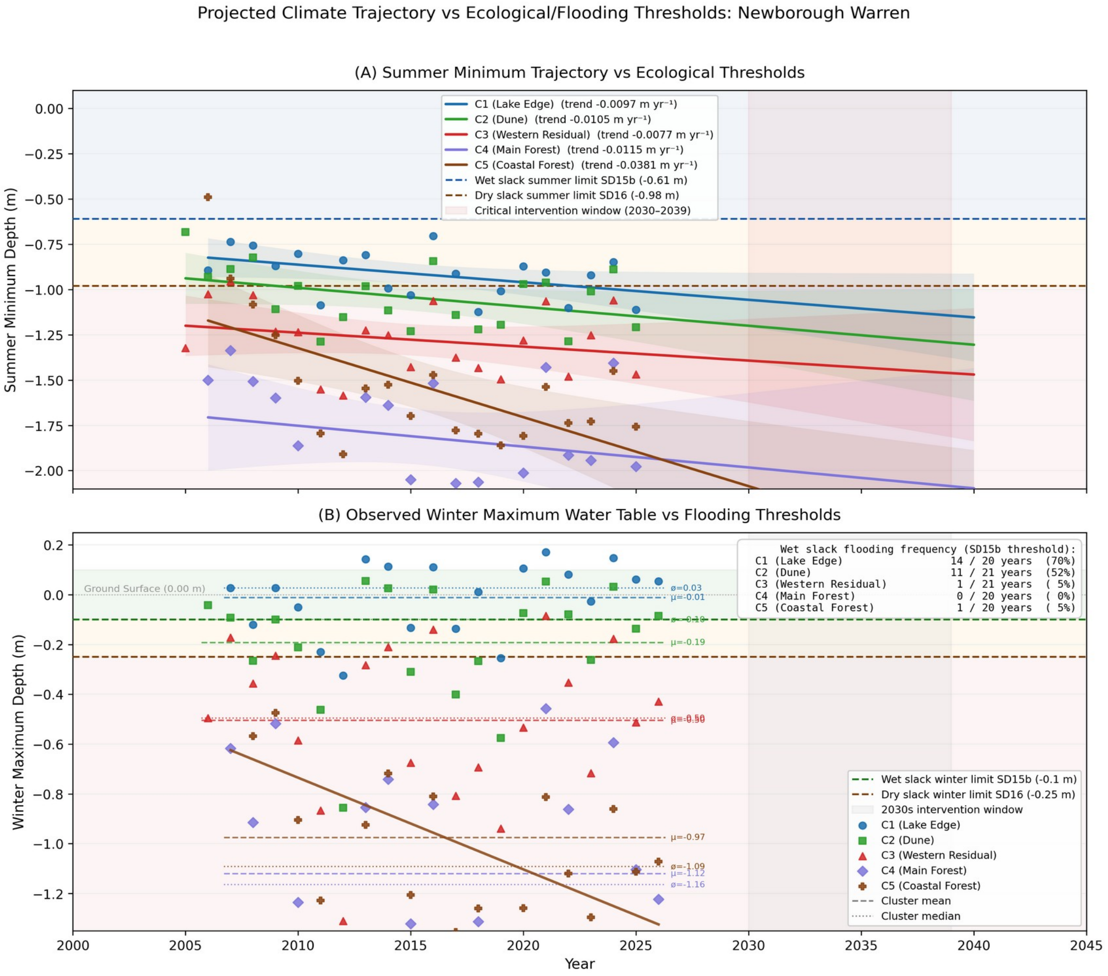

```
    Wet slack flooding frequency (SD15b threshold):
C1 (Lake Edge)                  14 / 20 years  (70%)
C2 (Dune)                       11 / 21 years  (52%)
C3 (Western Residual)            1 / 21 years  ( 5%)
C4 (Main Forest)                 0 / 20 years  ( 0%)
C5 (Coastal Forest)              1 / 20 years  ( 5%)
```

Figure 38. *Climate trajectory and threshold exceedance analysis for Newborough Warren, 2005–2025. Panel (a): observed annual summer minimum water table depth for all five clusters, with OLS linear trend lines and 95% confidence intervals. Dashed horizontal lines indicate the wet slack summer viability limit (SD15b, −0.61 m) and dry slack viability limit (SD16, −0.98 m) from Curreli et al. (2013). Panel (b): observed annual winter maximum water table depth for all five clusters, plotted against the wet slack winter flooding threshold (SD15b, −0.10 m) and dry slack winter flooding threshold (SD16, −0.25 m). The exceedance statistics box reports the proportion of hydrological years in which each cluster exceeded the wet slack flooding threshold over the full monitoring period. (14_climate_trajectory_stacked.png)*

4.9 Spatial Groundwater Analysis

The per-well SSM parameterisation provides a spatially resolved view of aquifer behaviour that complements the cluster-level summaries of Sections 4.2–4.4. This section presents the spatial coefficient structure across the 87-well classified network (66 reference wells with per-well SSM coefficients; 21 extended wells assigned cluster-level coefficients via the Pearson affinity audit of Section 3.3), the baseline head surface and groundwater flow field, the SSM water balance residual, and equilibrium scenario projections under prescribed climate and forest management perturbations. An interactive scenario viewer (Hollingham, 2026b) integrates these outputs with the Curreli et al. (2013) ecological thresholds and is available as a companion tool to this report.

4.9.1 Datum Sensitivity

Per-well drainage datum optimisation (Figure 40a), obtained by independently optimising the displacement datum D at each well to maximise SSM fit (Section 3.4), shows that the R^2-maximising datum varies systematically across the site — from shallow (1.5–2.5 m) at the Lake Edge wells to deep (3.5–5.0 m) in the western and forest interior. However, the R^2 gain from per-well optimisation over the uniform 3.7 m value is negligible across most of the network (Figure 40b), confirming the uniform datum assumption adopted in Section 3.4.

Per-well SSM coefficients interpolated across the site (Section 3.4.3) reveal spatially variable patterns whose broad structure aligns with the cluster boundaries (Figure 40a–d), although within-cluster heterogeneity is substantial. All coefficient surfaces are IDW-interpolated to a 50 m grid with dune ridge cells masked where the DEM raster exceeds the interpolated surface by more than 1 m (see Section 3.8).

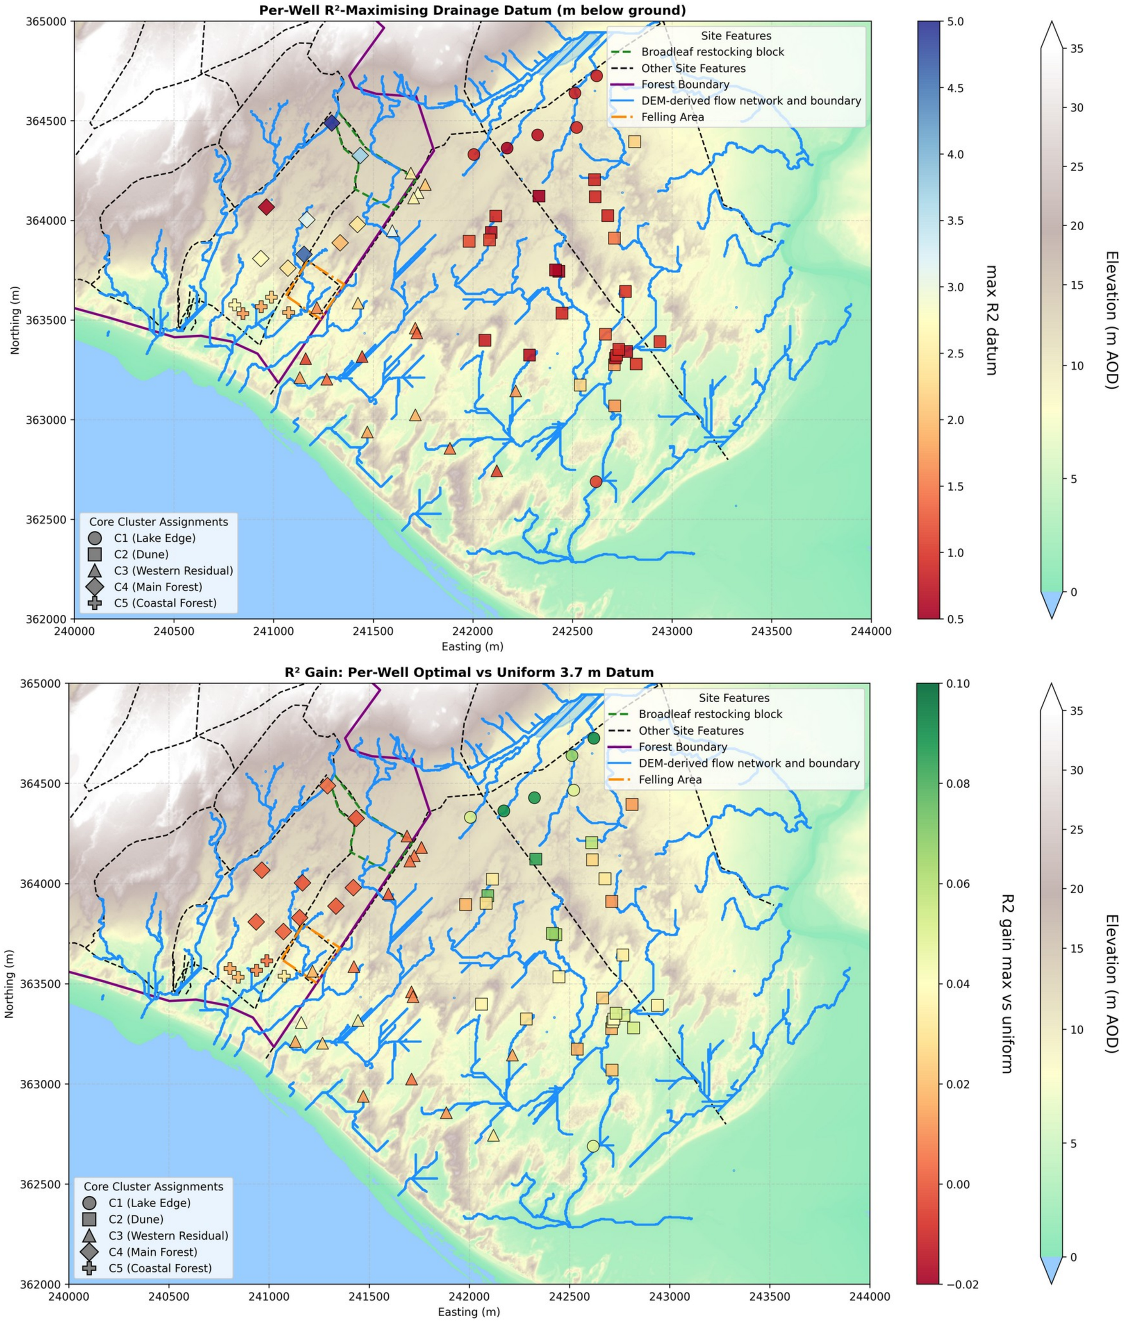

__Figure 39.__ Per-well optimal SSM drainage datum across the 66-well reference network, Newborough Warren. Each well's displacement datum D was optimised independently by sweeping from 0.5 to 5.0 m in 0.1 m steps and selecting the value maximising one-step R² (Section 3.4). (a) R²-maximising datum depth (m below ground surface). The east-to-west gradient from shallow (1.5–2.5 m at C1 Lake Edge) to deep (3.5–5.0 m in the western and forest interior) reflects the depth at which the SSM drainage term best captures lateral discharge at each well location. (b) R² gain from per-well optimal datum vs the uniform 3.7 m network datum. Gains are negligible (<0.01 R²) across most of the network, with meaningful improvement (0.02–0.04 R²) limited to C1 Lake Edge wells where the effective drainage base is shallowest, confirming that the uniform datum assumption does not materially compromise model performance. Both panels: IDW interpolation (power = 2) over 1 m LiDAR DEM hillshade, with ridge and non-aquifer areas masked. Well markers coloured by cluster assignment. LiDAR data: © NRW & OS; see Section 10 for full attribution.

4.9.2 Spatial Coefficient Structure

The R^2 surface (Figure 40a), derived from per-well SSM fits at each well's optimal drainage datum, shows the model explains 70–85% of monthly water table variance across the central and northern open dune, declining toward the forest (C4 and C5 mean 0.68) and some peripheral south-eastern wells. Since the R^2 gain from per-well datum optimisation is negligible (Figure 40b), these values also represent the effective fit quality under the standard pipeline datum. The spatial pattern is consistent with the SSM performing best where the climate signal is strongest and least attenuated by canopy or boundary effects.

The three coefficient surfaces (Figure 40b–d) reveal a coherent spatial story about how rainfall enters, is consumed, and drains from the aquifer. Recharge sensitivity (β_1) is highest in the north-east around the C1 Lake Edge wells closest to Llyn Rhos-ddu (>5 mm/mm), declines across the open dune (3–5 mm/mm), and drops to 2–3 mm/mm west of the forest boundary. Atmospheric draw (β_2) shows a different geography: the highest values occur in two zones — the C4 Main Forest interior, where canopy-driven transpiration amplifies the signal, and a band across the northern open dune adjacent to the ridge — while C1 Lake Edge wells show notably low values (mean 0.61 mm/mm), consistent with a system where head-dependent drainage dominates the loss pathway. C5 Coastal Forest shows moderate β_2 (mean 1.19), lower than C4 despite also carrying pine canopy, possibly reflecting coastal moderation of atmospheric demand. Drainage rate (β_3, logarithmic scale) shows the strongest spatial contrast of any coefficient: C1 and northern C2 wells drain fastest (>8% of displacement per month), the forest interior shows the lowest rates in the network (C4 mean 1.6%, with individual wells below 1%), and C5 falls between the two (mean 4.6%). One C4 well (CEH14) records a negative β_3 (−1.9%), where the displacement formulation does not adequately capture the boundary condition at the ridge flank.

The combined picture is of an aquifer whose eastern margin fills quickly, loses water quickly, and resets each season, while its western and forested interior fills slowly, holds water against evaporative demand, and drains sluggishly — the "shallow pan" versus "deep sponge" contrast identified at cluster level in Section 4.2.

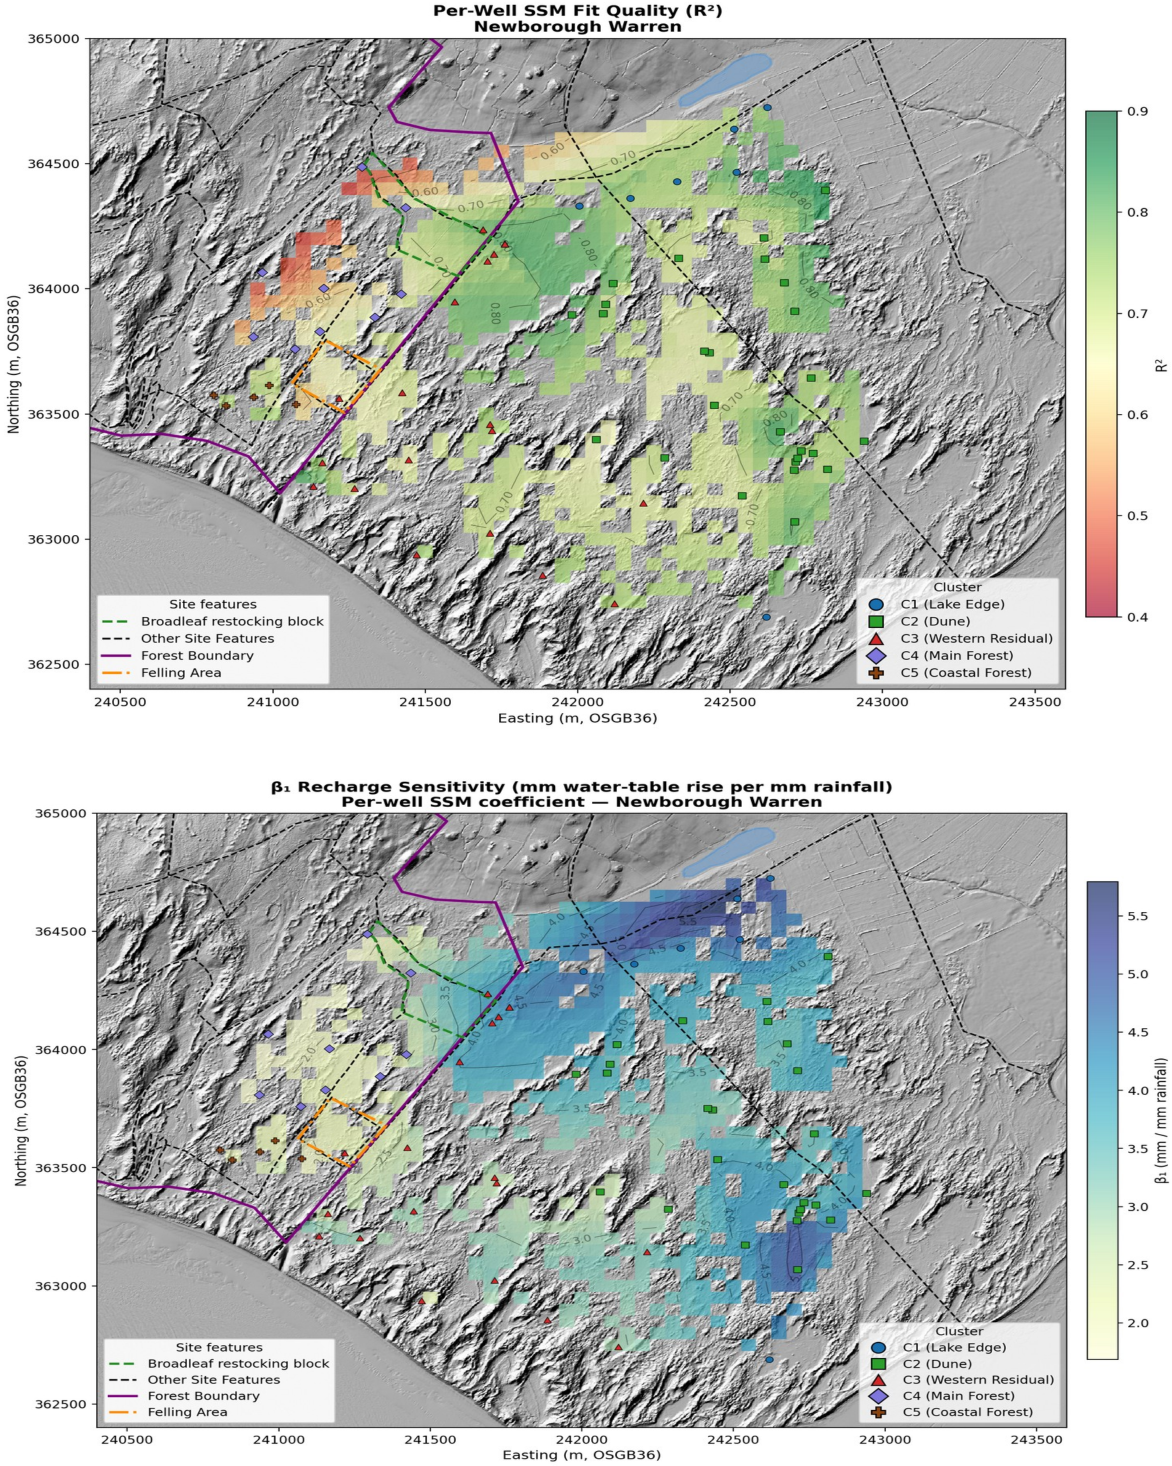

Figure 40. *Spatial SSM coefficient atlas for the 66-well reference network, Newborough Warren. Per-well coefficients interpolated using IDW (power = 2) over 1 m LiDAR DEM hillshade, with ridge and non-aquifer areas masked. Well markers coloured by cluster assignment. (a)* R^2 *one-step model fit quality — the SSM explains 65–85% of monthly water table variance across the open dune, declining to 43–72% in the forest interior (C4). (b)* β_1 *recharge sensitivity (mm water-table rise per mm rainfall) — highest at the Lake Edge (C1, 4–6 mm/mm), with a sharp drop at the forest boundary consistent with 24% canopy interception. Generated by 07_spatial_coefficients.py. (07_coeff_01–04_*.png) LiDAR data: © NRW & OS; see Section 10 for full attribution.*

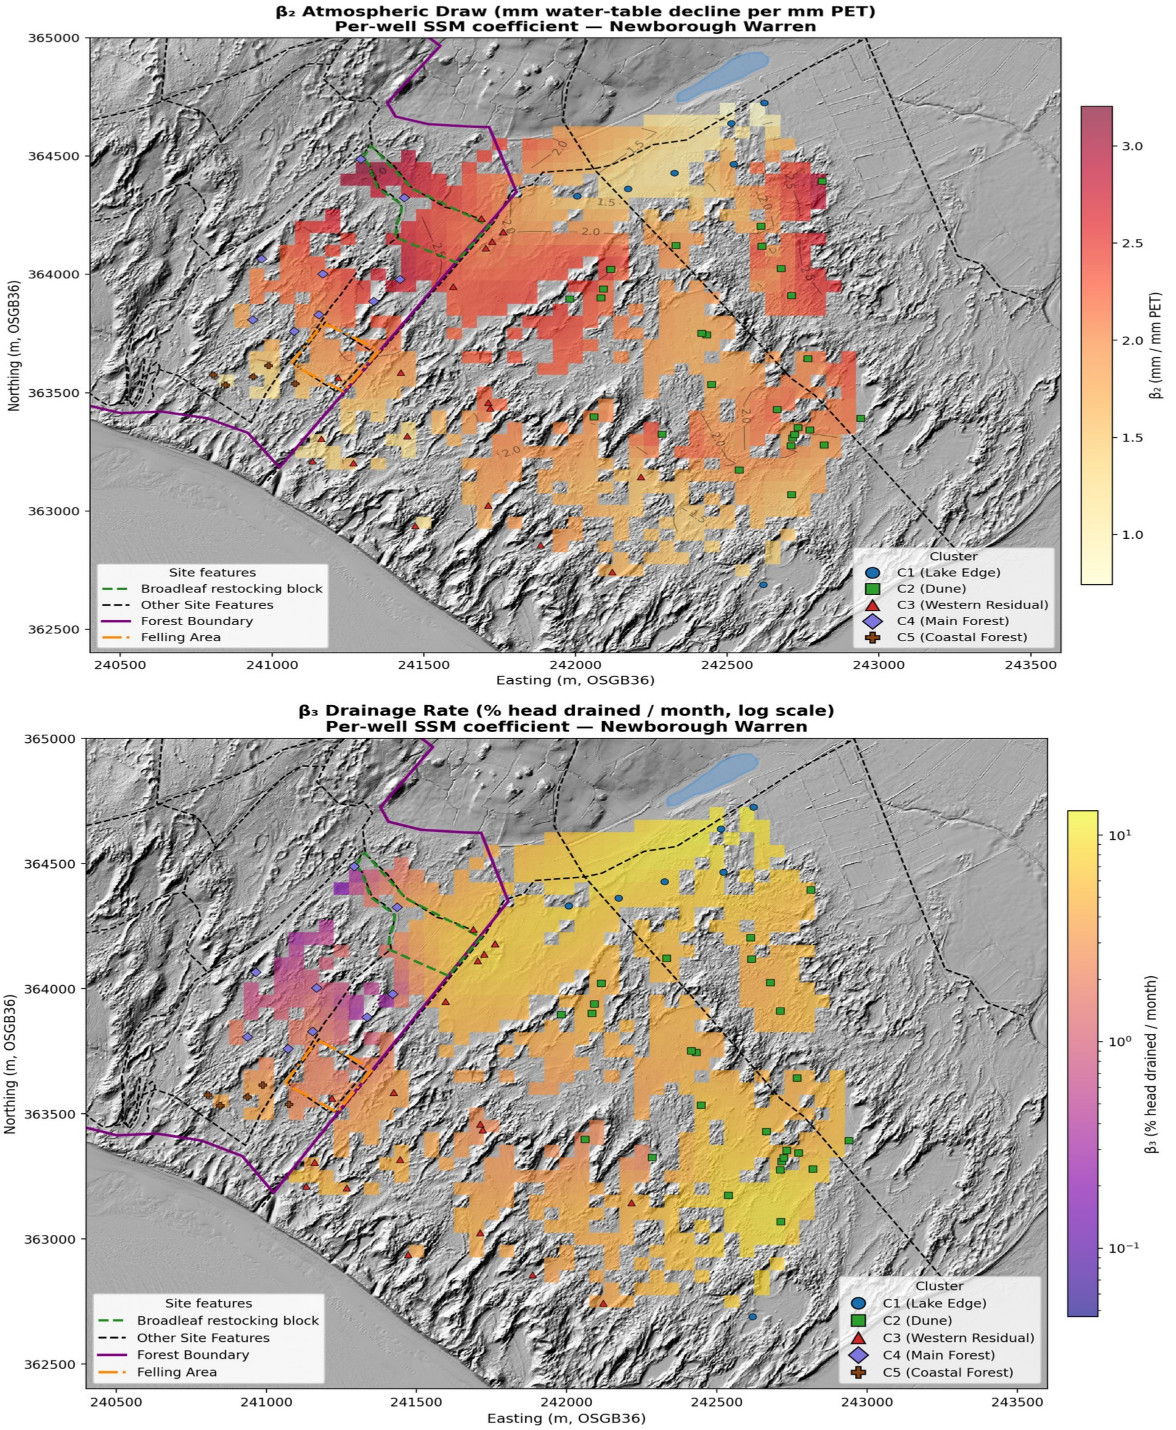

Figure 40. *Spatial SSM coefficient atlas for the 66-well reference network, Newborough Warren. Per-well coefficients interpolated using IDW (power = 2) over 1 m LiDAR DEM hillshade, with ridge and non-aquifer areas masked. Well markers coloured by cluster assignment. (c)* β_2 *atmospheric draw (mm water-table decline per mm PET) — highest in the forest interior (C4, 2.0–3.2 mm/mm), lowest at the Lake Edge (C1) and Coastal Forest (C5). (d)* β_3 *drainage rate (% displacement drained per month, logarithmic scale) — an order-of-magnitude contrast from the fast-draining Lake Edge (7–12% per month) to the near-stagnant forest interior (C4, <3% per month). CEH14, the sole well with negative* β_3*, sits at the ridge flank. Generated by 07_spatial_coefficients.py. (07_coeff_01–04_*.png) LiDAR data: © NRW & OS; see Section 10 for full attribution.*

4.9.3 Drainage Timescale and Forest Drawdown Propagation

The characteristic drainage timescale $\tau = Sy / \beta_3$ (Figure 41) synthesises two independently derived parameters — WTF-derived specific yield (Section 3.7.3) and SSM drainage rate (Section 3.4) — into a single measure of aquifer memory at each well. The eastern open dune (C1, C2) shows rapid turnover (τ = 1.8–5.0 months), the western dune (C3) intermediate values (5–11 months), and the forest interior (C4) the longest timescales in the network (up to 34 months), with a pronounced hotspot on the elevated plateau where β_3 is lowest. C5 Coastal Forest (5–13 months) resembles C3 despite carrying pine canopy, reinforcing the substrate and elevation control identified in Section 4.9.4. CEH14 (negative β_3) and CEH13 (near-zero β_3, $\tau \approx 124$ months) are excluded as ridge-flank wells where the displacement formulation does not apply.

The model benchmarking diagnostic (Section 4.4; Figure 14) complements the timescale map. The iterative NSE improvement from SSM over TLM (ΔNSE) correlates positively with β_3 ($r = 0.63$, $p < 0.001$) and negatively with β_2 ($r = -0.67$, $p < 0.001$). The TLM fails most severely at C1 Lake Edge (median ΔNSE = +1.29, median TLM NSE = −0.57), where head-dependent drainage to the lake is strongest and a fixed drainage rate cannot capture the large seasonal variation in discharge. Conversely, the TLM performs relatively well at C4 Main Forest (median ΔNSE = +0.23, median TLM NSE = +0.42), where β_3 is so weak that misspecifying the drainage term has little impact on overall model performance. ΔNSE therefore identifies where drainage memory is indispensable — the eastern, lake-proximal aquifer — and where atmospheric draw dominates the water balance regardless of drainage formulation — the forest interior.

When τ and ΔNSE are plotted together with point size scaled by Sy (Figure 43), the five clusters occupy distinct regions of the diagnostic space. C1 Lake Edge (low τ, high ΔNSE, low Sy) represents a shallow, fast-cycling aquifer over flat clay/till with a short, highly head-dependent drainage path to the lake. C4 Main Forest (high τ, low ΔNSE, low Sy) represents a sluggish system where drainage is impeded by irregular buried bedrock and β_3 is too weak for its misspecification to matter. C5 Coastal Forest separates cleanly from C4 on both axes despite sharing the same pine canopy — occupying intermediate τ and moderate-high ΔNSE consistent with deeper, more uniform sand at lower elevation. C2 and C3 span the middle ground, with C3's higher τ and Sy reflecting the deeper western sand body. The separation of clusters in this independently derived diagnostic space provides a model-based characterisation of aquifer architecture that is consistent with, but independent of, the Ward's clustering applied to the raw hydrograph data. This is the per-well quantification of the "shallow pan" versus "deep sponge" contrast identified at cluster level in Section 4.2.

The spatial reach of forest management perturbations is shown directly by the drawdown propagation analysis (Figure 42), which models the steady-state head perturbation from forest canopy interception as an exponential decay with flow-weighted cost-distance from the forest boundary (λ = 245 m). Estimated drawdown exceeds 50 mm only within ~100 m of the forest edge, decays to 28 mm at CEH23 and 10 mm at CEH6, and is negligible (<1 mm) at all C1 and eastern C2 wells. Estimated drawdown exceeds 50 mm only within ~100 m of the forest edge, decays to 28 mm at CEH23 and 10 mm at CEH6, and is negligible (<1 mm) at all C1 and eastern C2 wells. These are steady-state values; the timescale to reach equilibrium is governed by the drainage timescales mapped in Figure 41, ranging from months in the eastern dune to years in the forest interior (Section 5.4.3). Forest management perturbations are therefore structurally confined to the forest zone and its immediate western margin; they do not propagate to the eastern slacks where the ecological need is greatest. Forest management perturbations are therefore structurally confined to the forest zone and its immediate western margin; they do not propagate to the eastern slacks where the ecological need is greatest.

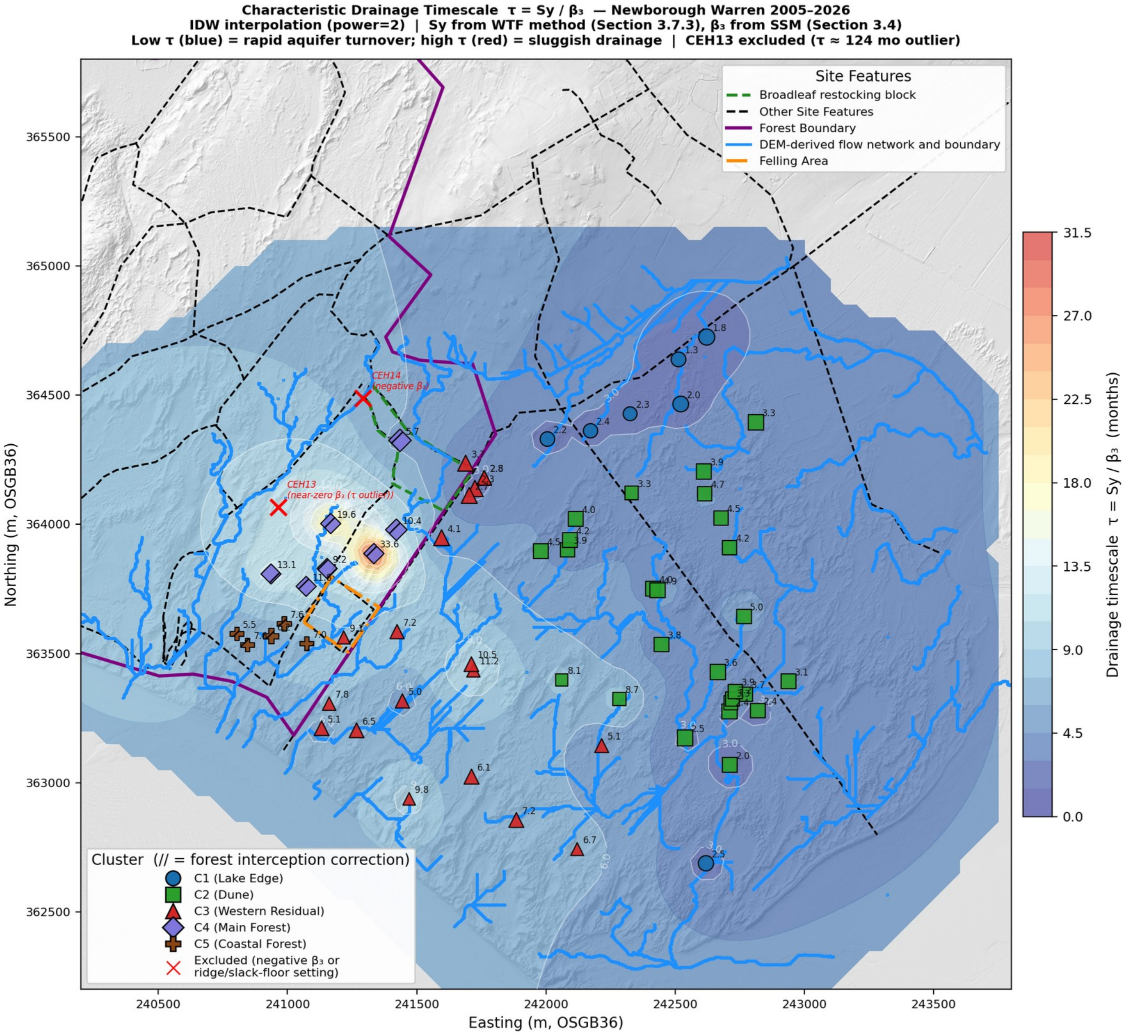

Figure 41. *Characteristic drainage timescale (τ = Sy / β_3, months) across the 66-well reference network, Newborough Warren 2005–2026. τ synthesises independently derived specific yield (WTF method, Section 3.7.3; forest wells interception-corrected per Freeman, 2008) with SSM drainage rate (Section 3.4) to quantify the time required to drain one unit of displacement at each well location. Low τ (blue) indicates rapid aquifer turnover; high τ (red) indicates sluggish drainage with long memory. The eastern open dune (C1, C2: τ = 1.8–5.0 months) contrasts sharply with the forest interior hotspot (C4: up to 34 months), where low β_3 on the elevated plateau produces drainage timescales of two to three years. C5 Coastal Forest (5–13 months) resembles C3 Western Residual despite carrying pine canopy, consistent with the substrate and elevation control identified in Section 4.9.4. Two wells are excluded: CEH14 (negative β_3, τ undefined) and CEH13 (near-zero β_3, τ ≈ 124 months), both ridge-flank settings where the displacement formulation does not apply. IDW interpolation (power = 2) over 1 m LiDAR DEM hillshade, with ridge and non-aquifer areas masked. Well markers coloured by cluster; hatched markers indicate forest interception correction applied to Sy. LiDAR data: © NRW & OS; see Section 10 for full attribution. (18_wtf_05_drainage_timescale_map.png)*

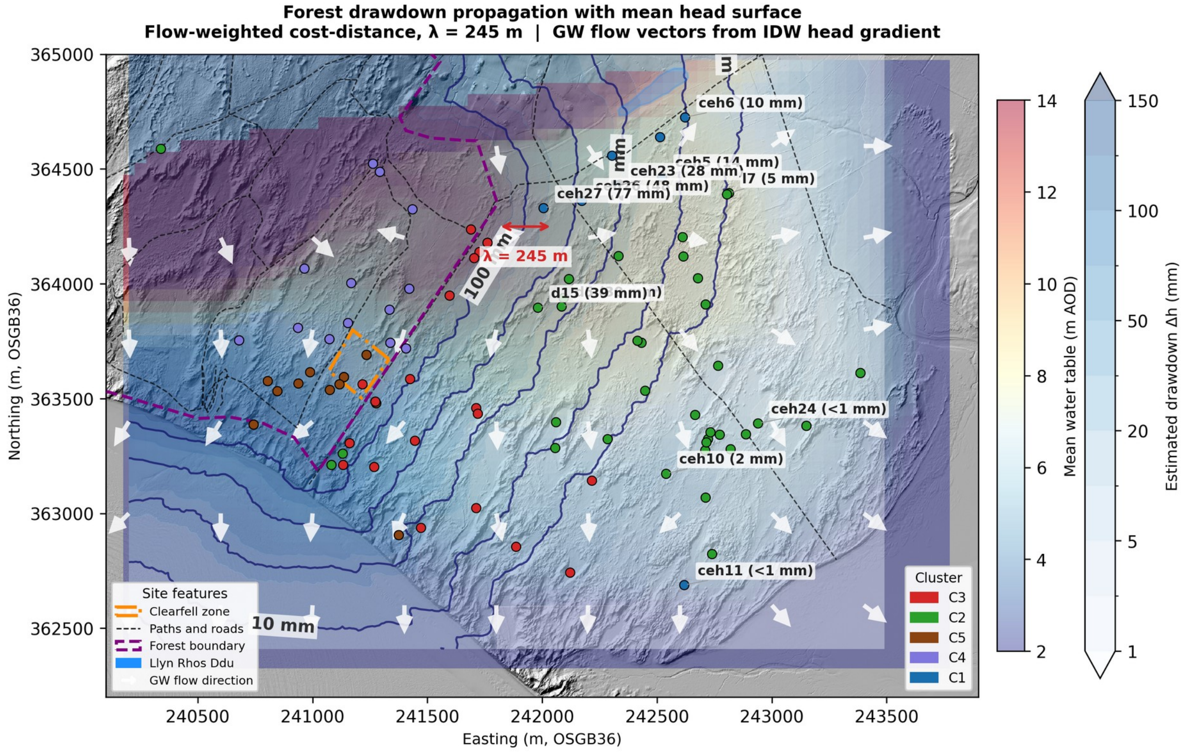

Figure 42. *Forest drawdown propagation across the Newborough Warren dune aquifer, showing the estimated steady-state head perturbation (mm) attributable to Corsican pine canopy interception. Drawdown at each well is modelled as an exponential decay with flow-weighted cost-distance from the forest boundary, using a fitted characteristic length scale λ = 245 m. The background surface shows mean annual water table elevation (m AOD) with normalised Darcy flow direction vectors (white arrows) derived from the IDW head gradient field. Labelled wells show estimated drawdown in parentheses: the signal exceeds 50 mm only within ~100 m of the forest edge (D15: 39 mm), decays to 28 mm at CEH23 and 10 mm at CEH6, and is negligible (<1 mm) at all C1 Lake Edge and eastern C2 Dune wells (CEH24, CEH10, CEH11). The 10 mm contour approximately coincides with the 100 m cost-distance ring from the forest boundary. Forest management perturbations are structurally confined to the forest zone and its immediate western margin. .Well markers coloured by cluster. Background: mean annual water table surface (m AOD; Figure 44, Section 4.9.5) with normalised Darcy flow direction vectors from the IDW head gradient field. LiDAR data: © NRW & OS; see Section 10 for full attribution. (20_drawdown_propagation.png)*

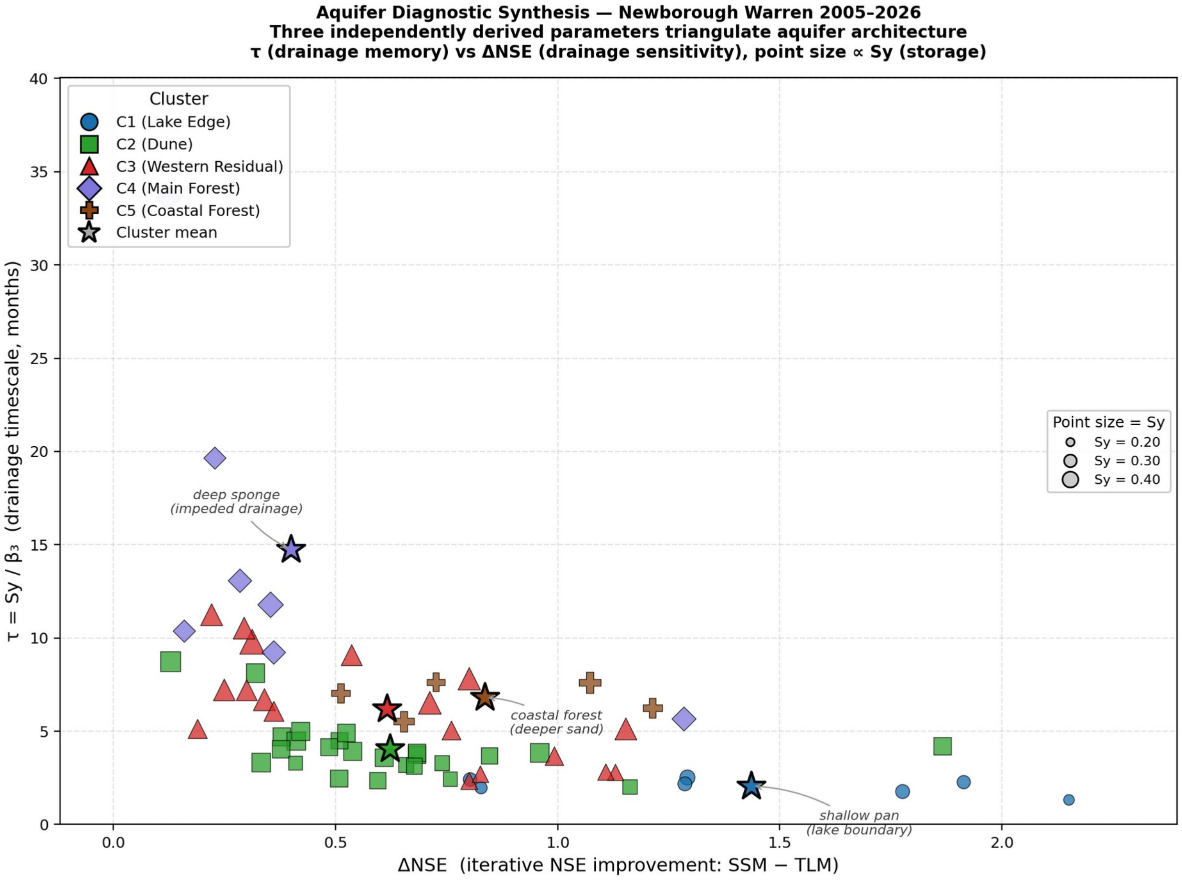

Figure 43. *Aquifer diagnostic synthesis across the reference well network, Newborough Warren 2005–2026. Drainage timescale (τ = Sy / β_3, months) plotted against iterative NSE improvement (ΔNSE = NSE_SSM − NSE_TLM), with points coloured by cluster and sized by WTF-derived specific yield (Sy). Large outlined diamonds show cluster means. The five clusters occupy distinct regions of the diagnostic space: C1 Lake Edge (low τ, high ΔNSE) represents a shallow, fast-cycling aquifer where head-dependent drainage is critical for prediction; C4 Main Forest (high τ, low ΔNSE) represents a sluggish system where drainage is impeded by irregular bedrock and atmospheric draw dominates. C5 Coastal Forest separates from C4 on both axes despite sharing pine canopy, consistent with the substrate and elevation control identified in Section 4.9.4. The three parameters — τ (drainage memory), ΔNSE (drainage sensitivity), and Sy (storage capacity) — are each derived independently, yet together characterize aquifer architecture without requiring borehole data. CEH14 (negative β_3) and CEH13 (near-zero β_3) excluded. (18_wtf_06_aquifer_diagnostic_synthesis.png)*

4.9.4 Per-Well Forest Zone Analysis

The spatial coefficient atlas (Figure 40b–d) shows that the C4 and C5 forest wells span nearly the full site-wide range of β_2 values despite sharing the same Corsican pine canopy. To test whether this within-forest variation is driven by substrate properties or canopy differences, per-well SSM coefficients for the 14 forest zone wells (9 in C4, 5 in C5) were regressed against elevation, distance from the ridge and Easting (Section 3.4.4).

Elevation is overwhelmingly the dominant spatial predictor (Table 17). The atmospheric draw coefficient β_2 shows the strongest relationship ($r = 0.975$, $p < 0.001$; $R^2 = 0.951$), with higher-elevation wells on the ridge returning systematically higher β_2 values. Adding distance from the ridge contributes virtually no additional explanatory power ($R^2 = 0.952$). The drainage coefficient β_3 shows a significant but weaker elevation dependence ($r = -0.771$, $p = 0.001$; $R^2 = 0.594$), with distance from the ridge adding meaningful information (combined $R^2 = 0.721$): wells further from the ridge and at lower elevation drain faster. The recharge coefficient β_1 shows no significant relationship with elevation ($r = 0.223$, $p = 0.444$); Easting is the best single predictor ($r = 0.590$, $p = 0.026$) but explains only 35% of variance. This elevation–drainage relationship is visible directly in the drainage timescale map (Figure 41), where C4 wells show τ values up to 34 months compared with 5–13 months at C5.

C4 and C5 form two clearly distinct groups rather than a continuum. They separate almost entirely on the β_2 axis: C4 wells cluster at $\beta_2 = 2.2$–3.8 while C5 wells sit at $\beta_2 = 0.9$–1.5, with zero overlap. Mean elevations differ significantly (C4: 10.7 m, C5: 5.7 m; $p < 0.001$), with a clean gap between 6.5 m and 8.7 m containing no wells. This corresponds to the transition from the elevated ridge and plateau terrain — where a deep unsaturated zone allows strong atmospheric draw and weak drainage — to the low-lying coastal fringe where the shallow water table limits atmospheric influence and enhances lateral drainage. Both zones carry Corsican pine canopy; the difference is substrate and topographic position.

Two individual wells merit comment. CEH14, at the highest elevation in the forest zone (14.4 m), has the highest β_2 in the network (3.759) — entirely consistent with the elevation regression rather than being an outlier. Its negative β_3 (−0.019) and low R^2 (0.568) suggest that at this ridge-proximal position, lateral recharge from the bedrock ridge at high head overwhelms the drainage signal, reversing the expected relationship between displacement and head change. NW10, the broadleaf block well, has the highest β_1 in the forest zone (3.526, 1.2 SD above the C4 mean) and the highest β_3 among C4 wells (0.046). Its elevated recharge sensitivity likely reflects the deciduous canopy intercepting less winter rainfall than the surrounding Corsican pine (15% vs 24% interception), though this cannot be confirmed with a single well.

Coefficient	vs Elevation r (p)	vs Dist. ridge r (p)	vs Easting r (p)
β_1 (recharge)	0.223 (0.444)	−0.528 (0.052)	0.590 (0.026)
β_2 (atm. draw)	0.975 (<0.001)	−0.889 (<0.001)	0.740 (0.003)
β_3 (drainage)	−0.771 (0.001)	0.577 (0.031)	−0.406 (0.150)

Table 17. *Per-well spatial predictors of SSM coefficient variation within the forest zone (C4 Main Forest and C5 Coastal Forest, n = 14). Pearson correlation coefficients and p-values for three spatial variables. Elevation dominates* β_2 *(*$R^2 = 0.951$*);* β_1 *shows no significant elevation dependence. Source: 10c_forest_zone_analysis.py (10c_04_forest_zone_summary.txt).*

The per-well analysis confirms that the C4/C5 partition in the Ward's solution reflects a genuine substrate and elevation discontinuity rather than an arbitrary split within a continuous canopy population. The mechanistic interpretation of this finding — and its implications for forest management — is developed in Section 5.6.1.

4.9.5 Mean Water Table Surface and Groundwater Flow Direction

The mean annual water table surface was constructed by IDW interpolation of per-well mean maOD values (2005–2026) to a 50 m grid on the British National Grid, masked to the rectangular sea-boundary extent (Figure 44). No DEM-height mask is applied because the water table is a continuous surface independent of the overlying topography. The highest mean water levels (~14 m AOD) occur in the north-western forest zone adjacent to the bedrock ridge, declining to 7–9 m AOD across the central dune and reaching 2–3 m AOD at the southern and eastern coastal margins. Contours run approximately NE–SW across the site, consistent with the topographic configuration. Normalised Darcy flow direction vectors, computed as $Q = -K \cdot b \cdot \nabla h$ (Bear, 1972; Freeze and Cherry, 1979) from the head gradient field using $K = 6.0\ \mathrm{m\ day^{-1}}$ (Betson et al., 2002), show a radial pattern diverging from the north-western high point: predominantly southward and south-westward in the western half of the site, and south-eastward in the eastern half. This pattern is broadly consistent with the DEM-derived flow network (blue lines in Figure 44; Conrad et al., 2015; Natural Resources Wales LiDAR DTM, 2023). The flow vectors are reported as directional context for the interpolated head field rather than as output of a calibrated continuous-flow model.

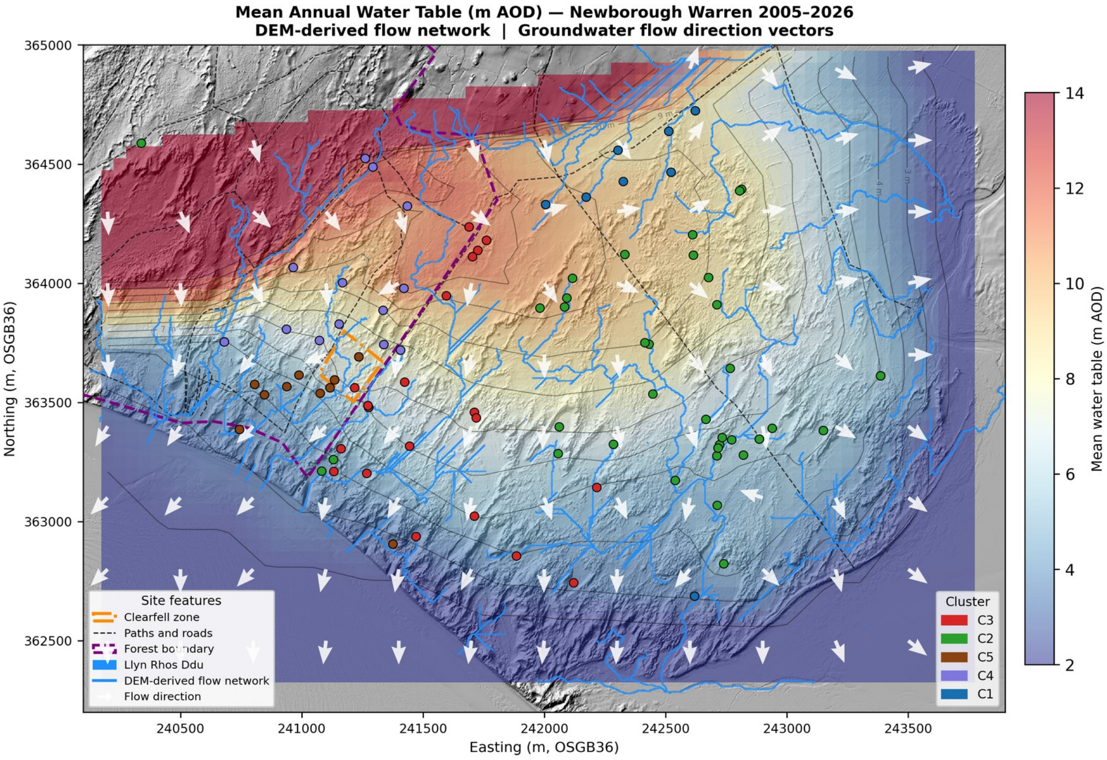

***Figure 44**. Mean annual water table elevation (m AOD) across the Newborough Warren dune aquifer, 2005–2026, derived by IDW interpolation of per-well mean annual heads to a 50 m grid. The highest water levels (~14 m AOD) occur in the north-western forest zone adjacent to the bedrock ridge, declining toward the southern and eastern coastal margins (~2–3 m AOD). Normalised groundwater flow direction vectors (white arrows; direction only, magnitude not scaled) are derived from the head gradient field $Q = -K \cdot b \cdot \nabla h$, using $K = 6.0$ m/day (Betson et al., 2002) and the IDW aquifer thickness surface. Well symbols are coloured by cluster. Background: LiDAR greyscale hillshade (LiDAR data: © NRW & OS; see Section 10 for full attribution). Generated by 20_spatial_figures.py. (20_head_surface_streams.png)*

4.9.6 Water Balance Residual Field

The SSM water balance residual field (Figure 45) quantifies the degree to which the three-term SSM closes the long-term water balance at each well. The residual at each reference well is computed as $\alpha = (\beta_2 \cdot \overline{\mathrm{PET}} + \beta_3 \cdot \bar{h}\text{\_disp}) - \beta_1 \cdot \bar{P}\text{\_eff}$, where overbars denote long-term monthly means, $\bar{h}$_disp is mean displacement above the 3.7 m drainage datum, and $\bar{P}$_eff applies the canopy interception correction at forest wells (Section 3.8). Positive values indicate locations where modelled losses exceed modelled recharge, implying an unmodelled input such as lateral recharge from the bedrock ridge. The field shows a broad spatial gradient from positive residuals in the north-west to near-zero or weakly negative residuals in the south-east. No DEM-height mask is applied. The strongest positive residuals (>+0.02 m/month) are concentrated along the northern forest margin and ridge flank, where measured rainfall-driven recharge alone does not fully account for the observed water levels. The largest individual residual occurs at CEH14 (+0.031 m/month), consistent with the lateral recharge from the bedrock ridge inferred from its anomalous negative β_3 (Section 4.9.4). The magnitude declines southward and eastward across the open dune, where residuals approach zero — indicating that the three-term SSM closes the water balance adequately without additional input. A number of wells in the south-eastern margin show weakly negative residuals, indicating slight over-prediction by the model in these lower-head positions. The independently computed Darcy flow vectors derived from the mean head gradient field are consistent with a predominantly south-eastward lateral flux from the ridge toward the coast, providing directional context for interpreting the residual pattern.

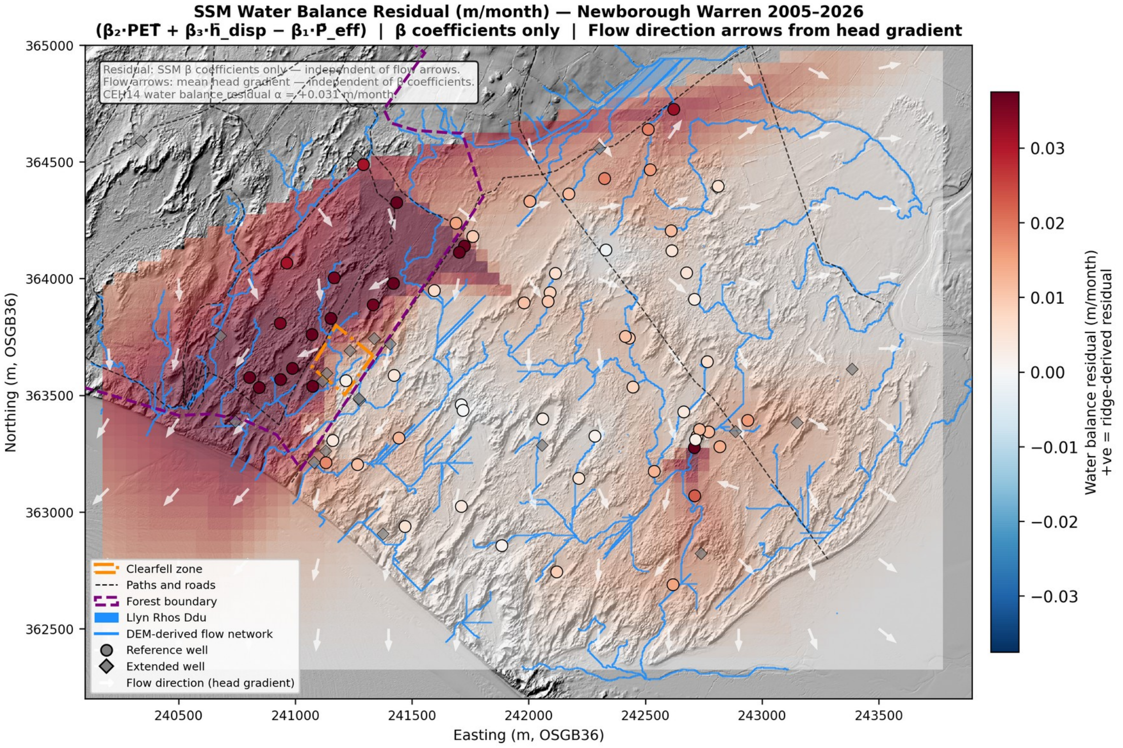

Figure 45. *SSM water balance residual across the Newborough Warren dune aquifer, 2005–2026. The residual at each reference well is computed from β coefficients only as* $(\beta_2 \cdot \bar{PET} + \beta_3 \cdot \bar{h}\_disp) - \beta_1 \cdot \bar{P}\_eff$ *and IDW-interpolated to a 50 m grid with sea boundary anchors at zero and dune ridge cells masked. Positive values (red) indicate locations where measured rainfall-driven recharge does not fully account for observed water levels; negative values (blue) indicate slight over-prediction by the model. The strongest positive residuals are concentrated along the northern forest margin and ridge flank; residuals decline to near zero across the open dune and become weakly negative at some south-eastern margin wells. White arrows show normalised Darcy flow direction vectors derived from the mean head gradient field (independent of the β coefficient residual calculation). Well symbols are coloured by residual value; circles indicate reference wells, diamonds indicate extended wells. The largest individual residual occurs at CEH14 (+0.031 m/month), consistent with the lateral recharge from the bedrock ridge inferred from its anomalous negative* β_3 *(Section 4.9.4).Background: LiDAR greyscale hillshade. (LiDAR data: © NRW & OS; see Section 10 for full attribution.) Generated by 20_spatial_figures.py (20_residual_ssm.png).*

4.10 Scenario Analysis

The per-well SSM coefficients derived in Section 4.9 were used to project equilibrium water table responses under prescribed climate and forest management perturbations, following the framework of Section 3.8. Six scenarios were evaluated: baseline; UKCP18 2050s and 2080s climate trajectories; full clearfell; 50% thinning; and broadleaf conversion. Scenario difference maps were computed as Δh = h(scenario) − h(baseline) for each well, season and cluster. The full per-scenario, per-season, per-cluster Δh output is available through the Newborough Warren Hydrogeological Modelling Suite (Hollingham, 2026b), an interactive web-based scenario viewer that renders the per-well Δh field across the site using IDW interpolation, with a complementary depth-below-surface view masked to hydrologically meaningful slack positions and overlaid with the Curreli et al. (2013) SD15b/SD16 ecological thresholds. The cluster-level results are summarised below.

4.10.1 Climate Scenario Projections

Climate scenario responses are spatially coherent and physically consistent with the per-cluster β coefficient structure. Under UKCP18 2050s forcing, the site-wide mean annual Δh is −0.018 m/month, with cluster-level annual responses ranging from −0.026 m/month at C4 (Main Forest) to −0.009 m/month at C1 (Lake Edge). Under UKCP18 2080s forcing, the site-wide mean annual response deepens to −0.032 m/month, with a maximum of −0.046 m/month at C4 and a minimum of −0.017 m/month at C1. Both UKCP18 scenarios show the forested clusters (C4, C5) carrying the largest annual declines, driven by the combination of high β_2 (atmospheric draw sensitivity) and reduced net recharge under interception. The annual response conceals a pronounced seasonal asymmetry. Under both trajectories, winter Δh is positive at all clusters (site-wide mean +0.028 m/month at 2050s and +0.055 m/month at 2080s) as the increased winter rainfall under RCP8.5 Wales outpaces the concurrent increase in winter PET. Summer Δh is strongly negative and carries the annual signal: site-wide means of −0.064 m/month at 2050s and −0.120 m/month at 2080s, reaching −0.131 m/month at C2 under the 2080s trajectory. This summer-dominated signature reflects the same physical mechanism established in the observed record (Section 4.8.1): at Newborough, hydrological deterioration under climate change operates through intensified summer atmospheric demand rather than declining winter recharge.

4.10.2 Forest Management Scenarios

Forest management scenarios produce responses at the two forested clusters (C4 and C5), with no effect at C1–C3. Full clearfell is modelled by setting canopy interception to 0% and adjusting β_2 at C4 and C5 wells. Fifty percent thinning halves interception from 24% to 12% with proportional β_2 adjustment. Broadleaf conversion replaces the Corsican pine canopy interception fraction of 24% (Freeman, 2008) with an annual-mean deciduous interception of 15% (Komatsu et al., 2011), and applies a seasonally varying β_2 adjustment to reflect deciduous phenology: β_2 is held at baseline during the winter dormant season (leaves absent November–March) and increased by a summer multiplier of ×1.113 during the growing season, reflecting the comparable transpiration rates of mature deciduous and Corsican pine canopy in full leaf.

Under full clearfell with the BACI-corrected upper-bound β_2 multiplier (×1.108; Section 3.8), the framework predicts a modest annual water table rise at both forest clusters (C4: +0.028 m/month head, equivalent to +4.7 mm w.e./month volumetric; C5: +0.034 m/month head, +8.8 mm w.e./month). The annual response is driven primarily by the winter recharge gain from canopy removal (C4: +0.043 m/month head; C5: +0.043 m/month head). The summer response is positive but small at both clusters (C4: +0.013 m/month; C5: +0.024 m/month), reflecting that the winter interception gain marginally exceeds the summer β_2 cost across the year at both substrates, with C5's lower β_2 sensitivity (median 1.155 vs C4's 2.618) producing the larger positive net summer response. Fifty percent thinning produces a similar pattern at reduced magnitude (C4: +0.014 m/month annually, +2.3 mm w.e./month; C5: +0.017 m/month, +4.4 mm w.e./month). These are steady-state predictions from the SSM coefficient structure; the predicted magnitudes are consistent with the BACI monitoring record, which detects a significant mean monthly recovery (+93 mm Impact, +153 mm Edge against the Forest control; Section 4.6.2) but no corresponding improvement in summer minima (Section 4.6.4; further interpreted in Section 5.5).

Broadleaf conversion produces 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 (C4: +0.026 m/month; C5: +0.021 m/month) reflecting the near-zero interception under leafless canopy during the recharge season, but the summer response is negative at C4 (−0.007 m/month) and near-zero at C5 (+0.004 m/month) as the growing-season transpiration penalty under mature deciduous canopy (β_2 summer multiplier ×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 additional winter recharge steepens the local hydraulic gradient and drains away through the highly transmissive sand before the summer minimum is reached (further interpreted in Section 5.5.2).

Site-wide means under forest management are small (annual +0.006 m/month clearfell, +0.003 m/month thinning) because only 14 of 66 reference wells are in the two forest clusters and the framework applies no spatial propagation from the forest zone into other clusters. Broadleaf conversion produces no meaningful site-wide signal.

5. Discussion

The 27-step analytical pipeline applied to 21 years of monitoring data across 89 wells yields a consistent picture of a coastal dune aquifer whose water balance is governed by the interplay of rainfall recharge, atmospheric evaporative demand and head-dependent drainage, with land cover operating as a surface boundary condition on an otherwise uniform sand substrate. Five hydrogeological clusters emerge from the data, each with distinct mechanistic coefficients but broadly convergent specific yield once canopy interception is corrected. The ecological viability of the dune slack system depends on a narrow 37 cm window of summer minimum water table depth (Curreli et al., 2013; Section 4.7.4), and summer minimum depth — not winter rainfall — is the dominant predictor of winter flooding potential across the western and forest clusters (Table 13). This seasonal asymmetry is the organising principle of the discussion that follows.

The cluster structure links groundwater behaviour to geology and land cover, establishing the physical contrasts that determine how each part of the aquifer responds to climate forcing (Section 5.1). The SSM framework and its benchmarking against the traditional linear model demonstrate that head-dependent drainage memory is essential for long-range prediction (Section 5.2), and the water balance decomposition shows how recharge, atmospheric draw and drainage partition differently across the five clusters under the displacement formulation (Section 5.3). These foundations set the context for evaluating the two management interventions in the monitoring record.

Dune scraping produced measurable improvements in summer minimum depth at CEH36, with a benefit that persisted through the post-felling period; the durability of these gains is set by the background climate trajectory and by interactions with adjacent interventions (Section 5.4). The December 2017 clearfell produced a significant mean water table recovery against the Forest control (+93 mm Impact, +153 mm Edge; Section 5.5). The recovery is concentrated in the mean monthly signal: summer minima at the felled wells tracked the unfelled Forest controls, consistent with the canopy acting simultaneously as a hydrological sink (recharge suppression through interception) and as a microclimatic buffer on atmospheric demand at the water table (Section 5.6). The spatial framework and BACI record together indicate that forest management perturbations propagate south-westward toward Caernarfon Bay rather than south-eastward into the C1 and C2 clusters (Sections 5.7, 5.8).

The discussion brings these threads together into a management synthesis (Section 5.9): the operational domain for direct hydrological intervention, the role of canopy management in retaining the plantation's buffering function while partially recovering the recharge penalty, and the climate trajectory that sets the binding constraint on all outcomes.

5.1 Linking Groundwater Behaviour to Geology and Land Cover

The analytical framework rests on two foundational datasets: a 95-year climate record from RAF Valley and a 21-year dipwell network spanning 89 wells across the full geomorphological range of the dune system. The climate record establishes that

Newborough Warren operates under a persistent annual surplus (mean P/PET = 1.36; Table 1) in which inter-annual water balance variability is driven almost entirely by precipitation rather than by evaporative demand — PET shows strong seasonality but negligible inter-annual trend (Section 4.1.1, Figure 3). This surplus is not seasonally uniform: winter rainfall exceeds PET by a factor of three or more, while summer PET exceeds rainfall, concentrating atmospheric demand on the aquifer during the period of lowest recharge input. The seasonal opposition is the primary climatic driver of dune slack hydrology, and the post-2013 period has run substantially warmer than the pre-2013 summer temperature baseline (Section 4.1.1, Figure 4), intensifying the summer deficit. The monitoring network itself is unusually dense for a dune site — 66 reference wells with a median record length of 12.5 years, supplemented by 21 extended-network wells — providing the spatial resolution needed to distinguish cluster-scale responses from site-wide climate forcing. The network mean water level explains approximately half the variance from cumulative climate forcing alone (Figure 3), with the residual variance reflecting the spatial heterogeneity, boundary subsidies and land cover effects that the subsequent analyses decompose.

The hierarchical clustering solution (Section 3.2; Section 4.2, Figures 6 and 7) identifies five spatially coherent and physically interpretable hydrogeological units that collectively explain the principal axes of groundwater variability at Newborough Warren. The fundamental binary divide between the Eastern Block (C1, C2) and Western Residual (C3) reflects a well-documented contrast in sub-surface stratigraphy. Shallow till and estuarine deposits in the east constrain storage capacity and drive rapid, flashy water table responses; conversely, deep sand accumulation in the west supports a capacious, heavily buffered aquifer (Stratford et al., 2007; Grootjans et al., 2004). The LCSC values (Table 3) quantify this physical contrast: the Eastern Block clusters show low LCSC values reflecting their constrained storage, the Western Residual is intermediate, and the Forest clusters show the highest values reflecting the combined effect of deep sand substrate and canopy interception.

LCSC is closely related to specific yield but is not identical to it: the SSM's β_1 coefficient absorbs canopy interception losses alongside true storage response, so that at the forested clusters the elevated LCSC partly reflects the 24% rainfall interception by the Corsican pine canopy (Freeman, 2008) rather than a difference in aquifer drainability. After applying canopy interception correction, the WTF-derived cluster-aggregate specific yields for the forested clusters fall within the range of the open dune clusters (Section 4.2.4, Table 4c), indicating that the forest clusters' elevated LCSC reflects canopy boundary conditions at the surface rather than differences in subsurface storage architecture. The LCSC framework translates a qualitative geological description into a management-relevant metric expressing the precipitation depth required to achieve a unit rise in the water table: open-dune wells respond to rainfall more efficiently than forested wells, the latter reflecting canopy interception rather than storage capacity (Table 3).

The plantation-covered forest clusters (C4 and C5) share the deep sandy substrate of the Western Block, yet segregate into distinct hydrogeological units (Figure 12). Land cover therefore modifies the hydrological signature of the underlying geology sufficiently to produce a distinct cluster signature. C4 occupies the highest mean water table elevation in the network yet shows the weakest drainage decay coefficient in the network (Table 3), while C5 shows an intermediate drainage rate consistent with its position between the forest interior and the coastal margin. C4's combination of high topographic position and slow drainage reflects two reinforcing mechanisms: canopy interception dampens the recharge pulse so that the water table never rises sufficiently above its suppressed equilibrium to generate the strong head-dependent drainage response seen in the open dune clusters, and the irregular bedrock topography on the elevated ridge flank impedes lateral drainage, with buried ridges creating within-cluster heterogeneity in β_3 (Section 4.9.4). The displacement formulation (Section 3.4.1), which references drainage to a datum 3.7 m below the pipe top rather than to the ground surface, resolves C4's β_3 to a clearly significant value (Table 3), indicating that head-dependent drainage operates at all five clusters when the hydraulic reference level is correctly specified.

The cluster-level physical reality is independently cross-validated by the state-space model parameters (Section 4.2.2, Table 3). The Western Block's lower recharge sensitivity reflects the deep unsaturated sand buffer that attenuates the rainfall-to-water-table transfer. The Eastern Block's higher β_1 and β_3 values are consistent with a fill-and-spill mechanism in which shallow till promotes surface ponding and rapid lateral overflow rather than slow drainage through the aquifer matrix. The structural divide is evident in the winter flooding record: C1 and C2 reached the wet slack flooding threshold in most years, C3 in a small fraction, and C4 and C5 in none across the monitoring record (Section 4.8.2, Figure 38). The consequences of this asymmetry for summer minimum trajectory and forecasting are developed in Section 5.2; the anomalous C5 decline is discussed in Section 5.7.2.

The implication for conservation management is that Eastern Block slacks respond sharply to individual rainfall events but retain little water between events, making them sensitive to both exceptional winter rainfall and prolonged summer drought in ways that the buffered Western Residual is not. The post-2018 amplification of C1 seasonal amplitude (Section 4.2.1) is consistent with this sensitivity: the post-2018 period contained two of the three drought summers identified in the record (2018, 2022), producing deeper summer minima at the lake-buffered cluster. Because the entire ecological gradient of dune slack vegetation is compressed into a 37 cm window of summer minimum depth (Curreli et al., 2013; Section 4.7.4), any management intervention must be evaluated primarily on its effect on summer minima rather than on annual-mean water balance metrics.

The cluster structure is independently characterised by the aquifer diagnostic synthesis (Section 4.9.3, Figures 42 and 44), in which three parameters derived independently of the clustering — drainage timescale ($\tau = Sy/\beta_3$), SSM forecasting gain over the TLM (ΔNSE), and WTF-derived specific yield — place the five clusters in distinct, non-overlapping regions of diagnostic space. C1 Lake Edge occupies the low-τ, high-ΔNSE, low-Sy corner (shallow, fast-cycling, drainage-dominated), while C4 Main Forest occupies the opposite extreme (high-τ, low-ΔNSE, low-Sy — sluggish, impeded by irregular bedrock). C5 Coastal Forest separates cleanly from C4 on both axes despite sharing the same canopy, consistent with substrate as a major contributor to the diagnostic signature within the forest zone (Section 5.3.1). This model-based characterisation of aquifer architecture is consistent with, but derived independently of, the Ward's clustering applied to the raw hydrograph data.

5.1.1 Cluster Boundaries and the Pearson Affinity Validation

The gradational nature of the aquifer between these structural units is shown by the Pearson affinity analysis (Section 3.3; Section 4.3). Of the 66 reference wells, the great majority confirmed their primary cluster membership; three Spy wells showed strongest temporal affinity with a neighbouring cluster at borderline correlation deltas, consistent with gradational boundaries rather than misclassification. The Multi-Cluster Affinity flag was met at a substantial fraction of wells, reflecting the high mutual correlation among the five centroids in a contiguous, highly transmissive sand aquifer (Rao and Srinivas, 2006). The spatial concentration of MCA flags along the C1/C2/C3 transitional zone, the C2/C3/C5 boundary where the Dune, Western Residual and Coastal Forest signatures converge, and the C3/C4/C5 forest–dune boundary (Figure 12) identifies the areas where geological and vegetative controls overlap. Hydrogeologically, these are candidate locations for rapid state change in response to sustained conservation management or climatic intervention (Grootjans et al., 2004).

The persistence of Forest cluster affinity at the clearfell treatment wells (Section 4.3) admits at least three interpretations: residual suppression by the surrounding intact canopy extending beyond the felling boundary; incomplete hydrological recovery within the felled footprint itself; or, most fundamentally, 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. Distinguishing between these mechanisms requires continued monitoring beyond the current record.

5.2 Boundary Conditions, Model Limitations and Forecasting

The SSM is a lumped mass-balance model fitted independently at each well. It does not resolve lateral groundwater flow, and its three-term structure — recharge ($\beta_1 \cdot P$), atmospheric draw ($\beta_2 \cdot PET$) and head-dependent drainage ($\beta_3 \cdot h\_disp$) — necessarily omits any flux that does not covary with rainfall or PET at monthly resolution. Where the three-term balance does not close, the residual appears as a structural mismatch between predicted and observed water levels.

5.2.1 Water Balance Residuals and Boundary Subsidies

The water balance residual field (Section 4.9.6, Figure 45) makes this mismatch spatially explicit. Positive residuals — indicating that measured recharge alone does not fully account for observed water levels — are concentrated along the northern forest margin and ridge flank. Residuals decline to near zero across the central and southern open dune, where the SSM closes the water balance adequately, and become weakly negative at some south-eastern margin wells. The independently computed Darcy flow vectors derived from the mean head gradient field (Section 4.9.5, Figure 44) are consistent with a predominantly south-eastward lateral flux from the ridge toward the coast, providing directional context for interpreting the residual pattern. The residual field is a structural diagnostic of where the SSM's three terms are insufficient; it should not be read as a quantified map of lateral inflow, because the residual also absorbs nonlinear soil-moisture storage dynamics, β coefficient sensitivity and any systematic bias in the Thornthwaite PET estimate. The spatial coincidence of the largest positive residuals with the ridge boundary makes ridge-derived lateral input the most parsimonious candidate for the unmeasured flux, but two independent diagnostic tests — a cross-correlation lag analysis testing for distance-dependent transport from the ridge (Spearman $\rho = +0.010$, $p = 0.947$, $n = 50$; Supplementary Note S5) and a seasonal climatology analysis testing for unmodelled summer evaporative demand (network mean $r = -0.028$ against sunshine hours; Supplementary Note S6) — returned null results for both candidate mechanisms. The spatial pattern is real; the mechanistic attribution is unresolved.

The seasonal residual diagnostic (Script 24) provides a further constraint. The dominant phase of the SSM residual peaks in winter and early spring (47 of 63 wells), with no wells peaking in summer. This is not the signature of unmodelled summer evapotranspiration — if Thornthwaite PET were systematically underestimating summer losses, the residual would peak in summer months. The winter phase is consistent with either nonlinear threshold recharge behaviour — in mid-winter with saturated soils, rainfall reaches the water table with higher efficiency than the cluster-mean β_1 represents, producing a positive residual during the months of greatest recharge — or ridge-derived lateral input, which would also be expected to peak during the months of greatest rainfall on the ridge itself. The sunshine-hours correlation (network mean $r = -0.028$, all 63 wells within the Bartlett null band) indicates that the residual is not dominantly unmodelled ET; any systematic bias in the Thornthwaite method that correlates with sunshine hours has been absorbed by β_2. The residual's mechanistic attribution remains unresolved, but the seasonal diagnostics narrow the field: the unmeasured flux operates in winter, not summer, which is consistent with either ridge-derived lateral recharge or a threshold recharge mechanism but not with missing evaporative demand.

5.2.2 Displacement Formulation and Depth-Dependent PET

The displacement formulation (Section 3.4.1) — which references drainage to a datum 3.7 m below the pipe top rather than to the ground surface — captures head-dependent drainage at all five clusters, with β_3 significant at $p < 0.01$ or better throughout the network (Table 3). Under the pre-displacement specification, C4's drainage term was borderline; the displacement formulation resolves this by correctly specifying the hydraulic reference level against which drainage operates (Section 5.1). The per-well datum sensitivity analysis (Section 4.9.1, Figure 39) shows that the uniform 3.7 m datum produces negligible R^2 loss relative to per-well optimised datums across most of the network, with meaningful gains limited to C1 Lake Edge wells where the effective drainage base is shallowest.

A depth-dependent extension of the SSM, in which the β_2 coefficient declines exponentially with water table depth according to $\beta_2(d) = \beta_2 \cdot \exp(-\lambda d)$, was evaluated at cluster level (Section 3.7.1). The depth-dependent formulation improves iterative NSE at all five clusters (ΔNSE +0.025 to +0.182), consistent with capillary connectivity between the water table and root zone diminishing with depth — a real physical effect that the standard SSM's fixed β_2 absorbs as a depth-averaged effective parameter. The improvement is strongest at C1 Lake Edge ($\lambda = 1.90\ \text{m}^{-1}$, $\Delta\text{NSE} = +0.182$), where shallow water-table evaporation near the lake is the dominant ET pathway, and weakest at C4 Main Forest ($\lambda = 0.20\ \text{m}^{-1}$), although C4's narrow depth range means even weak coupling produces a measurable NSE gain (+0.104). The standard SSM is retained as the primary operational model because its fixed β_2 preserves the closed-form threshold equations and scenario algebra on which the operational outputs depend. The mechanistic implications — particularly the finding that C4 shows the weakest depth coupling despite having the highest β_2 — are discussed in Section 5.6.1.

5.2.3 Predictive Skill and Operational Forecasting

The physical distinction between the SSM and the traditional linear model (TLM) is most directly revealed by long-range iterative simulation (Section 3.4.2; Section 4.4). In one-step diagnostic mode both models perform almost identically (median R^2 0.91 and 0.92 respectively; Table 5), because each prediction is corrected with the observed water level from the previous month. The divergence is substantial under iterative simulation: the TLM achieves a median iterative NSE of 0.14, while the SSM achieves 0.76 — a median gain of +0.62 NSE units. The SSM returned a positive iterative NSE at 65 of 66 wells compared with 42 of 66 for the TLM (Table 5; Figure 14). The explicit β_3 term in the SSM captures storage decay memory, preventing the cumulative drift that otherwise causes TLM divergence. The Main Forest cluster (C4) shows the smallest improvement (+0.229), consistent with two non-exclusive readings: the canopy may smooth the recharge signal into a near-linear climate response that the TLM can handle adequately without a storage decay term, and the structurally suppressed water table at C4 carries less nonlinear storage dynamics for the SSM to capture. The two readings are not distinguishable from the NSE benchmark alone.

This predictive capacity has three practical consequences for conservation management. First, the SSM can generate forward projections of water table trajectory under specified climate scenarios (Section 4.4), enabling managers to anticipate when summer minima may cross the ecological viability thresholds identified by Curreli et al. (2013) rather than relying solely on backward-looking trend extrapolation. The seasonal prediction equations (Table 13) and P_flood threshold equations (Table 16) are algebraic inversions of the SSM's fitted coefficients — they are not separate models but direct consequences of the SSM parameterisation, and their predictive credibility rests on the SSM's demonstrated iterative skill.

Second, the spatial pattern of SSM improvement delineates the aquifer's structural boundaries without invasive coring (Figure 14). The largest NSE gains concentrate in the C1 Lake Edge sector, where substantial seasonal head variation and head-dependent drainage to Llyn Rhos-Ddu make the storage decay term essential. The smallest gains occur in C4 Main Forest. The improvement map is itself a diagnostic: it identifies which parts of the aquifer are governed by nonlinear storage dynamics and which approximate a linear climate response.

Third, the SSM framework is operationally transferable. The model requires only three fitted parameters per well (β_1, β_2, β_3), a uniform drainage datum, and monthly rainfall and temperature from a single climate station. It does not require hydraulic conductivity measurements, aquifer thickness estimates, or lateral flow modelling — the parameter-sparsity constraints that produced calibration difficulties in the Betson et al. (2002) MODFLOW model at this site. Any coastal dune aquifer with a monitoring network and a nearby climate station can be parameterised in the same framework. The displacement formulation (Section 3.4.1), which references drainage to a fixed datum rather than the ground surface, is the methodological choice that makes this possible: it captures the depth-dependence of the drainage response without requiring site-specific calibration of the drainage geometry.

5.2.4 Seasonal Prediction and the Intervention Window

The demonstration that winter rainfall amount is a poor predictor of peak flood conditions in the Western Residual (P_winter $p = 0.057$; Table 13), whilst antecedent summer minimum depth is dominant ($p = 0.010$), is the most consequential single finding for conservation management. It means that managers cannot reliably forecast slack flooding from seasonal rainfall forecasts alone — what matters is how deep the aquifer drains each summer. The seasonal prediction equations — which regress peak winter water table against antecedent summer minimum depth and cumulative winter rainfall (Section 3.6.2; Table 13) — were fitted with one equation per cluster under the $k = 5$ partition: winter rainfall was significant in the two eastern clusters (Lake Edge $p = 0.002$, Dune $p = 0.005$) but non-significant in the Western Residual ($p = 0.057$), Forest ($p = 0.710$) and Coastal Forest ($p = 0.361$). Antecedent summer minimum depth was the dominant predictor in the three western and forest blocks (Western Residual $p = 0.010$, Forest $p < 0.001$, Coastal Forest $p < 0.001$). The ecological stakes of this finding are clarified by the eco-hydrological thresholds established by Curreli et al. (2013), whose four-year monitoring programme showed that wet slack vegetation (SD15b) requires mean annual maximum water levels above approximately −0.10 m relative to ground surface, and dry slack communities (SD16) persist at winter maxima around −0.25 m. Summer minimum depth sets the precondition for achieving these winter maxima: a water table that has drained deeply through summer must receive proportionally greater winter recharge to recover to flooding height. The P_flood equations quantify this precondition directly (Section 3.6.3; Table 16).

The summer minimum trajectory (Section 4.8.2, Figure 38) confirms that this precondition is deteriorating. All five clusters show declining summer minimum trends, with C1 ($p = 0.048$) and C5 ($p = 0.003$) statistically significant, C2 marginal ($p = 0.084$), and C3 ($p = 0.273$) and C4 ($p = 0.371$) non-significant. The post-2013 period is running 0.94°C above the pre-2013 temperature baseline (Section 4.1.1), and the trend line extrapolations indicate C1 summer minima crossing the dry slack viability threshold around 2030–2032 under current trajectories. Recent ecological field surveys provide direct observational validation: Callaghan et al. (2021) found that despite active topographical intervention at Newborough, Petalophyllum ralfsii populations remained restricted due to limited water availability, and recorded a complete absence of winter flooding following the severe 2018 summer drought. The intervention window is convergent with projections from other west-coast UK sites: Curreli et al. (2013), applying their community thresholds to the Ainsdale water-table projection of Clarke and Sanitwong Na Ayutthaya (2010), concluded that conditions supporting wet slack vegetation may become non-viable by the 2030s. The intervention window is therefore likely to be measured in one to two decades rather than the longer timeframes typically assumed in dune restoration planning.

The Eastern Block (C1, C2), by contrast, is better characterised as near-perennial: the shallow till-floored aquifer regularly reaches or approaches the surface in wet winters and flooding is driven by within-season rainfall accumulation rather than antecedent drought depth.

The Forest (C4) prediction equation ($R^2 = 0.871$, summer minimum dominant, winter rainfall non-significant $p = 0.710$; Table 13) is consistent with two non-exclusive interpretations. Under the canopy interception interpretation, summer drought conditions were attenuated by the canopy, producing a more compressed range of summer minima and consequently more predictable winter peaks. Under an alternative interpretation, the forest water table sits sufficiently deep that the winter recharge pulse cannot reach flooding height across the rainfall range in the record, so the regression cannot detect a winter rainfall signal. The two interpretations are not mutually exclusive — both contribute to the statistical pattern — and the available data do not allow them to be partitioned. The P_flood unreachability finding (Section 5.8.1) — that at C4 summer minima of approximately 1.8 m below ground, the required winter rainfall to achieve flooding exceeds anything in the 95-year RAF Valley record — directly supports the second interpretation. The BACI summer minima null (Section 5.5) is also consistent with a system sitting far below threshold, where neither canopy presence nor canopy removal materially affects whether winter flooding is reached. The BACI monitoring record (Section 4.6) documents increased seasonal amplitude in the impact zone following the December 2017 clearfell, but this signal is confounded by the April 2015 scraping at CEH36 and the wet 2015–16 baseline, and the analysis was unable to isolate a clearfell-specific response above these confounds (Section 4.6.7). The implications of this result for forest management are developed in Section 5.5. The Coastal Forest (C5) prediction equation shows a similar pattern ($R^2 = 0.881$, h_min $p < 0.001$, P_winter $p = 0.361$), consistent with both interpretations operating at C5 as well.

5.3 Recharge, Loss and the Water Balance

The cluster-level water balance decomposition (Section 3.7.2; Section 4.2.3; Tables 3a, 3b) makes explicit a structural feature that is implicit throughout the state-space model results: under mean climate conditions, the three-term displacement-formulation SSM closes the monthly water balance at the cluster mean level, with residuals of less than 2.5% of total losses (Table 4a). This near-closure indicates that recharge, atmospheric draw and head-dependent drainage together account for the principal water balance fluxes at the cluster aggregate level, with any unmeasured boundary contributions absorbed into the fitted β coefficients (Section 5.3.1) rather than appearing as systematic residuals at the cluster mean. Spatial variation in local closure across the site is examined in Section 5.2.1.

5.3.1 Recharge Partitioning and Atmospheric Draw

Recharge into the water table ($\beta_1 \cdot \bar{P}$) declines systematically from the Lake Edge cluster (C1: 0.341 m/month) through the Dune (C2: 0.288 m/month) and Western Residual (C3: 0.266 m/month) to the forested clusters (C5: 0.181 m/month; C4: 0.187 m/month), reflecting increasing storage capacity and — in C4 and C5 — the additional suppression of throughfall by canopy interception, which is captured implicitly in the lower β_1 rather than as a separate model term. Total losses (atmospheric draw plus drainage) track recharge closely at every cluster, ranging from 0.186 m/month at C4 to 0.345 m/month at C1. The drainage partition varies markedly: drainage accounts for 85% of total losses at C1 (reflecting rapid lateral exchange with Llyn Rhos-Ddu) but only 26% at C4, where atmospheric draw dominates. C2 and C3 show drainage-dominated partitions (67%/33% and 63%/37% respectively), while C5 is intermediate (59%/41%).

It should be noted that in open-boundary systems, lumped parameter models of this form are known to absorb unmeasured boundary fluxes into the rainfall coefficient when no intercept term is fitted (Young, 2011; Knotters and van Walsum, 1997), which means the recharge term ($\beta_1 \cdot P$) should not be read as a direct measure of local precipitation reaching the water table. The water balance residual field (Section 4.9.6, Figure 45) shows that the three-term balance does not close uniformly across the site: positive residuals concentrate along the northern forest margin and ridge flank, declining to near zero across the open dune and becoming weakly negative at some south-eastern margin wells. This spatial pattern is consistent with a ridge-derived lateral input that the SSM's β_1 term partially absorbs but does not fully resolve. However, as discussed in Section 5.2, the mechanistic attribution of this pattern is unresolved — the spatial coincidence with the ridge boundary is suggestive but the lag-distance test did not corroborate a specific transport mechanism, and an independent test for unmodelled evaporative demand using sunshine hours also returned a null result.

Thornthwaite PET is an empirical index of atmospheric evaporative demand derived from temperature and day length; it does not define a specific reference surface in the manner of Penman-Monteith, so a formal crop coefficient decomposition does not apply. The β_2 coefficient under the displacement-formulation SSM therefore absorbs the combined effects of vegetation transpiration, canopy interception, specific yield and depth-coupling into a single effective scaling between atmospheric demand and water table response (Section 5.6.1). C4 and C5 carry broadly similar Corsican pine canopy; the principal difference between them is substrate and topographic position, with C5 occupying the coastal-fringe sand and C4 occupying the elevated ridge flank. The substantial β_2 difference between the clusters — C4 considerably higher than C5 — points to substrate as a major contributor to the cluster-level contrast. The WTF-derived Sy estimates (Table 4c) show C5's Sy higher than C4's, supporting the direction of this interpretation: a higher Sy at C5 means a given evaporative flux produces a smaller head change at the water table, which would lower the apparent β_2. The magnitude of the β_2 ratio is larger than the WTF Sy ratio alone would predict if substrate were the only factor, indicating that depth-coupling differences (Section 5.2.2) also contribute. C4's position on the elevated ridge flank, where irregular buried bedrock creates a heterogeneous substrate of thin sand and deeper troughs, is consistent with reduced effective drainable volume per unit head change. The drainage timescale analysis (Figure 41) reinforces this picture: C4 wells show the longest τ values in the network, indicating that lateral discharge is impeded by subsurface barriers rather than flowing freely through a uniform sand column (Section 5.4.2).

5.3.2 Drainage Decay and the Volumetric Water Balance

The drainage decay coefficient β_3 controls the rate at which the water table relaxes toward its equilibrium position, and its spatial distribution (Section 4.9.2, Figure 40d) shows the strongest contrast of any SSM parameter across the site. C1 Lake Edge drains fastest, driven by rapid lateral exchange with Llyn Rhos-Ddu. C2 and C3 show similar intermediate rates. C5 Coastal Forest is lower, and C4 Main Forest has the weakest drainage in the network (Table 3). The displacement formulation (Section 3.4.1) is essential to resolving the C4 drainage term: under a ground-surface reference C4's β_3 was borderline significant; referencing drainage to the 3.7 m datum resolves it to a clearly significant value (Table 3). The slow drainage at C4 is physically consistent with the canopy suppressing recharge pulses that would otherwise drive the water table above its equilibrium and generate a strong drainage response, combined with the impeded lateral discharge through C4's heterogeneous substrate (Section 5.3.1; Section 5.6.1). The contrast between fast-draining eastern clusters and slow-draining forest clusters shapes the volumetric water balance partition between drainage and atmospheric draw.

The indicative volumetric water balance (Table 4b) provides a complementary perspective. All clusters receive approximately 892 mm/yr of rainfall; for the forested clusters, canopy interception (24%; Freeman, 2008) reduces the throughfall reaching the ground surface to approximately 678 mm/yr. The ET/drainage partition is bracketed by two independent methods — the SSM β_2/β_3 ratio and a seasonal recession curve analysis — with the mid-point reported in Table 4b. The cluster partitions differ substantially: at C1, drainage dominates total losses, reflecting rapid lateral exchange with the lake; at C4, atmospheric draw dominates losses and drainage is weak, indicating that the Forest cluster's water balance is governed primarily by the competition between rainfall input and evaporative demand rather than by hydraulic drainage. C2 and C3 sit intermediate, with drainage-dominated partitions; C5 is closer to balanced.

The water table fluctuation (WTF) analysis (Section 4.2.4; Table 4c; Figure 10) addresses the question of whether the substrate's storage capacity differs across the cluster structure. Uncorrected forested cluster Sy medians exceed the open dune range, but this reflects the WTF method attributing the full rainfall total to recharge when in reality a fraction is intercepted by the canopy. Once corrected for canopy interception, the forested cluster Sy estimates fall within the open dune range, consistent with the uniform aeolian dune sand substrate and the well-mixed depositional history of the dune field. With the substrate's storage capacity broadly uniform across the site, the forest clusters' distinctive β-coefficient signature — lower β_1, higher LCSC — derives largely from canopy boundary conditions operating at the surface rather than from differences in subsurface storage architecture, though depth-coupling and lateral drainage geometry also contribute to within-forest β variation (Section 5.3.1).

The implication for management interpretation is that the surface-mediated canopy effects identified by the analysis are in principle responsive to silvicultural intervention, while the subsurface storage and transmission properties of the aquifer are fixed by geological history. The net hydrological response to silvicultural intervention at the cluster level is small relative to climate forcing over the same horizon (Sections 5.5.2, 5.7.1), and the spatial reach of forest management does not extend to the open dune clusters (Section 5.8); the surface-vs-subsurface distinction therefore concerns *which* of the canopy's effects are accessible to management rather than *whether* canopy management is the dominant restoration lever.

5.4 Dune Scraping

Three dipwells — CEH36, CEH18 and CEH21 — were subject to ground surface lowering by mechanical scraping, with CEH36 scraped in April 2015 and all three re-scraped in October 2023 (Section 4.5). CEH36 is located approximately 340 m seaward and downslope of the December 2017 clearfell footprint, a spatial relationship that becomes important when the interaction between the two interventions is considered below.

5.4.1 Measured Benefit and Durability

At CEH36, three independent methods (raw BACI, synthetic control, SSM forward residual) converge on a scraping benefit during the pure scraping era from April 2015 to December 2017 (Section 4.5; Figures 16, 17). The seasonal record reinforces the finding: the era-mean summer minimum shallowed substantially during the scraping era and held its position through the post-felling period, even as the CEH4 control deteriorated under the background drying trend (Section 4.5; Figure 19). Calibrated against the Curreli et al. (2013) eco-hydrological gradient (Section 5.2), the summer improvement at CEH36 represents a shift across more than half the gradient separating the SD15b and SD16 community thresholds.

Scraping creates a topographic low that acts as a local drain, raising the water table at the scraped site by collecting lateral inflow. The propagation analysis (Section 4.5) found no detectable drawdown at wells uphill of the scrape within the observation window, indicating that the hydraulic influence either does not extend beyond the immediate vicinity of the scraped area at the spatial scales resolvable by the monitoring network, or operates on timescales that the available record cannot resolve (Section 5.4.3). The benefit at the scraped site does not alter the rate at which the water table responds to climate forcing, so the background drying trend continues to erode the initial improvement; at CEH36, the improvement is diminishing but the gain has been retained through the post-felling period to date. CEH18 and CEH21, scraped only in October 2023, have shown limited and inconsistent responses that do not survive Tier 1 correction for background drift — consistent with their position closer to the coast where progressive erosion outpaces the localised topographic benefit (Section 4.5).

5.4.2 Interaction with the Clearfell

The interaction between scraping and the subsequent clearfell is instructive. CEH36 sits approximately 340 m downslope of the felling footprint, connected by the continuous downslope drainage pathway identified in the Darcy flow analysis (Section 4.9.5, Figure 44). The CEH36 BACI step during the felling pulse era was smaller than during the pure scraping era — the scraping benefit diminished but was not reversed (Section 4.5; Figure 22). The CEH4 control deteriorated across the same period under the background trajectory, and a substantial part of the apparent reduction in BACI step at CEH36 reflects CEH4's own deterioration rather than a loss at the scraped well. Any recharge benefit from the clearfell propagating downslope to CEH36 would operate on the multi-year to multi-decadal timescale developed in Section 5.4.3 and may not yet be detectable in the monitoring record.

The reverse interaction also warrants consideration: the April 2015 scraping at CEH36 created a topographic drain that may have drawn water away from the wells used to measure the clearfell response. WMC3 (the sole Impact well) and CEH31 (the nearest Edge well) sit closer to CEH36 than the other Edge wells, and the ANCOVA's distance-weighted scraping covariate (Table 7; Section 4.6.2) absorbed the scraping signal at these positions. That WMC3 returned a smaller clearfell step than the Edge tier average, despite sitting closer to the felling compartment, is consistent with the scraping drain partially suppressing the apparent clearfell response at the Impact well. The two interventions propagate through the same continuous aquifer on overlapping timescales, and the interaction illustrates the difficulty of isolating individual management signals in a system where perturbations are entangled.

Scraping at CEH36 produced a benefit in summer minimum depth that persisted through the post-felling period. The benefit diminished but was not reversed; the dominant threats to scraping durability identified by the monitoring record are the background climate drying trend and coastal erosion at the western boundary, rather than upstream canopy management.

5.4.3 Coastal Erosion Analogy and Cascade Risk

The mechanism by which scraping produces its measured benefit — lowering the ground surface to reduce the depth to water table — is physically identical to the mechanism of coastal erosion. Both reduce ground elevation relative to a fixed absolute water table position; both increase the displacement head (h_disp) at adjacent wells; and both consequently enhance drainage gradients away from the intervention or erosion zone. The Tier 1 CUSUM trajectories at CEH4 and CEH22 (Section 4.5; Figure 16) show this process operating naturally at the western coastline, with both control wells showing accelerating negative CUSUMs attributable to progressive coastal erosion of the dune edge.

This analogy carries a management implication that the SSM physics make explicit but the available monitoring record cannot test. Scraping along the western coastline, while locally beneficial in reducing the depth to water table at the scraped location, increases the hydraulic gradient between the coast and the dune interior. The SSM drainage term operates in both directions: a scraped coastal zone with elevated h_disp draws water seaward, deepening the water table at wells uphill of the scrape. The magnitude of this effect is governed by the same aquifer parameters that control the natural coastal erosion signal — and the results at CEH4 and CEH22 show that signal to be substantial and progressive.

The risk of a cascade follows from the mechanism. If coastal scraping deepens summer minima at interior wells, those wells may in turn be identified as candidates for scraping — lowering their ground surface, increasing their h_disp, and propagating the drainage enhancement further inland. Each successive scrape addresses the local symptom while intensifying the systemic cause. This chase-inland dynamic is the topographic analogue of the drainage acceleration that the clearfell analysis identified as a fitting artefact in Section 4.6.5: the aquifer is not gaining storage decay capacity, but the systematic lowering of ground surfaces produces a head trajectory that the model — and the manager — interprets as progressive deepening requiring further intervention. The BACI-corrected $\Delta\beta_3$ values at uphill wells (Section 4.5; Script 09b) showed no distance-decay from the scraping site within the observation window; this is the predicted result if the propagation timescale exceeds the available record (see below), not evidence that the cascade mechanism is inactive.

Most of the Western Residual already lies below the hydrological range for functioning dune slack habitat (Section 4.8; Figure 39), so the practical scope for this cascade is limited to the transitional margins. The principle that coastal scraping mimics coastal erosion in its effect on the inland water balance is well-founded in the SSM physics regardless of the available data's ability to test it within the observation window, and on the precautionary principle should inform the spatial planning of future scraping programmes. The most productive scraping targets identified by the analysis (Section 5.7) are the C1/C2/C3 transitional wells in the mid-eastern sector, where structural sensitivity maximises the retained benefit and where the scraping does not amplify an existing coastal drainage gradient.

The timescale over which these processes operate is governed by the aquifer's hydraulic diffusivity. Using representative values for clean dune sand and the interception-corrected WTF Sy ranges from Section 4.2.4, the characteristic propagation timescale at 340 m is years to a decade, and at 570 m is one to several decades; the literature values for unconfined dune sand (Fetter, 2001) shorten these ranges but do not alter the order of magnitude. A direct measurement of the inland reach of the coastal-erosion gradient from the 21-year monitoring record (Section 4.8.1; Figure 38b) gives a fitted strip-aquifer width of 869 ± 39 m under a Dupuit–Forchheimer steady-state interpretation, or an equivalent exponential decay length of 397 ± 63 m under a transient diffusive interpretation. Both fall within the diffusivity-derived range above, providing independent network-scale corroboration of the years-to-decades timescale regime. Whichever estimate is adopted, the propagation timescale is measured in years to decades. This has three consequences. First, it explains why the BACI-corrected $\Delta\beta_3$ at uphill wells showed no distance-decay: the propagation timescale exceeds the 31-month observation window. Second, the coastal erosion signal visible in the CEH4 and CEH22 CUSUMs is a lagged response to erosion that occurred years earlier; if the rate of coastal retreat is itself accelerating, consistent with regional projections under sea level rise (Pye and Saye, 2005) and the CEH22 trajectory from 2021, then the full hydrological consequences of recent erosion have not yet propagated to the interior wells. The current CUSUM trajectories are a leading indicator rather than a real-time measure of present inland conditions. Third, it provides a physical explanation for the clearfell's slow equilibration, consistent with the propagation timescale combined with C4's weak drainage. The aquifer's response to any boundary perturbation — whether erosion, scraping or land cover change — is measured in years to decades, and monitoring programmes need to be designed accordingly.

5.5 The Clearfell Experiment

The December 2017 clearfell of approximately 8.4 ha of Corsican pine plantation provides the most direct test of how canopy removal affects the water table in the dune slack system. The clearfell zone sits upslope of CEH36 (the scraped slack described in Section 5.4) and straddles the gradational boundary between C4 Main Forest and C5 Coastal Forest. The sole Impact well, WMC3, shows strongest Pearson affinity with C3 Western Residual (r = 0.965) with C5 a close second, and the extended-network FE wells within the felled footprint show similarly mixed affinities across C3, C4 and C5 (Section 4.3). The felled compartment therefore cannot be treated as a clean C4-only experiment; the clearfell zone spans the transition from forest to open dune, and its hydrological effects propagate downslope into the C5 catchment via the continuous drainage pathway identified in the Darcy flow analysis (Section 4.9.5, Figure 44).

The three-counterfactual ANCOVA-BACI (Section 4.6.2) returned a significant clearfell step against the Forest control — the most direct test, comparing felled wells against unfelled forest sharing the same canopy and substrate — of +93 mm at the Impact tier and +153 mm at the Edge tier. The Edge tier returned a larger step despite sitting further from the felling compartment, a configuration explained by the Edge tier's four-well averaging reducing sampling noise relative to the single Impact well. The Climate control comparison is larger again and captures the standing forest's pre-existing hydrological footprint alongside the December 2017 clearfell; the Forest control is the appropriate test for canopy-removal-specific attribution.

Four independent robustness methods support the ANCOVA result: SSM residual normalisation, synthetic control via donor weighting, synthetic extension of FE1 and FE2 within the felled compartment, and a clearfell transect regression (Section 4.6.7). All four return clearfell steps in the same direction as the ANCOVA and at comparable magnitudes. The robustness estimates sit consistently below the ANCOVA because they do not include the easting × time correction for the coastal erosion gradient, and so represent lower bounds on the clearfell-specific component. The transect regression returned no significant spatial gradient, although the post-felling spread in transect hydrographs is visibly wider than pre-felling, with WMC3 consistently above the pack — suggesting a distance-dependent response that the transect regression lacks the power to resolve. The summer minima analysis (Section 4.6.4) returned no significant clearfell effect at the Impact tier and a small negative effect at the Edge tier, opposite in sign to the monthly ANCOVA result — summer minima at the felled Edge wells deepened slightly relative to the Forest control after felling. The clearfell therefore registers in the mean monthly water level but not in the summer extremes, where felled and unfelled wells track each other. The seasonal partition is consistent with the canopy's recharge-suppression and atmospheric-shielding effects operating symmetrically: their removal raises the winter recharge contribution and the summer atmospheric demand by comparable amounts, lifting the mean while leaving the seasonal extremes approximately unchanged. The β_1 and β_2 contributions to this partition are quantified in Section 5.5.1, where the contribution of the site-wide β_1 decline operating disproportionately on summer lows is also considered.

Independent validation of the ANCOVA's easting × time correction against the network-scale coastal-retreat gradient model (Section 4.6.3, Table 8) shows the Forest Impact comparison sits within sampling error of the model prediction ($z = -0.35$), confirming that the +93 mm Impact-tier clearfell step is not an artefact of incomplete gradient correction. At the Edge tier the covariate absorbs more than coastal retreat alone justifies ($z = -5.3$), most plausibly because it is also absorbing residual canopy-edge transpiration enhancement or differential cluster-membership drift. The implication is that the +153 mm Edge clearfell step is a conservative lower bound on the felling-attributable signal at that tier rather than an inflated estimate: if the easting × time term is absorbing real coastal retreat plus additional spatial drift, the clearfell-specific component is at least the headline value. The Climate-tier comparisons show a similar over-correction pattern, consistent with the climate controls themselves carrying gradient-zone signal that the covariate absorbs into the spatial term.

The coefficient decomposition (Section 4.6.6; Table 11; Figure 32) provides additional context. Recharge sensitivity (β_1) declined across all five tiers of the BACI network — a site-wide phenomenon visible in controls and treatment wells alike, ruling out a clearfell-specific origin (mechanism discussed in Section 5.5.1). Atmospheric draw (β_2) increased at Forest Control wells, consistent with progressive canopy maturation, but showed mixed responses at Climate Controls. Storage decay (β_3) changes were small and spatially incoherent. The predicted clearfell effect from coefficient shifts alone was negative at the felled tiers, but because controls show comparable shifts, the net coefficient-attributable component is near zero. This sits alongside the ANCOVA's significant +93/+153 mm step. The two analyses ask different questions. The ANCOVA measures the gap between felled and control wells, so the site-wide β_1 decline — which affects both — cancels and the canopy-removal effect is visible. The per-well decomposition measures coefficient changes at each well individually; the β_1 decline registers as a negative contribution at every well, and because controls and felled wells show comparable shifts, no clearfell-specific signal emerges. The two are reconciled rather than contradictory.

5.5.1 Seasonal Asymmetry, the β_1 Decline, and the Confounded Baseline

The site-wide β_1 decline is a significant finding independent of the clearfell result. It represents the strongest candidate mechanism for the observed summer minimum decline documented in Section 4.8.1 and is independent of any management intervention. The decline is most parsimoniously attributed to changing within-month rainfall distribution. If the same monthly total arrives in fewer, more intense events separated by longer dry spells, water held in the unsaturated zone between events is increasingly drawn back out by ET before it can percolate to the water table. The net recharge per millimetre of rainfall therefore falls even where the gross monthly total is unchanged. This interpretation is consistent with the summer-dominated drying trajectory visible across all clusters, with winter maxima showing no significant trend.

Against this background, the clearfell produces an inherent seasonal asymmetry. Removing the canopy recovers the interception loss, raising β_1 and increasing winter recharge. The same removal increases β_2 — the water table's sensitivity to atmospheric demand — so a given monthly PET produces a larger summer drawdown. The two effects operate on different seasons and through different pathways; their net annual balance is the quantity of management interest, and is addressed in Section 5.5.2.

The CWB × clearfell interaction — testing whether the water table's sensitivity to cumulative climate forcing changed after felling — was non-significant in all six ANCOVA models (Section 4.6.2). The climate sensitivity of the felled wells did not detectably change after canopy removal: the ANCOVA detects a significant upward shift in mean water table level, while the rate at which the BACI gap responds to cumulative water balance deficit remains unchanged, with pre-felling and post-felling slopes approximately parallel (Figure 26). This is consistent with the canopy acting simultaneously as a hydrological sink (transpiration, interception) and a microclimatic buffer (shading, reduced wind exposure, organic forest floor). Canopy removal modified both functions: the recharge gain raised the mean water table, and the concurrent change in atmospheric draw operated on the same monthly water balance, so that the system's net sensitivity to cumulative climate forcing was preserved. The practical consequence is that the clearfell effect on the BACI gap is a fixed offset — it shifts the baseline from which the water table responds to climate — rather than a change in the rate of response to wet or dry years.

The summer minima record makes this seasonal partition explicit. At WMC3 — the sole Impact well with sufficient pre-felling baseline — summer minimum depth was −1.556 m (mean, n = 6 years) before the April 2015 scraping, −1.620 m during the scraping era (2015–17, n = 3), and −1.591 m post-felling (2018–25, n = 8). The Forest Control group deepened over the same period, from a tier-mean summer minimum of −1.305 m pre-scraping to −1.390 m post-felling, and the gap between Impact and Forest Control remained approximately constant (Figure 30). Both Impact and Edge tiers track the Forest Control trajectory in summer extremes, with the felled wells' summer minima moving in parallel with unfelled forest while their mean monthly levels sit above it. The climate-corrected spatial step-change mapping (Section 4.6.5; Figure 31) shows that the felled area resisted the site-wide drying trend by approximately 28 mm relative to the open dune.

The ANCOVA, The ANCOVA additionally includes an easting × time interaction term, which absorbs the progressive west-to-east deepening signal attributable to coastal erosion at the western boundary (Section 4.6.5). The felled compartment sits in the western half of the network, so without this correction the coastal erosion signal is partially attributed to the felled wells and offsets the clearfell recovery. With the easting × time term included, the ANCOVA recovers a clearfell-specific step of +93 mm at Impact and +153 mm at Edge — substantially larger than the uncorrected spatial map's +28 mm. The two estimates are not in disagreement: the spatial map and the ANCOVA differ in what they correct for, and the ANCOVA's additional correction is the one that isolates the canopy-removal signal from the coastal erosion signal. The +28 mm figure should be read as a lower bound that retains the erosion confound; the ANCOVA figures are the clearfell-attributable component.The clearfell produced a significant relative recovery against the Forest control in the mean monthly signal; the response in summer extremes is set by the site-wide β_1 decline and the background climate trajectory, which operate symmetrically across felled and unfelled wells.

The 1990s-era attribution of water table depression to the forest (Cottingham, 1994; Bristow and Bailey, 2001; Betson et al., 2002) was conducted during the driest decade in the 95-year RAF Valley record (mean annual rainfall 793 mm against a long-run average of 853 mm; Section 4.1.1) and without the kind of climate-partitioned analysis applied in this study. The apparent depression of the 1989–1995 water table baseline is consistent with a compound signal of forest influence and exceptional climatic dryness that the data available at the time could not separate. The present three-counterfactual design separates these contributions and additionally partitions the canopy effect by season. The forest does suppress recharge — β_1 is lower under canopy, and the BACI record confirms a significant mean water table recovery following clearfell. The earlier studies would have recorded this winter-recharge signal as a uniform water table depression; the present analysis shows that the canopy's effect on summer extremes is small relative to the site-wide β_1 decline operating across all tiers. The 1989–1995 baseline was therefore not solely a forest signal, nor solely a climate signal, but a compound of both — and a compound whose seasonal structure was not visible at the time.

The monitoring evidence therefore characterises canopy removal as a winter-recharge intervention. Its effect on summer water table depth — the metric that determines dune slack viability — is set by the same site-wide β_1 decline that affects unfelled forest and open dune wells alike.

5.5.2 Forest Management Scenario Responses

The scenario framework (Section 4.10.2) quantifies the seasonal trade-off at the cluster level under the BACI-corrected upper-bound β_2 multiplier (×1.108; Section 3.8). At C4, the clearfell scenario gives a net annual gain of +4.7 mm w.e./month, with the winter recharge gain exceeding the summer β_2 cost on the thin, high-β_2 substrate above the ridge. At C5, where the thicker substrate and lower β_2 attenuate the summer cost more strongly, the net annual gain is larger at +8.8 mm w.e./month. Fifty percent thinning produces a similar pattern at reduced magnitude (C4: +2.3, C5: +4.4 mm w.e./month). Both forest clusters respond in the same direction under clearfell and thinning, with C5 capturing approximately twice the recovery per unit area as C4. The cluster-level β_2 contrast (Section 5.3.1) — substrate as a major contributor, with depth-coupling and other factors also acting — translates here into per-area scenario response: the same canopy treatment produces different hydrological responses across the forest zone, and management prescriptions cannot assume uniform per-area benefit across the forest even where the underlying canopy is the same.

Broadleaf conversion produces a different seasonal balance. When the summer transpiration increase under deciduous canopy in full leaf is accounted for (β_2 summer multiplier ×1.113, reflecting comparable growing-season transpiration rates between deciduous and Corsican pine canopies; Komatsu et al., 2011), the interception reduction from 24% to 15% during the recharge season is offset by the summer evapotranspiration penalty. The net annual volumetric effect is −3.4 mm w.e./month at C4 and +0.1 mm at C5 — near-neutral at C5 and slightly negative at C4 where the thin substrate amplifies the β_2 cost. The phenological advantage of winter leaf-off (zero interception during the recharge season) is offset by the growing-season transpiration demand, and the net annual effect is substantially smaller than climate-driven changes over the same horizon.

The three scenarios therefore differ in seasonal partition rather than in any one being categorically more or less effective than the others. Clearfell and thinning produce small positive net responses concentrated in the winter recharge season; broadleaf conversion produces a near-neutral annual balance that reflects the seasonal trade between winter recharge gain and summer transpiration demand. All three magnitudes are small relative to the climate-driven Δh range projected over the same horizon (Section 5.7.1), and the cluster-level partition between C4 and C5 — substrate-related but not substrate-only — means that the same canopy treatment delivers a smaller per-area response at C4 than at C5.

5.6 Hydrological Effects of the Plantation

The Corsican pine plantation intercepts approximately 24% of incident rainfall before it reaches the aquifer (Freeman, 2008). This interception is a partition of the available PET energy budget, not an additional loss on top of PET (Section 3.7.2): intercepted water evaporates from wet canopy surfaces without stomatal control, consuming a fraction of the atmospheric energy that would otherwise drive transpiration and direct soil evaporation. The β_2 coefficient for the forested clusters therefore reflects the combined effect of this unmediated interception evaporation plus the residual stomatal transpiration and direct evaporation beneath the canopy. The interception fraction is well constrained by Freeman's (2008) site-specific measurement and is independently cross-validated by the β_2 ratio between C3 and C5, which yields an estimate of 22.6% (Section 5.3). The broader conifer plantation literature reports interception losses of 25–45% (Rutter et al., 1971; Gash, 1979; Gash et al., 1980), with the higher end deriving from dense Sitka spruce stands in UK upland catchments (Calder and Newson, 1979; Robinson and Dupeyrat, 2005) — a different canopy structure from the partially thinned Corsican pine at Newborough. Numerical modelling of comparable coastal pine settings in northern Poland produced interception losses of approximately 27% (Gumuła-Kawęcka et al., 2021, 2024), consistent with Freeman's figure.

The canopy has further hydrological effects beyond recharge suppression. It provides microclimatic shielding — reducing direct solar radiation, insulating the sand surface through the organic forest floor, and reducing wind-driven advective evaporation — and the forest floor and root-zone macroporosity produce a smoothed vertical recharge signal relative to the rapid, pulsed infiltration of the open dune (Calder, 2007). The lag diagnostic

(Script 03) provides independent support for this smoothing effect. At lag-0, C4 Main Forest has the lowest R^2 of any cluster (0.682), confirming that the contemporaneous rainfall-to-water-table transfer is weakest beneath the canopy. At lag-1, the pattern inverts: C4's lag-1 R^2 (0.377) is the highest in the network, with a significant and physically coherent β_1 of 1.607 ($p < 10^{-12}$). No other cluster retains this level of explanatory power at one-month lag — C1 Lake Edge drops to $R^2 = 0.037$ (β_1 non-significant). The canopy delays and smooths the recharge pulse so that a substantial fraction of each month's rainfall reaches the forest water table during the following month rather than contemporaneously. The lag-0 formulation is retained for all clusters because it maximises overall network R^2 and preserves the operational simplicity of a single-lag framework. The strong lag-1 signal at C4 indicates that the forest's lower apparent β_1 (2.518 vs 3.576–4.581 at the open dune clusters) partly reflects temporal redistribution of recharge rather than a reduction in total recharge reaching the water table.

These effects, and their interactions with the aquifer substrate, are examined in the subsections below.

5.6.1 Atmospheric Draw, Specific Yield and the β_2 Coefficient

The compound nature of β_2 and the dominance of specific yield in the C4/C5 contrast were established in Section 5.3.1. The per-well analysis within the forest zone confirms and extends this finding.

C4 Main Forest has the highest β_2 in the network. This does not mean the forest water table experiences the greatest atmospheric demand — it means that a given evaporative flux produces the largest head change at C4, because C4 has the lowest effective specific yield. C4 overlies irregular bedrock topography on the elevated ridge flank, where buried ridges impede lateral drainage and create a heterogeneous substrate of thin sand pockets and deeper troughs, reducing the effective drainable volume per unit head change. The WTF-derived Sy estimates (Table 4c) independently confirm the direction: C5's Sy is approximately a third higher than C4's, consistent with C5 occupying deeper, more uniform coastal sand — essentially C3-type substrate with forest canopy, whereas C4 is more analogous to C2-type substrate with forest canopy. The drainage timescale analysis (Section 4.9.3, Figure 41) reinforces this picture: C4 wells show the longest τ values in the network, indicating that lateral discharge is impeded by subsurface barriers rather than flowing freely through a uniform sand column.

C5 Coastal Forest has the second-lowest β_2 in the network despite carrying the same Corsican pine canopy as C4. This reflects C5's higher Sy — the thicker coastal sand provides more drainable storage per unit depth, so the same atmospheric demand produces a smaller head change. The per-well β_2 range within C4 alone spans nearly the full site-wide range (Section 4.9.4; Table 17), confirming that within-cluster substrate variation is at least as important as between-cluster canopy differences. Cluster-level centroid values summarise a spatially variable field; they should not be taken as implying uniform conditions across the forest.

Across the forest wells, elevation alone explains the bulk of the variance in β_2 (Section 4.9.4), with higher-elevation wells on the ridge showing systematically higher atmospheric draw sensitivity. The C4/C5 boundary corresponds to a clean elevation break between the elevated ridge and plateau terrain and the low-lying coastal fringe, with no overlap. Both zones carry the same Corsican pine canopy; the difference is substrate and topographic position. The implication for management interpretation is that β_2 values at any given forest well are dominated by position, not canopy: changes in β_2 following management must therefore be measured against a position-matched control if they are to be canopy-attributable. This is precisely the role of the Forest control tier in the clearfell ANCOVA (Section 4.6.2).

Although substrate dominates the spatial pattern of β_2, the post-clearfell coefficient analysis (Section 4.6.6; Table 11) shows that canopy shielding exerts a measurable effect on top of this baseline. β_2 increased at Forest Control wells and at the Edge tier following felling, confirming that the canopy was moderating the atmospheric demand signal reaching the water table. The increase was smaller at the Impact well and at the Coastal Control, and mixed at the Climate Controls. The shielding mechanism operates at the surface: intercepted water evaporates readily from wet canopy surfaces without stomatal control, consuming PET energy before it can act on the water table. Removing the canopy eliminates this interception pathway and exposes the water table directly to atmospheric forcing — a greater proportion of PET demand is then transmitted to the shallow aquifer through direct soil evaporation and grassland transpiration on the grazed clearfell surface. This shielding effect is superimposed on a substrate-driven β_2 baseline that varies by a factor of two across the forest, and the cross-cluster β_2 pattern reflects Sy differences rather than canopy differences.

The depth-dependent PET diagnostic (Section 4.4; Section 3.7.1) is consistent with this interpretation. C1 Lake Edge shows the strongest depth coupling, consistent with shallow water-table evaporation near the lake being the dominant ET pathway, while C4 Main Forest shows the weakest depth coupling despite having the highest β_2. If C4's high β_2 were driven by deep-root transpiration overriding capillary disconnection, C4 would show the strongest depth coupling, not the weakest. Together with the sunshine-hours residual diagnostic (Supplementary Note S6; Section 5.2), which returned a null result, this confirms that the standard SSM's fixed β_2 has already absorbed the bulk of the ET signal at the cluster level.

The practical consequence connects the plantation directly to the ecological outcome of most concern. Summer minimum water table depth is the dominant predictor of winter flooding potential in all western and forest blocks (Section 5.2; Table 13). Under a warming climate with increasing summer PET, any intervention that raises β_2 — including canopy removal — would amplify the head response to that PET, deepening summer minima per unit of evaporative demand and raising the rainfall threshold required for winter flooding. The post-clearfell β_2 increases at Forest Control and Edge tiers confirm this mechanism empirically. The seasonal trade-off at the forest clusters — winter recharge gain from interception recovery against summer β_2 increase from loss of shielding — is quantified by the BACI-corrected scenario projections (Section 4.10.2; Section 5.5.2): net annual effects are small and positive at both C4 and C5, and substantially smaller than the climate-driven Δh range over the projection horizon.

5.6.2 Summer Minimum Depths at the Forest Margin

The relatively shallow summer minima at NW2 and CEH2 sit within the zone where the water balance residual field shows the largest positive values in the network (Section 4.9.6, Figure 45). The mechanistic attribution of this residual is unresolved (Section 5.2.1) — the spatial pattern is real, but the candidate mechanisms (ridge-derived lateral input, threshold recharge behaviour) cannot be distinguished by the available diagnostic tests. The favourable position at NW2 and CEH2 reflects some combination of this locational subsidy and the canopy's moderation of atmospheric demand operating at those specific wells. Clearance of the plantation in this zone would eliminate the canopy contribution to the local water balance — the coefficient decomposition (Section 4.6.6) shows that β_2 increased at Forest Control wells following felling (mean $\Delta\beta_2$ +0.259), indicating that the canopy was moderating the atmospheric demand signal reaching the water table. Removing it at NW2/CEH2 would expose the water table directly to atmospheric forcing while leaving the locational subsidy (whatever its mechanism) operating against a higher β_2. The BACI analysis (Section 5.5) confirms that canopy removal in the adjacent clearfell zone did not raise water tables — the same outcome would be expected at NW2/CEH2 unless the locational subsidy at those wells dominates the canopy contribution by a sufficient margin to reverse the BACI result, and no evidence in the available data supports that assumption.

The summer minimum depth map identifies where water table levels are currently most actively sustained against the background climate drying trend. The favourable position at the forest-adjacent margin reflects both location and canopy condition; the available evidence does not support clearance as a route to preserving or improving it.

5.6.3 Broadleaf Conversion: The NW10 Natural Experiment

The monitoring network contains a direct observational test of broadleaf conversion effects. A block of mixed broadleaf woodland was established following clearfell in 1993 and restocked around 1996 on the north-western margin of the plantation, with dipwell NW10 situated within the block and monitored continuously from 2007. CEH2 and CEH34, both classified as C4 Forest under intact Corsican pine canopy, serve as the principal comparators; NW11 and NW13 provide a paired pine-warren boundary transect approximately 90 m apart at virtually identical elevation.

The per-well SSM coefficients at NW10 are informative. Its recharge sensitivity $\beta_1 = 3.526$ substantially exceeds the C4 per-well median of 2.338 and the adjacent Corsican pine well CEH2 ($\beta_1 = 2.732$), consistent with the broadleaf canopy intercepting less rainfall than dense pine and allowing more recharge to reach the water table per millimetre of rainfall. The seasonal interception pattern expected from a deciduous canopy supports this reading: during the October–March recharge season the leafless broadleaf stand intercepts less rainfall than the year-round Corsican pine canopy. The meta-analytical review of Komatsu et al. (2011) reports that broadleaf deciduous forests in high-winter-precipitation climates consistently transmit more rainfall to the sub-canopy than coniferous forests at the same sites, with the effect most pronounced during the dormant season.

The atmospheric draw coefficient $\beta_2 = 2.524$ at NW10 is lower than CEH2 ($\beta_2 = 2.891$) but close to the C4 per-well median (2.618). The broadleaf stand does not show an increase in atmospheric draw above the C4 range. This is consistent with the seasonal offset of deciduous leaf-off during October–March partially compensating for higher growing-season ET demand, such that the year-round integrated atmospheric draw at the water table is comparable to that under adjacent pine. The comparison is complicated by positional differences: NW10 carries a higher β_3 (0.046) than the C4 median (0.019), so drains faster than most C4 wells, and its ground elevation and distance from the ridge differ from CEH2 and CEH34. These positional differences mean that the β_2 comparison isolates species effects imperfectly.

NW10's behaviour relative to its pine comparators differs depending on which comparators are used and whether the comparison is climate-normalised. When normalised to the shared climate signal, NW10's summer minimum sits consistently shallower than the pine interior composite (CEH2, CEH32, CEH33, CEH34), with a mean anomaly of +247 mm across the full record and +300 mm during the post-felling period (2018–2025). The anomaly is diminishing at −46 mm/yr (Section 4.6.8), consistent with the broadleaf stand progressively deepening summer minima as canopy closure increases interception and transpiration demand — the same direction of trajectory observed for the standing pine canopy as it matures. The direct comparison against CEH2 (NW10's nearest pine neighbour, 100 m downslope) shows NW10 deeper than CEH2 across the monitoring record; however, CEH2 itself shows a progressive deepening trend relative to the network mean that is not present in other pine interior wells (Section 4.6.8), which complicates the local-pair comparison. The reconciliation of these two pictures is the principal limitation of the NW10 experiment, and is addressed below.

The scenario framework (Section 4.10.2; Section 5.5.2) provides an independent quantitative reference. When the summer transpiration increase under mature deciduous canopy is accounted for (β_2 summer multiplier ×1.113), broadleaf conversion produces a near-neutral net annual effect at C4 (−3.4 mm w.e./month) and at C5 (+0.1 mm). The convergence of two pieces of evidence — the diminishing climate-normalised anomaly at NW10 and the near-neutral scenario prediction at the forest clusters — points in the same direction, towards a broadleaf canopy whose annual hydrological effect at maturity is similar in net direction to the pine it replaced, even though the seasonal partition differs. The phenological winter recharge advantage at NW10 is real, but does not translate into a sustained shallower summer minimum on the available evidence.

NW10 is a single dipwell, and the positional confounders are of comparable magnitude to the observed contrasts. The spread between CEH2 (mean summer minimum −1.634 m) and CEH34 (−1.309 m), both Corsican pine wells within the same cluster, is 325 mm, exceeding the NW10 climate-normalised anomaly of +247 mm. No reliable attribution to canopy type is possible from a single dipwell, and a minimum of three to four instrumented wells distributed across the broadleaf compartment would be required before firm conclusions could be drawn. The 6–7 year extrapolation from the current −46 mm/yr trend should be treated similarly as an indicative trajectory rather than a quantitative projection, since the trend itself rests on a single well's post-felling record.

A related observation suggests the broadleaf block's hydrological footprint may extend beyond its physical boundary. CEH2, situated approximately 100 m downslope within the adjacent Corsican pine plantation, shows the progressive deepening trend relative to the network mean noted above (Section 4.6.8). This is consistent with the maturing broadleaf canopy progressively increasing winter recharge within the block, steepening the local hydraulic gradient and drawing down the water table in the adjacent downslope zone — the same mechanism that operates at the scraped slacks (Section 5.4). The observation is based on a single well comparison and cannot be attributed to canopy type with confidence, but it warrants targeted instrumentation downslope of the broadleaf block.

5.6.4 Stand Density, Thinning and the Dothistroma Natural Experiment

The preceding sections characterise the seasonal partitions produced by complete canopy removal and broadleaf conversion: clearfell produces a small positive net annual effect concentrated in the mean monthly water level, with the winter recharge gain offset by the summer atmospheric draw increase; broadleaf conversion produces a near-neutral annual balance where the phenological winter advantage is offset by growing-season transpiration. A third option — partial canopy retention through stand density management — has not been tested in the monitoring record but is mechanistically distinct from both alternatives.

No thinning experiment has been conducted within the study period and no thinned-stand wells exist in the network. The theoretical argument that an intermediate canopy density could retain a measure of microclimatic shielding whilst partially recovering the recharge penalty is mechanistically plausible — the β coefficient framework predicts that partial canopy retention would reduce β_2 less than complete removal while recovering a proportional fraction of the interception loss from β_1 — but this remains a prediction rather than an observed outcome. The scenario framework (Section 4.10.2; Section 5.5.2) estimates that 50% thinning would produce approximately half the annual net response of full clearfell at both forest clusters, with a similar seasonal structure. The thinning and clearfell seasonal profiles are compared in Figure 46 (Section 5.7).

On prograding Atlantic shorelines, stabilised back dunes naturally succeed to deciduous woodland — the recognised climax community of the UK psammosere (Ranwell, 1972). The geomorphological position of the Newborough plantation is therefore not unusual for back-dune terrain; what differs from natural succession at this site is the species choice and the stand density, which reflect commercial forestry management. Freeman's (2008) site-specific measurement of 24% interception sits at the low end of the 25–45% range typically reported for UK conifer plantations, reflecting the exposed coastal setting and the partial thinning that has already occurred. The plantation's effect on the water balance reflects stand density as much as species choice.

Progressive Dothistroma needle blight infection is already driving natural canopy thinning across parts of the Newborough plantation. Monitoring the β_2 response of wells adjacent to blight-affected stands of varying severity would provide a direct field test of the canopy shielding mechanism quantified in this study, and would constitute a natural experiment in progress requiring no additional silvicultural intervention. Whether the trajectory leads toward a hydrologically distinguishable equilibrium cannot be asserted from the present data, but it is a well-defined monitoring question for the 2026–2030 period.

5.6.5 Canopy Salt Interception and Post-Felling Geochemistry

A further geochemical consequence of canopy removal appears to have received limited attention in the dune restoration literature but follows from the LAI-deposition relationship documented in coastal forest hydrology. Pine plantation canopies intercept sea salt aerosols with greater efficiency than open dune vegetation: throughfall studies in coastal *Pinus radiata* plantations have reported chloride deposition rates up to 89% higher beneath the canopy than in adjacent open fields, attributable to the high leaf area index of the needle canopy (Deng et al., 2013), and independent soil studies in the same species have shown systematically higher exchangeable sodium and magnesium pools under pine than under adjacent pasture (Giddens et al., 1997). The Newborough plantation is a different pine species (*Pinus nigra* var. *maritima*), so the specific deposition magnitudes do not transfer directly, but the qualitative mechanism — canopy interception of marine aerosols at LAI characteristic of mature coniferous stands — applies at any site exposed to onshore marine airflow. The Newborough Warren plantation has been exposed to frequent onshore Atlantic airflow throughout its 75-year history.

If this mechanism operates at Newborough at any meaningful magnitude, the forest soils will have accumulated an elevated Cl^- and Na^+ load relative to the open dune system over the plantation's history. Forest removal simultaneously terminates this elevated deposition input and, by increasing aquifer responsiveness following canopy removal, accelerates the downward flushing of accumulated ions through the highly transmissive dune sand. The geochemical consequence is a progressive leaching of base cations from soils that previously received sustained buffering from canopy-concentrated deposition. Curreli et al. (2013) and Sival and Grootjans (1996) identify the maintenance of carbonate-rich groundwater buffering as critical for the persistence of basiphilous slack communities; this flushing dynamic represents a potential geochemical pathway by which large-scale clearance could alter the soil chemistry of base cation buffering, operating on a timescale longer than the water table response and potentially with a longer lag time before its effects become apparent in the vegetation.

A countervailing process should be acknowledged: pine needle litter produces organic acids that acidify the soil surface over time, and removal of the canopy would terminate this acidification pathway. The net geochemical effect of clearfell — salt-leaching versus relief from acidification — is likely to be site-specific and depth-dependent, and warrants targeted investigation through paired throughfall and soil porewater chemistry monitoring across the forest–dune boundary.

5.7 Spatial Groundwater Modelling: Findings and Limitations

A well-level scenario framework (Section 3.8; Section 4.9) was applied to translate the per-well SSM parameterisation into site-wide summaries of equilibrium head change under prescribed climate and forest management perturbations. Climate scenario responses are spatially coherent and physically consistent with the per-cluster β coefficient structure: UKCP18 2050s forcing produces a site-wide mean annual Δh of −0.017 m/month, deepening to −0.030 m/month under UKCP18 2080s. The forested clusters carry larger annual declines than most open-dune wells; for C4 the mechanism is its low effective specific yield, which amplifies the head response to any change in atmospheric demand, while at C5 the larger annual signal reflects reduced net recharge under canopy interception combined with the cluster's already-steep observed trajectory (Section 4.8.1). These outputs are presented as indicative spatial summaries of the SSM equilibrium framework rather than calibrated predictions.

5.7.1 Seasonal Asymmetry under Climate Projections

The annual responses conceal a pronounced seasonal asymmetry that is ecologically critical (Section 4.10.1). Under both UKCP18 trajectories, winter Δh is positive at all clusters — increased winter rainfall under RCP8.5 Wales outpaces the concurrent winter PET increase — while summer Δh is strongly negative and carries the annual signal (site-wide means of

−0.059 m/month at 2050s and −0.111 m/month at 2080s). This summer-dominated signature reflects the same mechanism established in the observed record (Section 4.8.1): hydrological deterioration under climate change operates through intensified summer atmospheric demand rather than declining winter recharge. Annual-mean Δh therefore understates the ecological consequence of summer-dominated climate forcing; the seasonal partition should be reported alongside any annual summary.

5.7.2 Forest Management Scenarios and Structural Isolation

For forest management scenarios, the per-well framework carries an implicit structural feature that is itself informative: because β coefficient changes under canopy removal, thinning, or broadleaf conversion are applied only at C4 and C5 wells, the IDW-interpolated Δh field produces a response concentrated in the forest zone. The drawdown propagation analysis (Section 4.9.3, Figure 42) estimates a characteristic e-folding distance of 245 m, with the perturbation attenuating to negligible levels at all C1 and eastern C2 wells. Site-wide annual means under forest management are correspondingly small (Section 4.10.2; Section 5.5.2), because only 14 of 66 reference wells are in the two forest clusters and the propagation signal does not reach the eastern slacks at measurable magnitude.

C5 is the cluster where clearfell would yield the largest predicted summer minimum benefit (Section 5.5.2), driven by its higher effective specific yield and thicker substrate (Section 5.3.1). C5 is also the cluster with the steepest observed decline in the network — 3.9 times the summer rate of C1, with significant trends in both seasons and an ongoing collapse in seasonal amplitude (Section 4.8.1). Whether the predicted benefit could arrest the observed decline depends on the mechanism driving the decline: if the decline is primarily climatic, the scenario gain offsets several months of trajectory per year; if a structural process such as coastal boundary retreat contributes substantially, the gain offsets only the climatic component, with the structural component continuing.

The spatial coefficient atlas (Section 4.9.2, Figure 40) provides independent characterisation of the forest zone's hydrological position. The β_1 surface shows a sharp step at the forest boundary — recharge sensitivity drops from open-dune values to substantially lower values beneath the canopy — and the β_3 surface shows an order-of-magnitude contrast between the fast-draining lake edge and the near-stagnant forest interior. These are not gradual transitions; they delineate a zone that runs its own local water balance with limited coupling to the open dune clusters. The bedrock ridge forms the northern boundary of the entire study area (Section 2), and all groundwater flows southward off the ridge toward the coast. The Darcy flow vectors (Section 4.9.5, Figure 44) show the forest zone receiving lateral recharge along the ridge margin and discharging south-westward toward Caernarfon Bay, while the open dune discharges south-eastward toward the Menai Strait. A topographic high in the central dune field separates these two drainage paths. A canopy-management perturbation in the C4/C5 zone therefore propagates south-westward toward the coast, not south-eastward into the C1 and C2 clusters. Hydrological connectivity along the coastward gradient is empirically confirmed by the April 2015 scraping response at CEH36 (Section 5.4.2); this connectivity operates between the forest zone and the south-western coastal margin rather than between the forest zone and the eastern slacks.

5.7.3 Framework Limitations and Continuous-Flow Modelling

The conclusion that forest management perturbations do not reach the eastern slacks requires qualification. The forest drawdown propagation analysis (Section 4.9.3, Figure 42) estimates the steady-state head perturbation attributable to the existing Corsican pine canopy: the signal exceeds 50 mm only within ~100 m of the forest edge and is negligible (<1 mm) at all C1 and eastern C2 wells. This quantifies the spatial reach of the canopy's current influence on the water table, and implies that removal of the canopy would produce a recovery of comparable spatial extent — confined to the forest zone and its immediate margins. However, the Darcy flow vectors (Section 4.9.5, Figure 44) and the Pearson affinity analysis (Section 4.3) show that C4 and C5 are hydrologically connected along a continuous downslope drainage pathway. The clearfell's downslope effects on CEH36 (Section 5.4) demonstrate that perturbations in the C4 zone can propagate into C5. Whether they propagate further into C3 and C1/C2 at magnitudes below the propagation model's resolution cannot be resolved by the current framework, and would require a calibrated continuous-flow model.

The practical limits of the spatial framework lie in what it deliberately does not attempt. The per-well IDW framework is a geometric aggregator of SSM responses rather than a physical model of groundwater flow, so it cannot represent cluster-to-cluster hydraulic coupling, lateral redistribution of localized recharge perturbations, or the diffusive response of the aquifer to boundary-flux changes. For the questions this study 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 — the framework is adequate and the BACI record confirms its conclusions. Answering a different class of question, such as how a hypothetical intervention along the northern forest margin would propagate through the aquifer, would require a calibrated continuous-flow model, and for that the critical unknowns are hydraulic conductivity — currently constrained by a single tracer test ($K = 6.0$ m/day; Betson et al., 2002) — and aquifer thickness, inferred from four borehole logs across 700 ha (Bristow, 2003). 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 such a model. These data would also support a HYDRUS-1D vadose zone analysis (Šimůnek et al., 2008) at the cluster level, providing an independent test of the vegetation-driven recharge suppression signal identified in C4 and C5 (Gumuła-Kawęcka et al., 2021).

5.7.4 Forest Scenario Predictions Against the Observed Record

The scenario framework's predicted seasonal responses can be tested against the observed post-felling record at WMC3 (Figure 46). Applying scenario-specific β coefficient adjustments to the observed C4 mean seasonal cycle under the BACI-corrected upper-bound β_2 multiplier (Section 3.8) shows that the three management options diverge in summer response. Full clearfell and 50% thinning both produce shallower summer minima at C4, with clearfell recovering approximately 80 mm and thinning approximately 40 mm at the September trough. The clearfell and thinning gains arise because the interception recovery from canopy removal outweighs the β_2 increase at C4 in the summer equilibrium calculation, consistent with the small positive annual-mean effects reported in Section 5.5.2.

Broadleaf conversion produces the opposite seasonal response: the summer transpiration increase under mature deciduous canopy deepens the September minimum by approximately 120 mm below the pine baseline. The reduced annual-mean interception under broadleaf delivers a visible winter recharge gain — the broadleaf line sits shallower than baseline from November through March — but this additional winter input drains through the highly transmissive sand before summer, and the growing-season transpiration penalty dominates the annual balance (Section 5.5.2). At WMC3, summer minimum depth shows no significant shift between pre-scraping and post-felling, consistent with the small predicted clearfell effect at C4 and the BACI observed trajectory (red shaded area in Figure 46) tracking close to the pine baseline throughout the summer window.

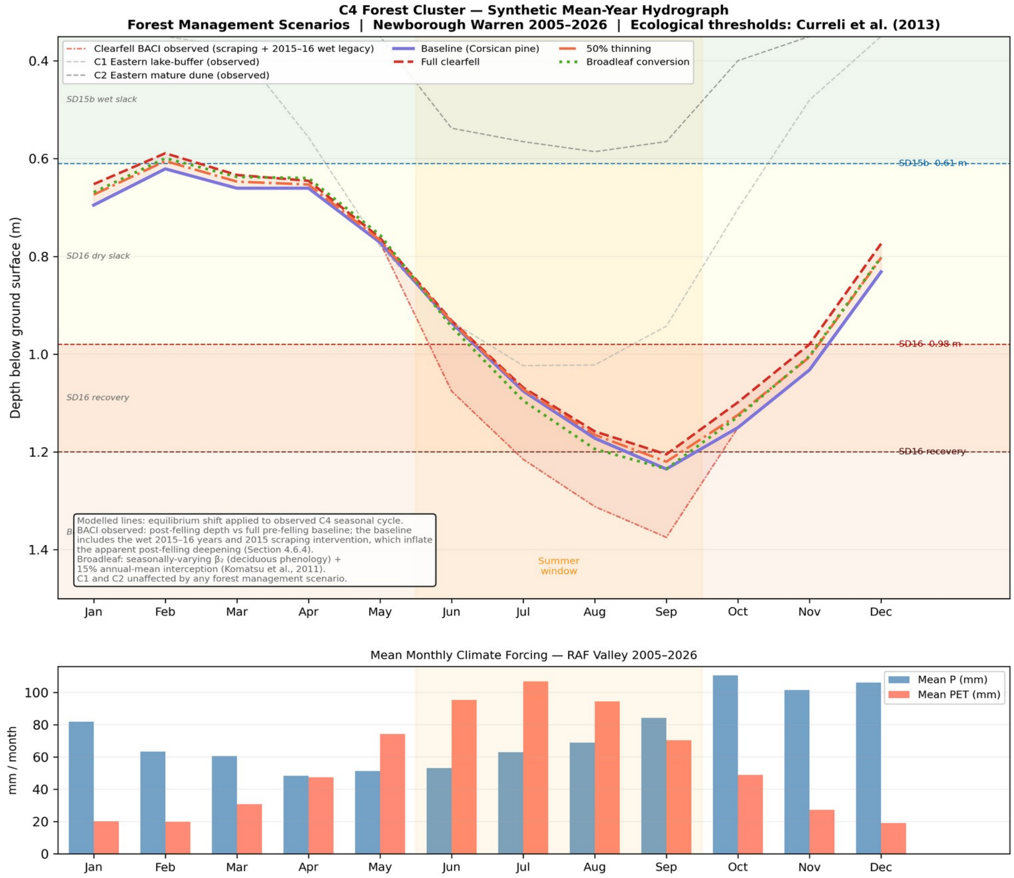

Modelled lines: equilibrium shift applied to observed C4 seasonal cycle.
BACI observed: post-felling depth vs full pre-felling baseline; the baseline includes the wet 2015–16 years and 2015 scraping intervention, which inflate the apparent post-felling deepening (Section 4.6.4).
Broadleaf: seasonally-varying β_2 (deciduous phenology) + 15% annual-mean interception (Komatsu et al., 2011).
C1 and C2 unaffected by any forest management scenario.

Figure 46. *Synthetic mean-year hydrograph for the C4 Forest cluster under four management scenarios: baseline Corsican pine, full clearfell, 50% thinning, and broadleaf conversion. Lines show equilibrium head shifts applied to the observed 2005–2026 C4 seasonal cycle, derived from scenario-specific* β_1 *(effective rainfall, via canopy interception) and* β_2 *(ET draw, for clearfell and thinning) adjustments. The BACI-observed post-felling raw displacement at WMC3 (summer minimum −1.591 m post-felling vs −1.556 m pre-scraping baseline; Section 4.6.2) is shown for context — the gap between the modelled clearfell line and the observed record reflects the site-wide* β_1 *decline that suppresses summer minima across all tiers equally (Section 4.6.6), masking the significant mean monthly recovery detected by the ANCOVA (+93 mm Impact,* $p = 0.019$*; Section 4.6.2). C1 and C2 observed mean seasonal cycles are shown for ecological context; no management scenario produces any detectable effect on these clusters. Broadleaf conversion: seasonally-varying* β_2 *applied to reflect deciduous phenology (lower winter ET, higher summer ET,* β_2 *summer multiplier ×1.113). Lower panel: mean monthly rainfall and PET from RAF Valley 2005–2026. Ecological zone colouring follows Curreli et al. (2013). Generated from 03_master_data.csv and 03_cluster_averages_maod.csv. (21_forestry_01_hydrograph.png).*

5.8 The Reach of Forest Management Interventions

The preceding sections have characterised the clearfell signal as concentrated in the mean monthly water level rather than in summer minima at the felled wells (Section 5.5), with the canopy acting both as a recharge-suppression mechanism through interception and as a microclimatic buffer on atmospheric demand at the water table (Section 5.6). A separate question is whether forest management in the C4/C5 zone could in principle improve water table conditions in the open dune clusters — C1 Lake Edge, C2 Dune and C3 Western Residual — where the ecological thresholds are most critical and the remaining wet-slack resource is concentrated.

The available evidence indicates that forest management in the C4/C5 zone does not measurably affect water table conditions in C1, C2 or C3. The spatial scenario framework (Section 4.10.2) produces no detectable Δh response at C1, C2 or C3 under any forest management scenario. The climate-corrected spatial step-change maps (Section 4.6.5; Figure 31) show that the felled area moved less than the site-wide drying trend — consistent with the significant ANCOVA result (Section 4.6.2) — but this effect was confined to the forest zone and did not propagate into the open dune clusters. Site-wide annual means under forest management are correspondingly small (Section 5.7.2), because only 14 of 66 reference wells are in the two forest clusters and the framework applies no spatial propagation from the forest zone into other clusters. Broadleaf conversion is near-neutral at the forest clusters themselves (Section 5.5.2) and produces no meaningful site-wide signal. The BACI monitoring record is consistent with this structural isolation: wells outside the forest zone showed no detectable felling-specific response once common-mode climate variability was accounted for (Section 4.6.7). The transect from uncleared plantation interior to clearfell core showed spatially uniform post-felling trajectories with no significant distance gradient (Section 4.6.5; Figure 33), indicating that the recharge recovery signal is real but local.

This structural isolation is not a framework artefact. The per-well analysis (Section 5.6.1) shows that the C4/C5 boundary corresponds to a substrate and elevation transition, and the drainage geography indicates that forest-zone perturbations propagate south-westward toward Caernarfon Bay rather than eastward into the C1 and C2 clusters (Section 5.7.2). The forest drawdown propagation analysis (Section 4.9.3, Figure 42) provides a quantitative upper bound: the steady-state head perturbation from the existing 75-year-old plantation is negligible at all C1 and eastern C2 wells, despite the plantation having had its full lifetime to reach equilibrium. Any recovery from clearfell would propagate outward on the same multi-year to multi-decadal timescale (Section 5.4.3), and the equilibrium recovery footprint would be spatially comparable. The P_flood equations (Table 16; Section 4.7.5) characterise the consequence at the forest clusters themselves: at Forest cluster summer minima of approximately 1.8 m below ground, the required winter rainfall to achieve flooding exceeds anything in the 95-year RAF Valley record regardless of canopy management.

The clusters where the ecological stakes are highest — C1 and C2, whose summer minima sit within 0.2–0.4 m of the SD15b and SD16 thresholds — show no detectable response to forest management in the C4/C5 zone, and their summer minimum trajectory is driven by the site-wide climate signal. Winter flooding in the Western Residual is controlled almost entirely by antecedent summer minimum depth (Table 13); winter rainfall amount is non-significant in C3, C4 and C5, and secondary even in C1 and C2 (Section 5.2.4). The primary restoration lever is therefore summer minimum depth, over which forest management in the C4/C5 zone has no detectable influence at C1 or C2.

5.8.1 P_flood Recoverability and the Operational Boundary

The P_flood threshold mapping (Section 4.7.4, Figures 27–29) characterises the recoverability of each well in the network under realistic rainfall conditions. Of the 87 classified wells, 62 (71%) are reachable under mildly wet winters ($\lambda < 1.5$) — these are concentrated in C1 and C2 and represent the zone where targeted scraping can shift a marginal well into reach of a mildly wet winter. A further 16 wells (18%) require the wettest winters in the instrumental record ($1.5 \leq \lambda < 2.0$), and the remaining 9 (10%) are structurally unreachable under the present climate. The unreachable wells are overwhelmingly in the forest zone: at C4 summer minima of approximately 1.8 m below ground, the required winter rainfall exceeds anything in the 95-year RAF Valley record regardless of canopy management. C3 Western Residual occupies the middle ground — most C3 wells sit within or just above the wet-winter-reachable band (median $\lambda = 1.28$) but the deeper-end wells approach $\lambda = 1.9$, and C3's near-zero winter flooding record (Section 5.1) indicates that even nominally reachable thresholds are rarely achieved in practice. The recoverability map delineates a clear operational boundary: the target zone for direct intervention is C1/C2 where λ values cluster close to 1.0 and a single scraping event can bring a well within reach of a mildly wet winter.

5.8.2 Climate Trajectory as the Binding Constraint

The clearfell produced a significant mean water table recovery against the Forest control (Section 5.5), but the recovery is concentrated in the mean monthly water level rather than in summer extremes, and the site-wide β_1 decline driving summer minimum deterioration appears in both felled and unfelled tiers (Section 5.5.1). The seasonal pattern at Newborough places the operational weight of canopy management on the winter recharge season, while the ecologically critical constraint — summer minimum depth — is set primarily by the climate trajectory operating across all clusters regardless of canopy cover.

Whether this seasonal asymmetry generalises to other coastal dune sites depends on the local water balance. The seasonal prediction equations at Newborough show that summer minimum depth dominates winter flooding prediction in the western and forest blocks, while winter rainfall is the dominant predictor only at the eastern lake-buffered and shallow till clusters (Table 13). Where this pattern holds at other sites, canopy management would address a different season from the ecologically critical one. Where the local water balance is winter-rainfall-limited rather than summer-minimum-limited, the calculation would differ. The Newborough finding cannot be directly transferred without site-specific seasonal partitioning of the flooding-control mechanism.

At Newborough specifically, the progressive redistribution of rainfall toward wetter winters and drier summers under a changing climate is visible in the declining summer minimum trends (Section 4.8.1). The corrected scenario projections (Section 4.10.2; Section 5.5.2) predict small positive net responses at the forest clusters under clearfell and thinning, and near-neutral or slightly negative responses under broadleaf conversion. These magnitudes are substantially smaller than the projected climate-driven changes over the same horizon

(Section 5.7.1). Forest management perturbations do not propagate to the C1/C2 zone where the ecological need is greatest (Section 5.8). The operational domain for direct hydrological intervention at Newborough therefore lies in topographic scraping in the C1/C2 zone where the P_flood threshold remains achievable under realistic winter rainfall (Section 5.8.1).

The management challenge at Newborough is hydrological adaptation rather than restoration to a prior state — maintaining flooding capacity in an increasingly seasonal climate, while the summer drought trajectory approaches the ecological viability thresholds identified by Curreli et al. (2013) within the next one to two decades.

5.9 Implications for Restoration and Monitoring

Summer minima are the primary restoration lever. Antecedent summer depth is the dominant predictor of peak winter water table in all western and forest clusters (Section 5.2; Table 13), with winter rainfall amount secondary or non-significant. Restoration success therefore depends on raising summer water table floors through topographic modification and reduced evaporative demand, not on maximising winter recharge. This reframing has direct implications for monitoring design: summer dipwell readings are more diagnostically informative than winter peak measurements for assessing restoration trajectory in every cluster.

Topographic scraping is the most directly effective intervention identified in the monitoring record, with a measured improvement in summer minimum depth at CEH36 during the pure scraping era representing more than half the gradient separating the SD15b and SD16 community thresholds (Section 5.4). The benefit operates on summer minima and is achievable within the C1/C2 zone where the P_flood threshold remains within reach of mildly wet winters (Section 5.8.1). Scraping steepens the local hydraulic gradient and draws the water table downward upslope (Section 5.4.3), so the spatial planning of any scraping programme must account for the cascade risk and avoid amplifying the existing coastal drainage gradient. The background climate drying trend continues to operate at the scraped site as elsewhere, and the structural benefit at CEH36 has diminished — though not reversed — over the post-felling period. Re-scraping intervals cannot yet be estimated from the available record.

Plantation clearance as implemented in 2017 produced a statistically significant mean water table recovery against the Forest control (Section 4.6.2), confirmed by synthetic extension within the felled compartment (Section 4.6.7). The clearfell signal is concentrated in the mean monthly water level rather than in summer minima, where felled and unfelled wells track each other (Section 5.5). The corrected scenario projections predict small positive net responses at the forest clusters under clearfell and thinning, and near-neutral responses under broadleaf conversion; these magnitudes are substantially smaller than the projected climate-driven Δh range over the same horizon (Section 5.5.2). The spatial framework and BACI record together indicate that forest management perturbations in the C4/C5 zone do not propagate to C1 and C2 where the ecological need is greatest (Sections 5.7, 5.8).

The prediction equations, P_flood thresholds and scenario framework are available as interactive tools hosted on the Newborough Warren Groundwater Research project page (Hollingham, 2026b; https://newbroman.github.io/Newborough_Hydrology/): a per-well flood probability forecaster driven by the P_flood threshold equations, a scenario viewer implementing the SSM equilibrium framework for exploring climate and management scenarios, and a network scatter plot of seasonal extremes referenced against the Curreli et al. (2013) thresholds. The three tools operate on different aspects of the analytical framework and are complementary.

Three priorities for future research are identified. First, slug tests at two to three representative wells per cluster and a ground-penetrating radar survey of aquifer thickness would enable a calibrated continuous-flow spatial model and directly test the finding that substrate thickness is the dominant control on within-forest coefficient variation (Section 5.6.1). These measurements would also lift the principal limitation on the spatial framework identified in Section 5.7.3, where the absence of calibrated hydraulic conductivity and aquifer thickness data prevents direct testing of cluster-to-cluster propagation. Second, the C5 Coastal Forest trajectory — the steepest decline in the network in both seasons, with three unresolved candidate explanations (Section 4.8.1; Section 5.7.2) — requires continued monitoring alongside targeted measurement of coastal retreat rates, since coastal boundary retreat is one of the candidate explanations and is the one most directly measurable in the field. Third, the relationship between canopy density and microclimatic shielding could be characterised through monitoring of β_2 responses at wells adjacent to blight-affected stands as the Dothistroma needle blight trajectory progresses (Section 5.6.4); this would be an opportunistic rather than designed study, but the natural experiment is already in progress and would require only sustained monitoring at existing or modestly extended sites.

6. Limitations

6.1 Temporal resolution and conservative volatility estimates

The reliance on monthly-averaged groundwater levels mathematically dampens the event-scale fluctuations inherent to highly conductive unconfined aquifers. Because daily peaks and troughs of individual storm events are smoothed over a 30-day window, the derived water balance parameters are optimised for seasonal trend analysis rather than event-based flood routing. This is a standard constraint in historic groundwater modelling (Taylor and Alley, 2001), and means the model captures the overall volumetric water balance while inherently underestimating daily transient spikes. The assessment of seasonal amplitude presented here is therefore conservative: the actual event-scale fluctuation — and any periodic severing of the capillary fringe from the root zone — is likely more pronounced than the monthly averages depict.

6.2 Thornthwaite PET and sensitivity of absolute thresholds

The Thornthwaite method was selected for its compatibility with the low-cost replication objective (Section 3.9), and its systematic biases cancel in all cluster-to-cluster and before-versus-after comparisons because the same PET series is applied uniformly across all clusters and time periods (Section 3.1.2). However, the absolute values produced by the prediction equations and P_flood thresholds are sensitive to the PET input. Thornthwaite PET overestimates evaporative demand by approximately 10–15% relative to Penman-Monteith in humid maritime climates (Donohue et al., 2010). Because the SSM's atmospheric draw term is $\beta_2 \times \text{PET}$, a 15% reduction in PET would reduce the fitted β_2 coefficients proportionally, with the effect concentrated in the summer months when PET is largest. The P_flood threshold — which depends on the winter recharge demand required to offset the preceding summer's atmospheric losses — would decrease, meaning the critical rainfall target would be easier to achieve than the current equations suggest. The scenario framework Δh estimates would be similarly affected: the absolute magnitude of forest management responses would be smaller if PET were lower, because the canopy interception and atmospheric draw perturbations are scaled against the PET baseline. The relative ranking of scenarios, the spatial pattern of responses, and the finding that forest management perturbations do not propagate to C1/C2 are all independent of the PET method because they depend on coefficient ratios and spatial structure rather than absolute magnitudes. Users of the P_flood equations and scenario framework should therefore treat the absolute thresholds as conservative upper bounds on rainfall demand, consistent with the upper-bound interpretation applied to the WTF-derived specific yield estimates (Section 4.2.4).

6.3 The NW10 broadleaf comparison — a single uncontrolled well

The comparison of NW10 (broadleaf block, established 1993–1996) with adjacent Corsican pine wells represents the only direct observational evidence of species conversion effects within the monitoring network, but it is subject to several limitations that preclude a fully controlled comparison. NW10 is a single well in a single block; replication across multiple broadleaf wells and blocks would be required to establish a statistically robust species effect independent of local variation. The comparison wells (CEH2, CEH34) differ from NW10 in ground elevation, distance from the ridge, and drainage rate (NW10 $\beta_3 = 0.046$ vs C4 median

0.019), and whilst the direction of the summer minimum contrast is consistent across all comparators the absolute magnitude is sensitive to these positional differences. Most significantly, the broadleaf stand was characterized by a dense bramble understorey throughout most of the monitoring period, which maintained high ground-level evapotranspiration demand and confounded the tree canopy signal. The stand has only begun to approach a closed canopy configuration from approximately 2022, leaving fewer than four years of post-closure record. The conclusions drawn from NW10 should therefore be treated as indicative observations from a natural experiment rather than confirmed causal attributions, and will require validation against a longer post-closure record and, ideally, against additional broadleaf wells installed within the block.

6.4 Spatial analysis — absence of calibrated continuous-flow modelling

The spatial framework used in this study (Sections 3.8, 4.9) is a per-well SSM scenario aggregator interpolated by inverse-distance weighting, not a calibrated continuous-flow model. It is well-suited to the questions the paper addresses — how the per-well β parameterisation aggregates spatially under scenarios, and whether the BACI-observed forest response propagates beyond the forest zone — but cannot represent cluster-to-cluster hydraulic coupling or the diffusive redistribution of localized recharge perturbations. Hydraulic conductivity is constrained by a single tracer test (K = 6.0 m/day; Betson et al., 2002) and aquifer thickness in the western interior is inferred from four borehole logs across 700 ha (Bristow, 2003). These are the same parameter-sparsity constraints that produced calibration difficulties in the Betson et al. (2002) MODFLOW model at this site. The priority field measurements required to enable a calibrated continuous-flow model are identified in Section 5.9.

6.5 BACI experiment — network design limitations

The clearfell zone and scraping sites were selected on different grounds: scraping was a targeted conservation intervention, while the clearfell was an experimental felling recommended by previous researchers and implemented as an operational forestry action. Neither was randomly assigned, which introduces the possibility that pre-existing site characteristics could confound the intervention signal. The monitoring network was not designed as a purpose-built propagation transect. The transect analysis in Section 4.6.7 (Figure 33) uses the available surveillance wells to test for an inland gradient in the post-felling signal and finds none, with post-felling step changes distributed uniformly from the plantation interior to the felling core. All 17 wells in the five-tier network have a minimum of eight years' pre-felling baseline (Section 3.5.5). Wells with shorter records (FE1–4, LIS1, NW8B) were excluded from the analysis. The network permits a test of whether clearfell-driven effects propagate into the adjacent warren, but does not provide fine-grained spatial resolution of any residual propagation signal at distances under 100 m from the clearfell edge, nor a purpose-built warren transect perpendicular to the clearfell boundary.

6.6 Coastal-erosion confound — first-order spatial control rather than time-lagged propagation

Both intervention analyses treat the coastal-erosion confound through covariates and control selection rather than through explicit time-lagged modelling. The clearfell ANCOVA includes an Easting × time interaction that captures average differential deterioration with proximity to the coast across the observation window, and the scraping analysis tracks coastal-driven drift through CEH4's CUSUM trajectory and through the synthetic-control and SSM-forward-residual robustness checks. Neither approach explicitly represents the multi-year to multi-decadal propagation timescale (L^2Sy/K, Section 5.4.3) over which coastal erosion signals travel inland through the aquifer. The result is that both intervention analyses estimate the felling-attributable and scraping-attributable steps cleanly within the available analysis windows, but cannot fully separate the contemporaneous management signal from the time-integrated coastal-erosion signal arriving from earlier historical retreat.

A direct test of the spatial component of this confound is now provided by the network-scale gradient regression of Section 4.8.1, which fits the strip-aquifer reach and coast-edge deepening rate directly to the 21-year monthly panel and corroborates the clearfell ANCOVA's easting × time covariate against an independently fitted physics-based model (Section 4.6.3, Table 8). At the headline Forest Impact comparison the BACI's easting × time absorption matches the gradient model's prediction within sampling error, indicating that the first-order spatial control is doing what coastal-retreat physics says it should. The Edge and Climate-tier comparisons absorb additional non-coastal-retreat drift (Section 5.5), so the spatial covariate's role in those comparisons is broader than coastal-retreat correction alone. The remaining limitation is the absence of an explicit transient boundary-condition flow model, which would refine the time-lagged propagation estimates and resolve the residual covariate absorption beyond what the steady-state gradient model captures. Direct measurement of present-day coastal retreat rates at the south-western dune front remains a monitoring priority

6.7 Water balance residuals — modelling artefact versus physical flux

The cluster-level water balance closes to within 2.5% of total losses at all five clusters under the displacement formulation (Table 4a). However, the water balance residual field (Section 4.9.6, Figure 45) shows that this near-closure is not spatially uniform: positive residuals concentrate along the northern forest margin and ridge flank, declining to near zero across the open dune and becoming weakly negative at some south-eastern margin wells. This spatial structure is real and reproducible but its physical interpretation is ambiguous. The residual could reflect a genuine unmodelled lateral flux (with the ridge aquifer as the leading geometric candidate), inadequacies of the linear lumped-parameter SSM in representing nonlinear dynamics, minor β coefficient sensitivity, or some combination. Two independent diagnostic tests on the SSM residuals (Supplementary Notes S5 and S6) returned null results for specific mechanistic attributions — the lag-distance test did not corroborate ridge transport, and the sunshine-hours test did not identify unmodelled evaporative demand. The depth-dependent PET diagnostic (Section 5.2) confirms that a real depth-coupling effect exists but is adequately absorbed by the standard SSM's fixed β_2 for operational purposes. Until the residual attribution is resolved through either a nonlinear or spatially explicit model, or through direct field measurement of the candidate source locations, any management strategy whose rationale depends on a specific mechanistic interpretation of the residual should be evaluated against this uncertainty.

6.8 Topographic proxies, aeolian deflation, and bedrock causation

The strong spatial correlation between DEM-derived drainage pathways and the hydrogeological cluster boundaries confirms that surface topography acts as a reliable operational proxy for subsurface flow divides. However, this correlation does not imply direct causation: surface sand deposition patterns do not dictate groundwater divides. Both are co-determined by the underlying architecture of glacial till and bedrock ridges, where present. In coastal dune systems this structural relationship is mediated by aeolian deflation — wind erosion is physically limited by the capillary fringe, as cohesive damp sand resists aeolian transport (Bristow and Bailey, 2001; Hesp and Thom, 1990), and because the regional water table is governed by subsurface geology, that geology indirectly controls the baseline limits of surface erosion (Haitjema and Mitchell-Bruker, 2005). Surface topography at Newborough Warren is therefore a physical expression of subsurface hydrological structure, supporting the use of the DEM as a proxy for regional groundwater routing in the absence of high-resolution bedrock mapping.

7. Conclusions

This study provides the first network-wide, multi-method quantitative characterisation of groundwater behaviour at Newborough Warren, integrating 89 wells, 21 years of monitoring and a 27-step analytical pipeline across eleven phases. The principal conclusions are:

1. **Five hydrogeological clusters govern site hydrology.** The Eastern Block (C1 Lake Edge, C2 Dune) responds rapidly to rainfall with high recharge sensitivity ($\beta_1 = 4.58, 3.87$) and fast drainage ($\beta_3 = 0.090, 0.063$). The Western Residual (C3) acts as a deep buffer (LCSC = 28.0%). The two Forest clusters (C4 Main Forest, C5 Coastal Forest) share elevated LCSC values (39.7%, 41.0%) reflecting canopy interception, but differ in substrate setting: C4 overlies irregular bedrock topography on the elevated ridge flank (high $\beta_2 = 2.504$, weak $\beta_3 = 0.021$), where buried ridges impede lateral drainage; C5 occupies deeper, more uniform coastal sand (lower $\beta_2 = 1.374$, faster $\beta_3 = 0.045$). Within the forest zone, elevation alone explains 95% of the variance in β_2, indicating substrate as a major contributor to the atmospheric draw coefficient, with depth-coupling and other factors also acting. Three independently derived parameters — drainage timescale, SSM forecasting gain and specific yield — place the five clusters in distinct, non-overlapping regions of diagnostic space (Figure 43), supporting the cluster structure independently of the Ward's linkage that defined it.
2. **Summer minimum depth, not winter rainfall, is the dominant control on dune slack ecological viability.** Antecedent summer minimum is the primary predictor of winter flooding potential in the Western Residual and Forest clusters, with winter rainfall non-significant (Table 13). The entire ecological gradient from SD15b wet slack to SD16 dry slack spans only 37 cm of summer minimum depth (Curreli et al., 2013). Any management intervention must be evaluated on its summer minimum effect.
3. **State-space modelling transforms forecast skill.** The median iterative NSE improvement of +0.62 over the traditional linear model — with only 1 of 66 SSM wells producing a negative iterative NSE compared with 24 of 66 for the TLM — indicates that explicit parameterisation of head-dependent drainage memory is essential for long-range prediction in this aquifer.
4. **Dune scraping delivers a measurable but spatially and temporally limited benefit whose placement within the drainage system is critical.** Three independent methods converge on a scraping benefit of 0.08–0.14 m at CEH36 (raw BACI +0.131 m, synthetic control +0.135 m, SSM forward residual +0.083 m), with a paired summer minimum shift of +195 mm ($p = 0.004$) — more than half the SD15b–SD16 ecological gradient. The benefit is geometric (permanent surface lowering) rather than a change in aquifer dynamics, retaining approximately 68% of the scrape depth as a long-term benefit. Scraping operates through the same physical mechanism as coastal erosion — lowering ground elevation to increase the displacement head and drainage gradient — and its spatial placement determines whether it complements or exacerbates the existing

coastal drainage regime. CEH21's limited response to the October 2023 intervention indicates that position within the coastal drainage pathway is a significant determinant of scraping efficacy. Priority targets are the C1/C2/C3 transitional wells where the aquifer base is stable and P_flood thresholds remain achievable (rainfall multiplier $\lambda < 1.5$).

5. **The clearfell experiment produced a statistically significant mean water table recovery against the Forest control.** The three-counterfactual ANCOVA-BACI found a significant clearfell step of +93 mm at the Impact tier ($p = 0.019$) and +153 mm at the Edge tier ($p < 0.001$), confirmed by synthetic extension of FE wells (+103 mm, $p = 0.011$). The recovery is concentrated in the mean monthly water level rather than in summer minima: the summer minima analysis returned no significant effect at the Impact tier and a small negative effect at the Edge tier (−72 mm, $p = 0.031$), opposite in sign to the monthly ANCOVA — summer minima at felled wells moved in parallel with the unfelled Forest controls. Four independent robustness methods yield lower estimates than the ANCOVA, consistent with the absence of the easting × time correction; they represent lower bounds on the clearfell-specific component. The corrected scenario projections predict small positive net annual responses to clearfell and thinning at the forest clusters (clearfell: C4 +4.7 mm w.e./month, C5 +8.8 mm w.e./month) and near-neutral responses under broadleaf conversion (C4 −3.4 mm, C5 +0.1 mm) — all substantially smaller than climate-driven changes over the same horizon. The +93 mm Impact-tier clearfell step is independently corroborated by the network-scale coastal-retreat gradient model (Section 4.6.3, Table 8), which shows that the ANCOVA's easting × time covariate is absorbing approximately the differential coastal-retreat signal that physics predicts ($z = -0.35$ between BACI absorption and model prediction for the Forest Impact comparison).
6. **Recharge efficiency (β_1) is declining site-wide, independently of any management intervention.** The coefficient decomposition reveals a β_1 decline across all five tiers of the BACI network — Impact, Edge, Forest Control, Coastal Control and Climate Control — consistent with changing within-month rainfall distribution, where the same monthly total arriving in fewer, more intense events separated by longer dry spells allows water held in the unsaturated zone between events to be drawn back out by ET before it can percolate to the water table. This decline is the strongest candidate mechanism for the observed summer minimum deterioration across the network, and it affects felled and unfelled wells equally. The monitoring evidence therefore characterises canopy removal as a winter-recharge intervention: it raised mean monthly water tables at the felled wells, but its effect on summer minimum depth is set by this site-wide β_1 decline that operates independently of canopy cover.
7. **Forest management perturbations do not propagate to the open dune clusters where ecological need is greatest.** The drawdown propagation analysis (Figure 42) shows that the steady-state head perturbation from the existing 75-year-old plantation is negligible (<1 mm) at all C1 and eastern C2 wells. The BACI transect analysis shows no felling-specific gradient (−0.4 mm/100 m, p =

0.26). The operational domain for direct hydrological intervention is the C1/C2 zone where P_flood thresholds remain achievable.

8. **WTF-derived specific yields converge across all clusters once canopy interception is corrected.** The corrected forest cluster Sy estimates (C4: 0.243, C5: 0.320) fall within the open dune range (0.210–0.341), indicating that the plantation's distinctive β-coefficient signature derives largely from canopy boundary conditions at the surface rather than from differences in subsurface storage architecture.
9. **The water balance closes to within 2.5% under the displacement formulation, but the residual field is spatially structured.** Positive residuals concentrate along the northern forest margin; the mechanistic attribution is unresolved. Two independent diagnostic tests returned null results for specific transport mechanisms.
10. **C5 Coastal Forest shows the steepest decline in the network in both summer (-0.038 m yr^{-1}, $p = 0.003$) and winter (-0.037 m yr^{-1}, $p = 0.035$).** At 3.9 times the summer decline rate of C1 — the only other cluster with a statistically significant summer trend — the magnitude of C5's divergence argues against a purely climate-driven explanation, since all clusters share the same climate forcing. Three candidate explanations — amplified head response per unit climate forcing due to thin saturated thickness at the coastal margin, coastal boundary retreat, and progressive canopy maturation — are not mutually exclusive, and the scale of the divergence is reinforced by a parallel 18–21% post-2018 collapse in seasonal amplitude (Section 4.2.1) indicating that recharge peaks are no longer reaching their pre-2018 heights. A network-scale non-linear panel regression on perpendicular distance to the eroding Caernarfon Bay shoreline (Section 4.8.1) attributes approximately 42% of C5's exceptional summer-minimum decline to a coastal-retreat gradient (coast-edge anomaly -30.6 mm yr^{-1} decaying linearly to climate background by 869 m inland; ΔAIC = 29 versus exponential decay). Progressive coastal boundary retreat is therefore quantified as a substantial but partial contributor to the C5 anomaly; the remaining -16 mm yr^{-1} residual at C5 is consistent with the other two candidate mechanisms operating in parallel. Direct measurement of present-day coastal retreat rates and continued monitoring of the CUSUM trajectory remain a priority.
11. **The intervention window is finite.** Summer minimum trends indicate C1 summer minima approaching the SD16 dry slack viability threshold around 2030–2032 under current trajectories. The analytical framework, prediction equations and interactive tools developed in this study provide the operational basis for evidence-based management decisions within that window.
12. **Progressive coastal boundary retreat is an independent and potentially accelerating threat whose full consequences are lagged.** The convergent deterioration of two spatially separated control wells (CEH4 and CEH22), independently of any management action, indicates that the aquifer is losing volume at its seaward edge. The aquifer's hydraulic diffusivity (L^2Sy/K of order 7–13 years at 570 m) means that the coastal erosion signal visible in the current

CUSUM trajectories represents erosion that occurred years earlier; if the rate of retreat is accelerating, as CEH22's trajectory from 2021 suggests, the full hydrological consequences have not yet propagated to interior wells. The network-scale gradient regression of Section 4.8.1 quantifies the steady-state coastal-retreat signal directly from 21 years of monitoring data: the coast-edge water-table anomaly above climate is approximately -30.6 mm yr^{-1}, decaying to zero over an inland reach of 869 m. C5 Coastal Forest sits within the inner third of this gradient zone; C4, C1 and C2 lie beyond the fitted reach. Direct measurement of present-day coastal retreat rates and continued monitoring of inland propagation through the CEH4 and CEH22 CUSUM trajectories remain a priority.

8. References

Arribas-Bel, D., Fleischmann, M., & others. (2020) *contextily: Context geo tiles in Python*. GitHub repository. https://github.com/geopandas/contextily

von Asmuth, J.R., Bierkens, M.F.P. and Maas, K. (2002) Transfer function-noise modelling in continuous time using predefined impulse response functions. Water Resources Research, 38(12), 1287. doi:10.1029/2001WR001136.

Bakker, M. and Schaars, F. (2019) Solving groundwater flow problems with time series analysis: you may not even need another model. *Groundwater*, 57(6), pp. 826–833. https://doi.org/10.1111/gwat.12927

Bear, J. (1972) Dynamics of Fluids in Porous Media. New York: American Elsevier.

Betson, M., Connell, L. and Bristow, C. (2002) Groundwater Modelling of Newborough Warren. In: Bristow, C. (ed.) The Impact of Forestry on Coastal Geomorphology at Newborough Warren/Ynys Llanddwyn NNR, SSSI, pSAC. Volume 5: Hydrogeology. Report to the Countryside Council for Wales, Contract FC 73-05-18. Department of Geological Sciences, UCL and School of Earth Sciences, Birkbeck University of London. July 2002.

Betson, M. and Scholefield, P. (2004) Implications for the Water Balance of Newborough Warren NNR of the Evapotranspiration of Herbaceous and Afforested Duneland Ecosystems. Unpublished report to Countryside Council for Wales. ADAS.

Beven, K. (2012) Rainfall-Runoff Modelling: The Primer. 2nd edn. Chichester: Wiley-Blackwell.

Bristow, C.S. (2003) The Impact of Forestry on Coastal Geomorphology at Newborough Warren / Ynys Llanddwyn NNR, SSSI, pSAC. Volumes 1–5. Contract FC 73-05-18, Final Report. Countryside Council for Wales. London: Birkbeck University of London.

Bristow, C.S. and Bailey, S.D. (2001) Non-invasive investigation of water table and structures in coastal dunes using ground-penetrating radar (GPR): implications for dune management. In: Houston, J.A., Edmondson, S.E. and Rooney, P.J. (eds) *Coastal Dune Management: Shared Experience of European Conservation Practice.* Liverpool University Press, pp. 408–417

Burnham, K.P. and Anderson, D.R. (2002) Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach. 2nd edn. New York: Springer.

Calder, I.R. and Newson, M.D. (1979) Land-use and upland water resources in Britain — a strategic look. Water Resources Bulletin, 15(6), pp. 1628–1639.

Calder, I.R. (2007) "Forests and water — Ensuring forest benefits outweigh water costs." *Forest Ecology and Management* 251, 110–120. doi:10.1016/j.foreco.2007.06.015

Caliński, T. and Harabasz, J. (1974) A dendrite method for cluster analysis. *Communications in Statistics*, 3(1), pp. 1–27.

Callaghan, D. A., van Willegen, L., Williams, G., Hollingham, M., & Jones, L. (2021) Pony trails, hydrology and habitat restoration: aspects of the ecology of *Petalophyllum ralfsii* in a Welsh oceanic dune system. *Journal of Bryology*, 43(2), 150-161. https://doi.org/10.1080/03736687.2020.1819719

Clarke, D. and Sanitwong Na Ayutthaya, S. (2010) Predicted effects of climate change, vegetation and tree cover on dune slack habitats at Ainsdale on the Sefton Coast, UK. Journal of Coastal Conservation, 14(2), pp. 115–125. doi:10.1007/s11852-009-0066-7.

Collenteur, R.A., Bakker, M., Caljé, R., Klop, S.A. and Schaars, F. (2019) Pastas: open source software for the analysis of groundwater time series. *Groundwater*, 57(6), pp. 877–885. https://doi.org/10.1111/gwat.12925

Conrad, O. et al. (2015) System for Automated Geoscientific Analyses (SAGA) v. 2.1.4. *Geoscientific Model Development*, 8, 1991–2007.

Cottingham, P. (1994) The effects of afforestation on the hydrology of Newborough Warren dune system. Unpublished MSc dissertation, University of Wales, Bangor.

Curreli, A., Wallace, H., Freeman, C., Hollingham, M., Stratford, C., Johnson, H. and Jones, L. (2013) Eco-hydrological requirements of dune slack vegetation and the implications of climate change. Science of the Total Environment, 443, pp. 910–919. doi:10.1016/j.scitotenv.2012.11.035.

Dargie, T.C.D. (1995) Sand dune vegetation survey of Great Britain: a national inventory. Part 3: Wales. Peterborough: Joint Nature Conservation Committee.

Davy, A.J., Grootjans, A.P., Hiscock, K. and Petersen, J. (2006) Development of Eco-Hydrological Guidelines for Dune Habitats — Phase 1. English Nature Research Reports, No. 696. Peterborough: English Nature.

Deng, Z., Priestley, S.C., Guan, H., Love, A.J. and Simmons, C.T. (2013). Canopy enhanced chloride deposition in coastal South Australia and its application for the chloride mass balance method. *Journal of Hydrology*, 497, 62–70. https://doi.org/10.1016/j.jhydrol.2013.05.038

Donohue, R.J., McVicar, T.R. and Roderick, M.L. (2010) Assessing the ability of potential evaporation formulations to capture the dynamics in evaporative demand within a changing climate. Journal of Hydrology, 386(1–4), pp. 186–197. https://doi.org/10.1016/j.jhydrol.2010.03.020

Fetter, C. W. (2001) *Applied Hydrogeology* (4th ed.). Prentice Hall.

Flyamer, I. M., Illingworth, R. S., & others. (2020) *adjustText: Automatically adjust text position in matplotlib plots to minimize overlaps*. GitHub repository. https://github.com/Phlya/adjustText

Freeman, S. (2008) *The hydrological impact of Corsican pine plantation on the dune system at Newborough Warren, Anglesey.* Unpublished BSc Honours dissertation, University of Wales, Bangor.

Freeze, R.A. and Cherry, J.A. (1979) Groundwater. Englewood Cliffs, NJ: Prentice-Hall.

Gash, J.H.C. (1979) An analytical model of rainfall interception by forests. Quarterly Journal of the Royal Meteorological Society, 105(443), pp. 43–55.

Gash, J.H.C., Wright, I.R. and Lloyd, C.R. (1980) Comparative estimates of interception loss from three coniferous forests in Great Britain. Journal of Hydrology, 48, pp. 89–105.

Giddens, K.M., Parfitt, R.L. and Percival, H.J. (1997) Comparison of some soil properties under *Pinus radiata* and pasture. *New Zealand Journal of Agricultural Research*, 40(4), pp. 443–450

Gillies, S., & others. (2013) *Rasterio: Geospatial raster I/O for Python programmers*. GitHub repository. https://github.com/rasterio/rasterio

Gotelli, N.J. and Ellison, A.M. (2004) A Primer of Ecological Statistics. Sunderland, MA: Sinauer Associates.

Grootjans, A. P., Adema, E. B., Bekker, R. M., & Lammerts, E. J. (2004) Why young coastal dune slacks sustain a high biodiversity. In M. L. Martínez & N. P. Psuty (Eds.), *Coastal Dunes: Ecology and Conservation* (Ecological Studies, Vol. 171, pp. 85-101). Springer-Verlag, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-06259-5_6

Grootjans, A.P., Geelen, H.W.T., Jansen, A.J.M. and Lammerts, E.J. (2002) Restoration of coastal dune slacks in the Netherlands. Hydrobiologia, 478(1–3), pp. 181–203.

Grootjans, A.P., Shahrudin, R., van der Craats, A., Kooijman, A., Oostermeijer, G., Peterson, J., Amatirsat, D., Bland, C. and Stuyfzand, P.J. (2017) Window of opportunity of *Liparis loeselii* populations during vegetation succession on the Wadden Sea islands. *Journal of Coastal Conservation*, 21, 45–55. https://doi.org/10.1007/s11852-016-0448-6

Gumuła-Kawęcka, A., Jaworska-Szulc, B., Szymkiewicz, A., Gorczewska-Langner, W., Pruszkowska-Caceres, M., Angulo-Jaramillo, R. and Šimůnek, J. (2021) Estimation of groundwater recharge in a shallow sandy aquifer using unsaturated zone modeling and water table fluctuation method. Journal of Hydrology, 605, 127283. https://doi.org/10.1016/j.jhydrol.2021.127283

Gumuła-Kawęcka, A., Jaworska-Szulc, B., Jefimow, M. (2024) "Climate change impact on groundwater resources in sandbar aquifers in southern Baltic coast." *Scientific Reports* 14, 11828. doi:10.1038/s41598-024-62522-0

Gupta, H. V., Wagener, T., & Liu, Y. (2008) Reconciling theory with observations: elements of a diagnostic approach to model evaluation. *Hydrological Processes*, 22(18), 3802–3813. https://doi.org/10.1002/hyp.6989

Haitjema, H.M. and Mitchell-Bruker, S. (2005) Are water tables a subdued replica of the topography? Groundwater, 43(6), pp. 781–786.

Hannah, D. M., Smith, B. P. G., Gurnell, A. M., & McGregor, G. R. (2000) An approach to characterize data streams in hydrology. *Journal of Hydrology*, 232(1-4), 71-91. https://doi.org/10.1016/S0022-1694(00)00214-9

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., ... & Oliphant, T. E. (2020) Array programming with NumPy. *Nature*, 585(7825), 357-362. https://doi.org/10.1038/s41586-020-2649-2

Healy, R.W. and Cook, P.G. (2002) Using groundwater levels to estimate recharge. *Hydrogeology Journal* 10, 91–109. doi:10.1007/s10040-001-0178-0

Hennig, C. (2007) Cluster-wise assessment of cluster stability. *Computational Statistics and Data Analysis*, 52(1), pp. 258–271.

Hesp, P.A. and Thom, B.G. (1990) Geomorphology and evolution of active transgressive dunefields. In: Nordstrom, K.F., Psuty, N. and Carter, R.W.G. (eds.) Coastal Dunes: Form and Process. Chichester: John Wiley and Sons, pp. 253–288.

Hollingham, M. (2026b) Newborough Warren Hydrogeological Modelling Suite Scenario Viewer. Available at: https://newbroman.github.io/Newborough_Hydrology/scenario_viewer.html [Accessed May 2026].

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90-95. https://doi.org/10.1109/MCSE.2007.55

Hypolite, J., Arvor, D., Bertrand, G., & Fovet, O. (2021) A state-space approach to groundwater level forecasting: Comparisons with lumped and machine learning models. *Journal of Hydrology*, 603, 127115. https://doi.org/10.1016/j.jhydrol.2021.127115

Jennings, T.A.H. (1990) The changing hydrology of the Newborough Dune System due to afforestation. Unpublished M.Sc. dissertation, University of North Wales, Bangor.

Jones, M.L.M., Sowerby, A. and Wallace, H.A. (2007) Better Understanding of Soil Resources — Dune Stabilisation and Rates of Soil Development on Welsh Dune Systems. CCW Contract Science Report No. 796. Bangor: Countryside Council for Wales / Centre for Ecology and Hydrology.

Jordahl, K., den Bossche, J. V., Fleischmann, M., Wasserman, J., McBride, P., Gerard, J., ... & Ward, B. (2020). *geopandas/geopandas: v0.8.1*. Zenodo. https://doi.org/10.5281/zenodo.3946761

Knotters, M. and van Walsum, P.E.V. (1997) Estimating fluctuation quantities from time series of water-table depths using models with a stochastic component. Journal of Hydrology, 197(1–4), pp. 25–46.

Knotters, M. and Bierkens, M.F.P. (2000) Physical basis of time series models for water table depths. Water Resources Research, 36(1), pp. 181–188. https://doi.org/10.1029/1999WR900288

Komatsu, H., Kume, T. and Otsuki, K. (2011) Increasing annual runoff — broadleaf or coniferous forests? *Hydrological Processes*, 25(2), 302–318. https://doi.org/10.1002/hyp.7898

Krause, P., Boyle, D. P., & Bäse, F. (2005) Comparison of different efficiency criteria for hydrological model assessment. *Advances in Geosciences*, 5, 89-97. https://doi.org/10.5194/adgeo-5-89-2005

Manis, J., Shanafield, M., Cook, P. G., Andersen, M. S., & McCallum, J. L. (2022) Quantifying internal drainage and vertical leakage in multi-layered aquifer systems using state-space representations. *Groundwater*, 60(4), 512–525. https://doi.org/10.1111/gwat.13171

McKinney, W. (2010) Data structures for statistical computing in Python. In *Proceedings of the 9th Python in Science Conference* (Vol. 445, pp. 51-56). https://doi.org/10.25080/Majora-92bf1922-00a

Met Office (2018) UKCP18 Regional Projections on a 12 km grid over the UK for 1980–2080. Centre for Environmental Data Analysis. Met Office Hadley Centre. Available at: https://catalogue.ceda.ac.uk/uuid/b4d24b3df3754b9d9028447eb3cbd878

Nash, J.E. and Sutcliffe, J.V. (1970) River flow forecasting through conceptual models. Part I — A discussion of principles. Journal of Hydrology, 10(3), pp. 282–290.

Obergfell, C., Bakker, M. and Maas, K. (2019) Identification and explanation of a change in the groundwater regime using time series analysis. *Groundwater*, 57(6), pp. 886–894. https://doi.org/10.1111/gwat.12891

Osenberg, C.W., Schmitt, R.J., Holbrook, S.J., Abu-Saba, K.E. and Flegal, A.R. (1994) Detection of environmental impacts: natural variability, effect size, and power analysis. *Ecological Applications*, 4(1), 16–30.

Page, E.S. (1954) Continuous inspection schemes. Biometrika, 41(1/2), pp. 100–115.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, É. (2011) Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12: 2825–2830.

Peters, E., Torfs, P.J.J.F., van Lanen, H.A.J. and Bier, G. (2003) Propagation of drought through groundwater — a new approach using linear reservoir theory. Hydrological Processes, 17(15), pp. 3023–3040.

Peterson, T.J. and Western, A.W. (2014) Nonlinear time-series modeling of unconfined groundwater head. *Water Resources Research*, 50(10), pp. 8330–8355. https://doi.org/10.1002/2013WR014800

Pye, K. and Saye, S. (2005) The Geomorphological Response of Welsh Sand Dunes to SeaLevel Rise Over the Next 100 Years and the Management Implications for SAC and SSSI Sites. CCW Contract Science Report No. 670. Bangor: Countryside Council for Wales.

Ranwell, D.S. (1959) Newborough Warren, Anglesey 1. The dune system and dune slack habitat. Journal of Ecology, 47, pp. 571–601.

Ranwell, D.S. (1972) *Ecology of Salt Marshes and Sand Dunes*. Chapman and Hall, London. The classic UK dune ecology text; not currently cited in the paper, so this is a new bibliography entry.

Rao, A. R., & Srinivas, V. V. (2006) Regionalization of watersheds by fuzzy cluster analysis. *Journal of Hydrology*, 318(1-4), 57-79. https://doi.org/10.1016/j.jhydrol.2005.06.004

Robins, N. S., Jones, M. L. M., & Farr, G. (2013) Shallow groundwater in the dune and slack environment: The implications for management. *Quarterly Journal of Engineering Geology and Hydrogeology*, 46(3), 261–265. https://doi.org/10.1144/qjegh2012-045

Robinson, E.L., Brown, M.J., Kay, A.L., Lane, R.A., Chapman, R., Bell, V.A. and Blyth, E.M. (2023) CHESS-SCAPE: high-resolution future projections of multiple climate scenarios for the United Kingdom derived from bias-corrected UKCP18 regional climate model output. Earth System Science Data, 15(12), pp. 5371–5401. https://doi.org/10.5194/essd-15-5371-2023

Robinson, M. and Dupeyrat, A. (2005) Effects of commercial timber harvesting on streamflow regimes in the Plynlimon catchments, mid-Wales. Hydrological Processes, 19(6), pp. 1213–1226.

Rutter, A.J., Kershaw, K.A., Robins, P.C. and Morton, A.J. (1971) A predictive model of rainfall interception in forests. I: Derivation of the model from observations in a plantation of Corsican pine. Agricultural Meteorology, 9, pp. 367–384.

Scanlon, B.R., Healy, R.W. and Cook, P.G. (2002) Choosing appropriate techniques for quantifying groundwater recharge. *Hydrogeology Journal*, 10(1), 18–39. https://doi.org/10.1007/s10040-001-0176-2

Seabold, S., & Perktold, J. (2010). statsmodels: Econometric and statistical modeling with python. In *Proceedings of the 9th Python in Science Conference* (Vol. 57, pp. 61-66). https://doi.org/10.25080/Majora-92bf1922-011

Šimůnek, J., Šejna, M., Saito, H., Sakai, M., & van Genuchten, M. T. (2008) *The HYDRUS-1D Software Package for Simulating the One-Dimensional Movement of Water, Heat, and Multiple Solutes in Variably Saturated Media, Version 4.0*. HYDRUS Software Series 3, Department of Environmental Sciences, University of California Riverside, Riverside, California, USA.

Sival, F.P. and Grootjans, A.P. (1996). Dynamics of seasonal bicarbonate supply in a dune slack: effects on organic matter, nitrogen pool and vegetation succession. *Vegetatio*, 126, 39–50.

Smokorowski, K.E. and Randall, R.G., (2017). Cautions on using the Before-After-Control-Impact design in environmental effects monitoring programs. *FACETS* 2: 212–232. https://doi.org/10.1139/facets-2016-0058

Stewart-Oaten, A., Murdoch, W.W. and Parker, K.R. (1986) Environmental impact assessment: 'pseudoreplication' in time? *Ecology* 67(4): 929–940. https://doi.org/10.2307/1939815

Stratford, C. (2006) Review of Hydrological Reports for Newborough Warren, Anglesey. Unpublished report. Wallingford: Centre for Ecology and Hydrology.

Stratford, C., Ratcliffe, J., Hughes, A.G., Roberts, J. and Robins, N.S. (2007) Complex interaction between shallow groundwater and changing woodland, surface water, grazing and other influences in partly wooded duneland in Anglesey, Wales. In: Ribeiro, L., Chambel, A. and Condesso de Melo, M.T. (eds) IAH 35th Congress, Groundwater and Ecosystems, Lisbon.

Stuyfzand, P. J. (1993) *Hydrochemistry and hydrology of the coastal dune area of the Western Netherlands*. (Doctoral dissertation, Vrije Universiteit Amsterdam)

Taylor, C.J. and Alley, W.M. (2001) Ground-Water-Level Monitoring and the Importance of Long-Term Water-Level Data. U.S. Geological Survey Circular 1217. Denver, CO: U.S. Geological Survey.

Thornthwaite, C.W. and Mather, J.R. (1957) Instructions and tables for computing potential evapotranspiration and the water balance. Publications in Climatology, 10(3). Centerton, NJ: Drexel Institute of Technology, Laboratory of Climatology.

Underwood, A.J. (1992). Beyond BACI: the detection of environmental impacts on populations in the real, but variable, world. *Journal of Experimental Marine Biology and Ecology*, 161(2), 145-178.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., ... & van der Walt, S.J. (2020) SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods, 17(3), 261–272. https://doi.org/10.1038/s41592-020-0772-5

Ward, J.H. (1963) Hierarchical grouping to optimise an objective function. Journal of the American Statistical Association, 58(301), pp. 236–244.

Waskom, M.L. (2021) seaborn: statistical data visualization. *Journal of Open Source Software* 6(60): 3021. https://doi.org/10.21105/joss.03021

Wu, X., Koch, M. and Xu, Y. (2021) Unsupervised learning for groundwater level clustering and pattern recognition in data-sparse aquifer systems. Hydrogeology Journal, 29(4), pp. 1507–1528.

Young, P.C. (2011) Recursive Estimation and Time-Series Analysis: An Introduction for the Student and Practitioner. Berlin: Springer.

Yuan, C. and Yang, H. (2019) Research on K-value selection method of K-means clustering algorithm. J, 2(2), pp. 226–235.

Zanotti, C., Rotiroti, M., Fumagalli, L., Stefania, G.A., Canonaco, F., Prevot, A.S.H., Leoni, B., Nava, V. and Bonomi, T. (2022) Groundwater and surface water quality characterisation through positive matrix factorisation combined with GIS approach. Water Research, 218, 118520.

9. Acknowledgements

The author would like to acknowledge the **Centre for Ecology & Hydrology (CEH)** for providing the materials required for the initial installation of the dipwell network. This study was otherwise conducted as an independent research project.

10. Data Availability

10.1 . Data Availability and Software

The complete dipwell monitoring network, comprising 117 wells across Newborough Warren and the adjacent dune systems, is mapped at: https://www.google.com/maps/d/edit?mid=1hXLAauiMeaVsXhBR_IoUTziAtjk (M. Hollingham, unpublished).

The analytical pipeline underpinning this study is documented at https://newbroman.github.io/Newborough_Hydrology/ and comprises sixteen modules covering data preparation, state-space modelling, and intervention analysis and mapping. All analyses were implemented in Python using NumPy (Harris et al., 2020), pandas (McKinney, 2010), SciPy (Virtanen et al., 2020), statsmodels (Seabold and Perktold, 2010), scikit-learn (Pedregosa et al., 2011) and seaborn (Waskom, 2021), with figures produced in matplotlib (Hunter, 2007). Spatial analysis used GeoPandas (Jordahl et al., 2020), rasterio (Gillies, 2013), contextily (Arribas-Bel et al., 2020) and adjustText (Flyamer et al., 2020).

Full diagnostic outputs for all 66 reference network wells and all 16 clearfell BACI wells are reproducible by running the published pipeline. Groundwater monitoring data supporting this study are also available.

Source Data Credits:

- **Climate Data:** Meteorological records for RAF Valley were obtained from the UK Met Office (MIDAS Open Government Licence).
- **Topographic Data:** Terrain analysis was performed using 1-metre resolution LiDAR composite datasets provided by Natural Resources Wales (NRW) and accessed via the Welsh Government's Lle Geo-Portal.
- **Groundwater Data:** The primary dipwell monitoring record was independently maintained by the author.