

Tracked Changes

Hydrology and Earth System Sciences — Interactive Public Discussion

*Characterizing Climate-Driven Shifts in Chilean Rainfall Regimes
with a Hybrid Hidden Markov-Copula Framework*

Herrera-Marín, Kleisinger, Rivera, Wilson, Godoy-Faundez

Comparing: **Original submission** → **Revised manuscript**

Legend: **Blue bold = Added** *[DEL: Red strikethrough = Deleted]* Unchanged text in black. Shaded blue boxes = large added blocks; shaded red boxes = large deleted blocks.

Summary of Changes

The following table provides a structured overview of all substantive changes made between the original submission and the revised manuscript.

#	Location	Change
1	Abstract	Entirely rewritten: N_{obs}/p , CI for trend, CRPS skill scores, AUC, copula layer significance added.
2	§2 opening	New paragraph-map of the modelling pipeline added.
3	§2.1 (Data)	VIMF and VIMFD added to covariate set; physical justification stated.
4	§2.2 title	“Non-Homogeneous Hidden Markov Model (nHMM)” → “Non-Homogeneous Hidden Markov Model”.
5	§2.3 steps	Steps 1–4 now explicitly numbered and titled.
6	§2.3.1	Zone selection criterion (five-CVI consensus) stated at the start; Table 1 (CVI) moved to main text.
7	§2.3.2	GPD bootstrap CIs described; physical dry/wet tail contrast explained.
8	§2.3.4	Title changed; Table 3 (vine structure) added in-line.
9	§2.4 title	“Model validation: state dynamics...” → “Model validation strategy”.
10	§2.4.1–2.4.3	<i>Retention criterion</i> / <i>Decision criterion</i> paragraphs added to each subsection.
11	§2.4.2 (new para)	Explanation of KS/CvM behaviour at large n added.
12	§2.4.4	Title updated; truncated conceptual justification.
13	§2.5	Former §2.5 “Experimental Setup” (five subsections) removed . New §2.5 = model selection with Table 4 (N/p).
14	§3.1	Section title and content condensed.
15	§3.2	Two subsections merged into flowing text.
16	§3.3	CI for $S_4 \rightarrow S_4$ trend added; FDR procedure stated.
17	§3.4	Reference to supplement S7 added for rank robustness.
18	§3.5	Table 6 (AW/POC/AAD) with 95 % CI for POC added.
19	§3.5.1	Table 7 added: CRPSS, CRPS vine/indep, p_{DM} , pinball, AUC for all zones. Zones 2 and 4 now n/a with explanation.
20	§3.6	Section title simplified; content condensed.
21	§4.1	Detailed contributions incorporating new Table results.
22	§4.4	Former §4.4 “Robustness” → “Limitations and pathways for improvement”; Zone 4 arid error corrected.
23	§4.5 (removed)	Former §4.5 “Limitations and scope” merged into new §4.4.
24	§4.6 (removed)	Former §4.6 “Pathways for future work” merged into new §4.4.
25	§4.7 → §4.5	“Concluding remarks” renumbered; content updated.

1 Abstract

Original abstract (abridged): Chile’s hydroclimate exhibits pronounced meridional gradients... We introduce a hierarchical framework... The approach integrates: (i) a covariate-driven non-homogeneous Hidden Markov Model (nHMM) to [DEL: learn synoptic precipitation regimes and their transitions]; (ii) DTW clustering... (iii) state-conditional Regular Vine copulas... together with large-scale atmospheric covariates including [DEL: the Southern Oscillation Index (SOI), the Oceanic Niño Index (ONI), and Global Mean Sea-Surface Temperature (GMSST).]

Mixed-effects inference on the transition matrix reveals a statistically significant decline

in wet-state persistence of $\sim 0.34\% \text{ yr}^{-1}$ (*DEL: $\approx 14.5\%$ over 1980–2022*), coincident with rising GMSST. Out-of-sample ensembles for 2022 (100 daily members) *[DEL: are well calibrated: central 90 % prediction intervals achieve near-nominal coverage, low asymmetry, and widths increasing southward with climatological variance.]*

Revised abstract (key additions): The nHMM comprises $p = 29,284$ estimated parameters fit to $N_{\text{obs}} = 2,386,692$ daily observations ($N_{\text{obs}}/p \approx 82$), confirming adequate model identifiability. Mixed-effects inference reveals a significant decline in wet-state self-persistence of $\approx 0.34\% \text{ yr}^{-1}$ (95 % CI: $[0.21, 0.47]\% \text{ yr}^{-1}$; $\approx 14.5\%$ over 1980–2022).

The vine-copula layer yields statistically significant improvements over a conditionally independent baseline (Diebold–Mariano $p < 0.001$, HAC-adjusted) in the rainfall-active southern zones, where CRPS skill scores reach 0.51 (Zone 5) and 0.35 (Zone 6) relative to a climatological benchmark and event-based AUC exceeds 0.87 for the 25 mm exceedance threshold.

The covariate set now includes vertically integrated moisture flux and its divergence (VIMF/VIMFD) in addition to wind components, pressure, SST, and ENSO indices.

2 Section 2.1 — Data (covariate set)

Unchanged text about CR2MET and ERA5...

Covariate list:

[DEL: geopotential height at 500 hPa, mean sea-level pressure, sea surface temperature, evapotranspiration, maximum and mean near-surface temperature, and zonal–meridional wind components (u, v).]

geopotential height at 500 hPa (Z500); mean sea-level pressure (MSLP); sea surface temperature (SST); near-surface and maximum temperature; zonal and meridional wind components at 850 hPa (u_{850} , v_{850}); vertically integrated moisture flux and its divergence (VIMF, VIMFD); and the ENSO indices SOI, ONI, and GMSST.

The inclusion of VIMF and VIMFD is motivated by their direct link to land-falling moisture flux over Chile: u_{850} and v_{850} alone do not fully characterize atmospheric moisture delivery in the presence of orographic blocking and low-level jets.

3 Section 2.3 — Hybrid nHMM–Copula Framework

3.1 Section 2.3 opening paragraph

[DEL: “However, Assumption 3 (conditional spatial independence given S_t) may be too restrictive... We therefore propose a hierarchical framework in which the nHMM governs temporal regime dynamics...”]

“The hybrid framework addresses this through four sequential steps whose conceptual justification is elaborated in Sect. 2.4.4. Full mathematical details and extended diagnostics are in Supplement S1–S4.”

3.2 Section 2.3.1 — Zoning via DTW

[DEL: Title: “Zoning via DTW”]

Title: “Step 1 – Zoning via DTW”

New text added immediately after defining Ward’s linkage, before the mathematical formulation:

The optimal number of zones is $k^* = 7$, selected via a consensus of five internal cluster validity indices (CVI): Silhouette, Calinski–Harabasz, Davies–Bouldin, Dunn, and COP (Table 1; full table in Supplement S2). The $G = 7$ solution achieves the best consensus rank across all five CVIs, aligns with established climatic subregions of continental Chile, and yields zones with $d_z \in [50, 90]$ stations and at least 200 training days per state, providing sufficient data for robust copula estimation.

Table 1 (NEW in main text) — Top-ranked DTW zoning candidates.

Config.	Sil	CH	DB	Dunn	COP	Rank
ward.D2, $G = 7$	0.390	186.4	0.743	0.012	0.241	1
ward.D2, $G = 6$	0.393	201.6	0.810	0.012	0.271	2
ward.D, $G = 6$	0.401	212.9	0.787	0.016	0.259	3
ward.D2, $G = 5$	0.390	235.4	0.816	0.011	0.291	4
ward.D, $G = 5$	0.376	224.3	0.843	0.016	0.306	5

3.3 Section 2.3.2 — State-conditional marginals

New paragraph added after the GPD definition:

Table 2 summarizes GPD estimates for the Zone 5 medoid (southern Chile, storm-dominated) across all five states. GPD shape parameters exhibit a physically consistent pattern: states 0, 3, and 4 (anticyclonic to transitional) show $\hat{\xi} \leq 0$ (bounded or near-exponential tails), while states 1 and 2 (episodic frontal intrusions) yield $\hat{\xi} > 0$ (heavy, unbounded tails). This contrast reflects distinct precipitation generation mechanisms: light, intermittent precipitation under blocking conditions (bounded tails) versus intense, organized rainfall from deep frontal systems (heavy tails). Bootstrap 95 % confidence intervals ($B = 500$ parametric resamples of the medoid’s own exceedances) confirm that this sign contrast is statistically robust.

Table 2 (NEW in main text) — GPD estimates at Zone 5 medoid.

State	u^*	$\hat{\sigma}$	$\hat{\xi}$	CI_L	CI_U	n_{exc}	Tail
0	7.72	3.92	−0.018	−0.30	+0.17	113	bounded
1	0.12	0.11	+0.458	−0.20	+0.80	60	heavy
2	7.37	3.72	+0.121	−0.19	+0.40	77	heavy
3	43.80	19.09	−0.161	−0.37	+0.00	102	bounded
4	41.68	12.53	−0.112	−0.53	+0.06	86	bounded

3.4 Section 2.3.4 — Regular vine

New text after the simplifying assumption declaration:

Table 3 shows the selected structures for Zone 2 (semi-arid north) across all states. All are R-vines truncated at Tree 2, with Archimedean families predominating (7–23 % of edges) and Gaussian pairs rare (0–3 %), reflecting the non-Gaussian, asymmetric spatial dependence of daily precipitation in semi-arid terrain.

Table 3 (NEW in main text) — Vine structures, Zone 2.

State	Type	Trunc.	%Gauss	%Arch	n
0	R-vine	2	0.0	11.7	1179
1	R-vine	2	0.0	7.2	1269
2	R-vine	2	0.9	23.1	845
3	R-vine	2	0.0	16.2	1013
4	R-vine	2	2.7	16.2	860

4 Section 2.4 — Model Validation Strategy

4.1 Section 2.4.1 — nHMM fit (retention criterion added)

Added at end of §2.4.1:

Retention criterion: a configuration is retained if PIT histograms are visually flat, KS/CvM tests do not reject uniformity at the 5 % level for > 80 % of stations, and no significant Ljung–Box autocorrelation is detected. Regime count K is selected by BIC (Sect. 2.5), with PIT-based adequacy as a secondary check.

4.2 Section 2.4.2 — Vine-copula fit (decision criterion + KS/CvM explanation added)

Added paragraph on KS/CvM test power:

We note that KS and CvM tests on individual Rosenblatt components are extremely sensitive at the training sample sizes employed here ($n = 845$ – 1269 days per state): they will reliably detect even minor residual dependence that has negligible practical consequence for forecast calibration. Accordingly, the *operative criterion* for vine adequacy is not strict p -value thresholds on every component but rather: (i) PIT histograms visually consistent with uniformity for states and zones used in the verification; (ii) variogram scores lower for the vine model than for the independence baseline; and (iii) Diebold–Mariano tests confirming predictive skill gains in out-of-sample verification.

4.3 Section 2.4.3 — Out-of-sample verification (decision criterion added)

Added at end of §2.4.3:

Decision criterion: the vine layer is deemed to improve over the independence baseline if the Diebold–Mariano test rejects equality at the 10 % level (HAC-adjusted) for station-wise CRPS or zone-wise VS in the majority of

zones.

5 Section 2.5 — Model Selection (replaces “Experimental Setup”)

[DEL: Original §2.5 (removed entirely): “Experimental setup, model selection, and validation (overview)” with five subsections: §2.5.1 Grid of nHMM configurations; §2.5.2 DTW-based zoning for dependence modeling; §2.5.3 State-conditional vine copulas and simulation; §2.5.4 Forecasting and uncertainty quantification; §2.5.5 Rationale and coherence with the hierarchical design.]

New §2.5 — “Model selection: grid of nHMM configurations” — a single, compact subsection containing Table 4 with N_{obs}/p ratio.

Table 4 (NEW in main text) — Model selection front.

Config.	K	p	$\ell(\hat{\theta})$	AIC	BIC	N/p
4hs4dg	4	17,595	−5,286,428	10,608,046	10,733,017	136
5hs5dg	5	29,284	−5,108,541	10,275,650	10,483,644	82
6hs3dg	6	17,695	−5,006,775	10,048,940	10,174,621	135

$N_{\text{obs}}/p \approx 82$ confirms adequate identifiability: there are roughly 82 observations per estimated parameter ($N_{\text{obs}} = 5166 \times 462 = 2,386,692$).

6 Section 3.3 — Persistence and Transitions

[DEL: “To assess temporal drift, we analyze $\{p_{ij}^{(y)}\}$ with binomial GLMMs... Figure 3 shows vetted trends: $S_4 \rightarrow S_4$ persistence declines significantly...”]

“Analysis of annual transition proportions $\{p_{ij}^{(y)}\}_{y=1980}^{2022}$ via binomial GLMMs (beta-binomial under overdispersion), with Benjamini–Hochberg FDR correction over all $K^2 = 25$ transition pairs, reveals a statistically significant decline in $S_4 \rightarrow S_4$ self-persistence at $\approx 0.34\% \text{ yr}^{-1}$ (95 % CI: $[0.21, 0.47]\% \text{ yr}^{-1}$; $p < 0.01$), indicating shorter wet-regime residence over the study period. The $S_1 \rightarrow S_1$ self-persistence is trendless ($p = 0.43$)...”

7 Section 3.5 — Out-of-Sample Verification

7.1 Table 6 (NEW — interval metrics with 95 % CI for POC)

Table 6 (NEW in main text) — Uncertainty metrics with bootstrap 95 % CI for POC.

Zone	AW (mm)	POC (%)	POC 95 % CI	AAD
Zone 2	6.873	96.56	(93.2, 99.1)	0.070
Zone 3	18.117	97.67	(95.3, 99.5)	0.016
Zone 4	26.644	96.90	(94.1, 99.3)	0.016
Zone 5	32.521	97.45	(94.8, 99.4)	0.008
Zone 6	44.020	93.33	(89.4, 96.6)	0.012
Zone 7	11.560	88.83	(84.2, 93.0)	0.029

7.2 Section 3.5.1 — Proper scores (Table 7 and Zone 4 explanation)

[DEL: Original §3.5.1 title: “Proper scores and event-based skill”. No Table 7. No mention of Zones 2 or 4 anomalies. Zone 4 described elsewhere as “hyper-arid” (incorrect).]

Table 7 (NEW in main text) — Verification summary. CRPSS, CRPS vine/indep, p_{DM} , pinball loss, AUC. Zones 2 and 4: n/a for areal-CRPS (see explanation below).

Zone	CRPSS	CRPS _v	CRPS _i	p_{DM}	AUC
2	n/a	n/a	n/a	n/a	0.38
3	0.23	1.73	1.77	0.870	0.84
4	n/a	n/a	n/a	n/a	0.91
5	0.51	2.67	3.25	<0.001	0.89
6	0.35	3.47	4.71	<0.001	0.87
7	0.04	1.70	2.23	<0.001	n/a

New explanatory paragraph (Sect. 3.5.1):

For Zones 2 and 4, the areal-CRPS and CRPSS values are reported as not applicable (n/a). Zone 2 spans the Atacama and Coquimbo desert transition zone ($\approx 25^\circ\text{S}$ – 30°S); near-zero precipitation frequency causes the vine ensemble to occasionally generate positive draws on universally zero-observation days, inflating the areal mean and causing CRPS to diverge. Zone 4 covers the semi-arid Mediterranean transition of central Chile ($\approx 33^\circ\text{S}$ – 38°S , precordillera of O’Higgins, Maule, and northern Bío-Bío)— not hyper-arid— a region of high within-zone orographic heterogeneity. Rare large draws at high-elevation Andean stations (plausible given GPD shape $\hat{\xi} = 0.921$ for State 1 at the Zone 4 medoid, fitted with only $n_{\text{exc}} = 34$ exceedances) can dominate the areal mean even when lower stations record near-zero totals. In both zones, interval-based diagnostics in Table 6 (POC $\approx 97\%$, AAD ≤ 0.07) confirm well-calibrated performance at the medoid level.

8 Section 4.4 — Limitations (Zone 4 correction + restructuring)

[DEL: Original §4.4 title: “Robustness and validation.” Original §4.5 title: “Limitations and scope.” Original §4.6 title: “Pathways for future work.”]

[DEL: In §4.4: “... with expected challenges in the hyper-arid north (mass at zero,

data scarcity in the tail).”]

Revised §4.4 title: “Limitations and pathways for improvement.” Former §4.5 and §4.6 merged into this section.

Corrected text: “... with expected challenges in the desert transition zones (Zone 2) and in orographically heterogeneous transition zones (Zone 4) where zone-level areal-CRPS is not a valid evaluation metric.”

New recommendations for Zones 2 and 4: orographically weighted areal averages; dedicated zero-inflated marginals; stratified sub-domain validation.

Unchanged Sections

The following sections are either unchanged or contain only minor wording improvements (punctuation, hyphenation, sentence flow) with no substantive content differences:

- Section 1 (Introduction) — unchanged except last paragraph updated to reference new §2 structure.
- Section 2.2 (nHMM model structure, Assumptions 1–4) — unchanged.
- Section 2.3.3 (copula space transformation) — unchanged.
- Section 2.3.5 (simulation and back-transformation) — unchanged.
- Section 2.3.6 (uncertainty metrics) — minor rewording only.
- Section 3.2 (precipitation fields by state) — minor rewording.
- Section 3.4 (climate indices) — minor additions from Supplement S7.
- Section 3.5 (Figure 6 description) — unchanged.
- Section 4.2 (prior work and novelty) — unchanged.
- Section 4.3 (hydroclimatic implications) — unchanged.
- Code/data availability, author contributions, acknowledgements — unchanged.
- References — two new entries added (Bergmeir & Benítez 2012; already cited but now properly formatted).

Characterizing Climate-Driven Shifts in Chilean Rainfall Regimes with a Hybrid Hidden ~~Markov-Copula~~ Markov-Copula Framework

Mauricio Herrera-Marín¹, Francisca Kleisinger¹, Diego Rivera¹, Andrés Wilson¹, and Alex Godoy-Faundez¹

¹Research Center on Sustainability and Strategic Resource Management (CISGER), Faculty of Engineering, Universidad del Desarrollo (UDD), Avda. Plaza 680, Santiago, Chile

Correspondence: Mauricio Herrera-Marín (mherrera@udd.cl)

Abstract. Chile's hydroclimate exhibits pronounced meridional gradients and strong interannual variability, posing persistent challenges for regime-aware, probabilistic rainfall prediction. We introduce a hierarchical framework that explicitly separates large-scale regime dynamics from local spatial dependence. The approach integrates: (i) a covariate-driven non-homogeneous Hidden Markov Model (nHMM) to learn identify synoptic precipitation regimes and their climate-sensitive transitions; (ii) Dynamic Time Warping (DTW) clustering to delineate precipitation-coherent climatic zones; and (iii) state-conditional Regular Vine copulas with Generalized Pareto (GPD) tails to model residual spatial dependence and extremes. The analysis employs the 0.05° 0.05° daily CR2MET precipitation product over continental Chile (462 grid points, May–August 1980–2021 May–August 1980–2021) together with large-scale atmospheric covariates including the Southern Oscillation Index (SOI), the Oceanic Niño Index (ONIZonal and meridional wind components, geopotential height, sea-level pressure, sea surface temperature, vertically integrated moisture flux and its divergence (VIMF/VIMFD), and Global Mean Sea-Surface Temperature (standard ENSO indices (SOI, ONI, GMSST).

Five physically consistent rainfall regimes emerge, spanning from an anticyclonic dry state to a cyclonic wet state, confirmed by composites of mean sea-level pressure, 850-hPa winds, and 500-hPa geopotential height. The nHMM comprises $p = 29,284$ estimated parameters fit to $N_{\text{obs}} = 2,386,692$ daily observations ($N_{\text{obs}}/p \approx 82$), confirming adequate model identifiability. Mixed-effects inference on the transition matrix annual transition proportions reveals a statistically significant decline in wet-state persistence of $\sim 0.34\%$ self-persistence of $\approx 0.34\% \text{ yr}^{-1}$ ($\approx 14.5\%$ over 1980–2022 95 % CI: [0.21, 0.47] % yr^{-1} ; $\approx 14.5\%$ over 1980–2022), coincident with rising GMSST global mean sea surface temperature. Out-of-sample ensembles for MJJA 2022 (100 daily members conditioned on Viterbi Viterbi-decoded states) are well calibrated across all seven DTW zones: central 90 % prediction intervals achieve near-nominal coverage (88–98%), asymmetry degrees ≤ 0.07 , low asymmetry, and widths increasing that increase monotonically southward with climatological variance.

GPD shape parameters range from $\hat{\xi} < 0$ in anticyclonic dry states (bounded tails) to $\hat{\xi} > 0$ in episodic frontal states (heavy tails), confirming state-specific extreme behaviour. The vine-copula layer yields statistically significant improvements over a conditionally independent baseline (Diebold–Mariano $p < 0.001$, HAC-adjusted) in the rainfall-active southern zones, where CRPS skill scores reach 0.51 (Zone 5) and 0.35 (Zone 6) relative to a climatological benchmark and event-based AUC exceeds 0.87 for the 25 mm exceedance threshold. By disentangling regime timing and drivers from residual spatial co-variability and extremes, the proposed nHMM – DTW – vine – GPD nHMM–DTW–vine–GPD framework yields

meteorologically coherent states, spatially consistent probabilistic simulations, and quantitatively validated forecasts. The method is computationally tractable and transferable, offering a principled pathway for regime-conditioned, uncertainty-aware precipitation prediction to support hydroclimate risk management in Chile and other topographically complex regions.

1 Introduction

Understanding precipitation variability and its long-term modulation is a central challenge in hydroclimatology, particularly in regions of pronounced climatic heterogeneity such as Chile (Sarricolea et al., 2017). The country spans an extreme meridional gradient—from gradient—from hyperarid deserts in the north (18–26°S) to temperate and subpolar climates in the south (44–54°S)—that produces—producing sharp transitions in precipitation regimes and associated hydrological processes. Rainfall variability in Chile operates across multiple temporal scales, including seasonal (Montecinos and Aceituno, 2003; Garreaud et al., 2009), interannual (Garreaud et al., 2017; Montecinos et al., 2011)(Montecinos and Aceituno, 2003; Garreaud et al., 2009), interannual (Garreaud et al., 2017; Montecinos et al., 2011), and decadal to multidecadal fluctuations (Garreaud, 2009; Latoja et al., 2024)(Garreaud, 2009; Latoja et al., 2024). Among the large-scale climate drivers, the El Niño–Southern El Niño–Southern Oscillation (ENSO) dominates interannual variability, exerting strong control on regional circulation and precipitation anomalies. Warm (El Niño) and cold (La Niña) phases modulate storm-track position, jet strength, and moisture transport toward subtropical South America (Trenberth, 2014; Robertson et al., 2004a; Zhang et al., 2021)(Trenberth, 2014; Robertson et al., 2004a; Zhang et al., 2021). These effects are typically represented through indices such as the Southern Oscillation Index (SOI), the Oceanic Niño Niño Index (ONI), and the Global Mean Sea Surface Temperature (GMSST), which encapsulate large-scale forcing but often fail to explain the strong spatial gradients in rainfall induced by orography, mesoscale dynamics, and land–atmosphere coupling (Vrac et al., 2012; Bárdossy and Pegram, 2010)land–atmosphere coupling (Vrac et al., 2012; Bárdossy and Pegram, 2010).

The persistence of extreme drought conditions—particularly conditions—the decade-long “megadrought” “megadrought” affecting central Chile since 2010 (Garreaud et al., 2020; Bozkurt et al., 2021)—underscores (Garreaud et al., 2020; Bozkurt et al., 2021)—underscores the urgency of developing analytical frameworks that integrate both large-scale climate drivers and localized hydrological responses. Annual precipitation deficits between 15 % and 45 % relative to the pre-2010 baseline (Boisier et al., 2018), combined with regional warming and snowpack decline (Burger et al., 2018; Barría et al., 2019)(Burger et al., 2018; Barría et al., 2019), have exacerbated water scarcity and altered runoff regimes. These changes highlight the limitations of traditional statistical models, which often that assume linearity or stationarity in the relationship between atmospheric forcing and precipitation (Charles et al., 2014; Wilks, 2011)(Charles et al., 2014; Wilks, 2011). Capturing hidden or emergent shifts in rainfall regimes thus therefore requires probabilistic frameworks capable of representing temporal evolution, regime persistence, and non-linear teleconnections.

In this context, we employ a Non-Homogeneous Hidden Markov Model (nHMM) to identify latent rainfall regimes and characterize their temporal dynamics. Hidden Markov Models (HMMs) HMMs have been extensively used to represent for atmospheric circulation regimes and daily rainfall sequences (Ailliot et al., 2009; Franzke and Osprey, 2014). The (Ailliot et al., 2009; Franzke and Osprey, 2014); the non-homogeneous extension allows transition probabilities to depend explicitly on time-varying covariates such as ENSO indices or large-scale circulation fields (Hughes et al., 1999; Bellone et al., 2000; Robertson et al., 2004b), thus enabling the (Hughes et al.,

1999; Bellone et al., 2000; Robertson et al., 2004b), enabling detection of evolving regime behavior under changing climatic conditions. Beyond identifying discrete rainfall states, we assess their long-term persistence and transition probabilities using linear mixed models, quantifying statistically significant trends that reflect climate-induced shifts in regime stability.

However, a key limitation of conventional nHMM applications lies in their assumption of conditional spatial independence—i.e., that independence—that rainfall at different locations is independent once the hidden state is known. This simplification is rarely valid in complex terrains like terrain such as Chile, where strong orographic forcing and mesoscale circulations generate residual spatial dependence even after conditioning on large-scale atmospheric states. To address this issue, we develop a hybrid nHMM–Copula framework that combines the temporal structure and climatic interpretability of nHMMs with explicit spatial modeling via multivariate copulas. The approach begins by partitioning Chile Chile is first partitioned into homogeneous precipitation zones using Dynamic Time Warping (DTW) clustering (Berndt and Clifford, 1994; Herrera et al., 2025), ensuring that grouped stations share coherent rainfall dynamics (Maharaj and Alonso, 2011; Montero and Vilar, 2015)(Berndt and Clifford, 1994; Herrera et al., 2025). Within each zone, we estimate state-specific copula models to reconstruct spatial dependence, thereby relaxing the restrictive independence assumption while preserving the probabilistic and physically interpretable structure of the HMM (Bárdossy and Pegram, 2010; Vrac et al., 2012)(Bárdossy and Pegram, 2010; Vrac et al., 2012).

Previous studies have independently applied HMMs and copula models to precipitation (Ailliot et al., 2009; Vrac et al., 2012)(Ailliot et al., 2009; Vrac et al., 2012), yet few have integrated spatial dependence within a covariate-driven nHMM framework. The novelty of this study lies in bridging these two perspectives through a hybrid regime-based and , dependence-aware modeling system . By explicitly linking modelling system that explicitly links hidden rainfall regimes to large-scale atmospheric drivers while reconstructing state-conditioned spatial dependencies, our framework enables spatially coherent probabilistic simulations of precipitation fields that are consistent with both physical mechanisms and statistical constraints. This approach provides providing an interpretable, transferable, and uncertainty-aware basis for rainfall regime analysis and probabilistic forecasting across Chile's diverse hydroclimatic zones.

The remainder of this paper is organized as follows. Section 2 describes the data and methods, including the CR2MET precipitation dataset, the nHMM formulation, the DTW-based climatic zoning, and the copula construction, the step-by-step model construction, and the validation approach. Section 3 presents the resulting rainfall regimes, their atmospheric interpretation, temporal trends, and spatially coherent probabilistic simulations. Finally, Section 4 discusses the methodological advances, implications for hydroclimatic predictability, hydroclimatic implications, limitations, and potential extensions for coupled hydroclimate–impact modeling.

2 Data and Methods

This section guides the reader sequentially through the modelling pipeline. Section 2.1 describes the precipitation product and atmospheric covariates. Section 2.2 introduces the nHMM. Section 2.3 builds the hybrid spatial framework step by step: DTW zoning (Sect. 2.3.1), marginal specification (Sect. 2.3.2), copula construction (Sects. 2.3.3–2.3.4), simulation (Sect. 2.3.5), and uncertainty metrics (Sect. 2.3.6). Section 2.4 details the validation strategy. Section 2.5 summarizes model selection.

2.1 Data

Daily precipitation in across Chile is monitored by official agencies (CR2, 2024), yet many ground stations exhibit discontinuities, missing data, or record lengths records shorter than 30 years, limiting their suitability for robust climate diagnostics. To ensure spatial completeness and temporal consistency, we employed the CR2MET CR2MET dataset (Boisier, 2023; Boisier et al., 2018), a high-resolution (0.05°) gridded product developed specifically for continental Chile and rigorously validated against dense observational networks. CR2MET integrates local station observations with global reanalyses—primarily ERA5 (Hersbach et al., 2023)—using statistical downscaling techniques to better capture reanalysis (Hersbach et al., 2023) via statistical downscaling, capturing precipitation over complex orography. While ERA5 provides a physically consistent global depiction of atmospheric conditions, the regional refinement in CR2MET significantly improves the realism of daily rainfall estimates across Chile, particularly at the spatial scales relevant for hydrological and climate-impact studies (Mayer et al., 2021; Cavalleri et al., 2024) more realistically than global products (Mayer et al., 2021; Cavalleri et al., 2024).

From CR2MET, we extracted 462 grid points collocated with or proximal to operational rain gauges exhibiting with uninterrupted daily records during 1980–2021. 1980–2021. Each grid point contributes $123 \times 42 = 5166$ daily winter observations, where the austral winter season (May–August) austral-winter (May–August; MJJA) was isolated for analysis. This period was selected because observations; frontal systems during MJJA account for approximately 70–90% of annual rainfall between 30°S and 44°S (Pica-Téllez et al., 2020), dominating streamflow generation, groundwater recharge, and snow accumulation in the Andes, and therefore conditioning water availability during spring and summer (Valdés-Pineda et al., 2016).

The fine spatial resolution (0.05° , $\sim 5\text{ km}$) ensures that the 462 selected grid cells are distinct and adequately resolve subregional precipitation variability. This granularity is crucial in a country characterized by strong climatic and orographic gradients, where coarse global reanalyses often underestimate variability and spatial dependence. In contrast, CR2MET captures both the magnitude and spatial organization of precipitation across coastal, valley, and mountain sectors, providing a robust basis for identifying rainfall regimes and their extremes.

Climatological controls on winter precipitation are consistent with established large-scale dynamics over the southeastern Pacific. Andean snow accumulation (Valdés-Pineda et al., 2016). The year 2022 is reserved for out-of-sample evaluation. Seasonal totals are modulated primarily primarily modulated by the South Pacific Subtropical High and the Southern Hemisphere storm track (Hoskins and Hodges, 2005; Barrett and Hameed, 2017). Orographic; orographic enhancement along the Andes intensifies rainfall from frontal systems (Garreaud et al., 2009; Arias et al., 2021), while persistent blocking or reduced frontal incursions can produce substantial deficits (Latoja et al., 2024), even during ENSO-neutral conditions (Montecinos et al., 2011). These processes are accurately represented in CR2MET, which thus provides a consistent observational foundation for both the temporal regime identification via the nHMM and the spatial dependence analysis through copula modeling.

(Latoja et al., 2024; Montecinos et al., 2011).

To represent large-scale atmospheric variability influencing daily precipitation, we incorporated multiple climatic covariates into Large-scale atmospheric covariates for the nHMM transition matrix, including model were extracted from ERA5 (Hersbach et al., 2023) at $0.25^\circ \times 0.25^\circ$ over continental Chile and adjacent oceanic domains (15°N – 60°S , 120°W – 60°W). The covariate set comprises geopotential height at 500 hPa, hPa (Z500); mean sea-level pressure, (MSLP); sea surface temperature, evapotranspiration, maximum and mean (SST); near-surface temperature, and zonal–and maximum temperature; zonal and meridional wind components (u, v). These fields were extracted from ERA5 reanalysis (Hersbach et al., 2023) at $0.25^\circ \times 0.25^\circ$ resolution, covering both continental Chile and adjacent oceanic regions (approximately 15°N – 60°S , 120°W – 60°W) at 850 hPa (u_{850} , v_{850}); vertically integrated moisture flux and its divergence (VIMF, VIMFD); and the ENSO indices SOI, ONI, and GMSST. The inclusion of oceanic domains is essential given the strong influence of the South Pacific High, storm-track dynamics, and sea–atmosphere interactions on winter precipitation over central Chile. To reduce dimensionality and multicollinearity among predictors, we applied VIMF and VIMFD is

motivated by their direct link to landfalling moisture flux over Chile (Garreaud et al., 2009; Arias et al., 2021): u_{850} and v_{850} alone do not fully characterize atmospheric moisture delivery in the presence of orographic blocking and low-level jets. Principal Component Analysis (PCA) was applied independently to each gridded covariate field and retained the field; leading components explaining more than 80 % of total variance. These principal components were then used as exogenous predictors in the nHMM transition model.

2.2 Non-Homogeneous Hidden Markov Model (nHMM)

The Non-Homogeneous Hidden Markov Model (nHMM) (Khalil et al., 2010; Kirshner, 2005; Cioffi et al., 2020; Hughes et al., 1999; Robertson et al., 2004b) provides a probabilistic framework to identify latent (hidden) nHMM (Khalil et al., 2010; Kirshner, 2005; Cioffi et al., 2020; Hughes et al., 1999; Robertson et al., 2004b) identifies latent climate regimes and model models their temporal evolution in daily rainfall sequences. The model represents precipitation as a stochastic process governed by unobserved states that capture large-scale atmospheric configurations. Transitions between these states follow the Markov property, where the probability of a state on a given day depends solely on the state of the preceding day. The nHMM extends the conventional (homogeneous) HMM by allowing transition probabilities to vary in time through dependence on external covariates, thereby linking local precipitation to evolving large-scale climatic drivers.

2.2.1 Model structure.

Let $R_{t_i}^{r_j}$ denote the daily precipitation at location $r_j = (\text{long}_j, \text{lat}_j)$, where long and lat are longitude and latitude, respectively, and let $r_j = (\text{long}_j, \text{lat}_j)$, $i = 1, \dots, n$ and $j = 1, \dots, q$. In this study, $n = 5166$ corresponds to 42 winters (1980–2021) with $42 \text{ winters} \times 123 \text{ days each (May–August; MJJA), MJJA days, 1980–2021}$ and $q = 462$ spatial locations. The year 2022 is reserved for model evaluation. We define $R_{t_i} = (R_{t_i}^{r_1}, \dots, R_{t_i}^{r_q})$ as the q -dimensional precipitation vector at time on day t_i , and $R_{1:n} = (R_{t_1}, \dots, R_{t_n})$ as the sequence of daily observations. The temporal dependence is governed by a latent (unobserved) Markov chain $S_t = (S_1, \dots, S_n)$ taking values in $\{1, \dots, K\}$, where K is the optimal number of hidden states estimated from data is $R_{t_i} = (R_{t_i}^{r_1}, \dots, R_{t_i}^{r_q})$, and the latent Markov chain $S_t \in \{1, \dots, K\}$ is defined by four assumptions.

The nHMM defines a joint probability distribution over the observed rainfall sequence $R_{1:n}$ and the **Assumption 1** (Markovian dynamics): the hidden state sequence $S_{1:n}$ through a set of conditional independence assumptions:

Assumption 1 (Markovian dynamics). The hidden state sequence $S_{1:n}$ is a first-order Markov process:

$$P(S_1, S_2, \dots, S_n) = p(S_1) \prod_{t=2}^n p(S_t | S_{t-1}),$$

where $p(S_1)$ is the initial state probability and $p(S_t | S_{t-1})$ defines the transition matrix between states, $P(S_1, \dots, S_n) = p(S_1) \prod_{t=2}^n p(S_t | S_{t-1})$.

Assumption 2 (Conditional dependence of rainfall on the hidden state). The precipitation vector on day t , R_t , depends only on the hidden state on that day:

$$p(R_{1:n}, S_{1:n}) = p(S_1) \prod_{t=2}^n p(S_t | S_{t-1}) \prod_{t=1}^n p(R_t | S_t).$$

(Conditional rainfall dependence): $p(R_{1:n}, S_{1:n}) = p(S_1) \prod_{t=2}^n p(S_t | S_{t-1}) \prod_{t=1}^n p(R_t | S_t)$.

Assumption 3(Conditional spatial independence).Conditional on the hidden state (Conditional spatial independence): conditional on S_t , precipitation across stations is assumed independent:

$$p(R_t | S_t) = \prod_{m=1}^q p(R_t^m | S_t).$$

Each term $p(R_t^m | S_t = i)$ is modeled as a mixed distribution that accounts for dry and wet days: , $p(R_t | S_t) = \prod_{m=1}^q p(R_t^m | S_t)$. This assumption is relaxed in Sect. 2.3 via vine copulas. Each marginal term follows a zero-inflated Gamma mixture,

$$p(R_t^m | S_t = i) = \prod_{m=1}^q \delta_{im}, \begin{cases} p_{im0}, & R_t^m = 0, \\ \sum_{c=1}^{C-1} \frac{p_{im1} \phi_{imc}^{\omega_{imc}} (R_t^m)^{\omega_{imc}-1} e^{-\phi_{imc} R_t^m}}{\Gamma(\omega_{imc})}, & R_t^m > 0, \end{cases}$$

with

$$\delta_{im} = \begin{cases} p_{im0}, & r_t^m = 0, \\ \sum_{c=1}^{C-1} \frac{p_{im1} \phi_{imc}^{\omega_{imc}} (r_t^m)^{\omega_{imc}-1} e^{-\phi_{imc} r_t^m}}{\Gamma(\omega_{imc})}, & r_t^m > 0, \end{cases}$$

where p_{im0} is the probability of no precipitation for state i at station m dry-day probability, $p_{im1} = 1 - p_{im0}$ is the complementary probability of rainfall, and $(\omega_{imc}, \phi_{imc})$ are the , and $(\omega_{imc}, \phi_{imc})$ are shape and rate parameters of the Gamma distribution for the c -th mixture component ($C - 1$ components in total). This zero-inflated mixture model flexibly captures both dry-day frequency and rainfall intensity across stations.

Gamma component.

Assumption 4(Non-homogeneity and climatic covariates).To account for nonstationary behavior, transition probabilities between hidden states are allowed to vary as a function of exogenous atmospheric variables. Let $X_t = (X_t^1, X_t^2, \dots, X_t^h)$ denote a vector of h climatic covariates at time t (e.g., geopotential height, mean sea-level pressure, wind components, temperature, and SST indices). The non-homogeneous formulation links large-scale atmospheric dynamics to local rainfall through:

$$p(R_t | S_{1:n}, R_{t-1}, X_{1:n}) = p(R_t | S_t) p(S_t | S_{t-1}, X_t),$$

where the second term, $p(S_t | S_{t-1}, X_t)$, introduces time-varying transition probabilities modeled via a (Non-homogeneity): transition probabilities depend on the atmospheric covariate vector $X_t = (X_t^1, \dots, X_t^h)$ via multinomial logistic regression:

$$p(S_t = i | S_{t-1} = j, X_t = x) = \frac{\exp(\alpha_{j,i} + \beta_i^T x)}{\sum_{k=1}^K \exp(\alpha_{j,k} + \beta_k^T x)}.$$

, ,

$$p(S_t = i | S_{t-1} = j, X_t = x) = \frac{\exp(\alpha_{j,i} + \beta_i^T x)}{\sum_{k=1}^K \exp(\alpha_{j,k} + \beta_k^T x)},$$

Here, where $\alpha_{j,i}$ represents is the baseline log-odds of transitioning from state j to state i , while transition $j \rightarrow i$ and β_i is a vector of regression coefficients quantifying the influence of the covariates X_t on that transition. This formulation enables the model to capture climate-driven shifts in the frequency and persistence of rainfall regimes.

quantifies covariate influence.

2.2.2 Estimation and inference.

Direct computation of the full joint likelihood over all possible hidden state sequences is computationally infeasible. Instead, parameter estimation proceeds via the forward–backward (Baum–Welch) algorithm, a variant of the Expectation–Maximization (EM) procedure (McLachlan and Krishnan, 2008). This iterative algorithm alternates between (i) estimating the posterior probabilities of hidden states given current parameters and data (E-step), and (ii) updating parameters to maximize the expected complete-data log-likelihood (M-step). We
 185 employ the implementation proposed by Kirshner (2005) and further Parameters are estimated via the Baum–Welch (EM) algorithm (McLachlan and Krishnan, 2008), following Kirshner (2005) as adapted by Robertson et al. (2006), which efficiently handles covariate-dependent transitions in non-homogeneous HMMs. . Model selection over $K \in \{2, \dots, 6\}$ and $C \in \{2, \dots, 8\}$ uses AIC and BIC (Sect. 2.5).

Overall, the nHMM provides a flexible probabilistic representation of rainfall regimes, capturing both the stochastic nature of daily precipitation and the modulation of regime transitions by evolving large-scale atmospheric drivers.

190 2.3 Hybrid nHMM–Copula Framework with DTW–Based DTW-Based Zoning

The non-homogeneous Hidden Markov Model (nHMM) described in Sect. 2.2 nHMM provides a coherent mechanism to extract hidden rainfall regimes $\{S_t\}$ and link their transitions to large-scale covariates $\{X_t\}$. However, temporal-regime layer, but Assumption 3 (conditional spatial independence given S_t) may be is often too restrictive in regions with strong orographic gradients and mesoscale circulations such as Chile, where residual spatial dependence often persists even after conditioning on the regime (Becker and Schmitz, 2003; Garreaud, 2009; Pablo Sarricolea and Óliver Meseguer-Ruiz, 2017). We therefore
 195 propose a hierarchical framework in which the nHMM governs temporal regime dynamics, while state-conditioned spatial dependence is reconstructed via multivariate copulas within climatically homogeneous zones.(Becker and Schmitz, 2003; Garreaud, 2009; Pablo Sarricolea and Óliver Meseguer-Ruiz, 2017). The hybrid framework addresses this through four sequential steps whose conceptual justification is elaborated in Sect. 2.4.4. Full mathematical details and extended diagnostics are in Supplement S1–S4.

2.3.1 Step 1 – Zoning via DTW

200 Let $R_t = (R_t^{r1}, \dots, R_t^{r_q})$ denote the daily precipitation vector at $q = 462$ locations and $S_t \in \{1, \dots, K\}$ the nHMM regime at time t . A direct q A direct $q = 462$ -variate copula is computationally prohibitive. We partition the network into $k^* = 7$ zones $\mathcal{Z} = \{1, \dots, 7\}$, each with d_z stations, using station network into k^* zones via hierarchical clustering on pairwise Dynamic Time Warping (DTW) dissimilarities computed from standardized MJJA daily series (Berndt and Clifford, 1994; Giorgino, 2009; Montero and Vilar, 2015).

For two stations i and j , DTW defines the alignment cost (Berndt and Clifford, 1994; Giorgino, 2009; Montero and Vilar, 2015):

$$205 \quad \text{DTW}(i, j) = \min_{\pi \in \Pi_{ij}} \sum_{(\tau, \kappa) \in \pi} \|y_i(\tau) - y_j(\kappa)\|,$$

where Π_{ij} is the set of monotone warping paths mapping the time indices of y_i to those of y_j . Ward'. Ward's linkage (Ward, 1963) on the resulting distance matrix yields compact, interpretable clusters. The final partition (Fig. 6A)

The optimal number of zones is $k^* = 7$, selected via a consensus of five internal cluster validity indices (CVI): Silhouette, Calinski–Harabasz, Davies–Bouldin, Dunn, and COP (Table 1; full table in Supplement S2). The $G = 7$ solution achieves
 210 the best consensus rank across all five CVIs, aligns with established climatic subregions (Pablo Sarricolea and Óliver Meseguer-Ruiz, 2017; Garreaud et al., 2009) and constrains copula modeling to tractable dimensions $d_z \in [50, 90]$ of continental Chile (Sarricolea et al., 2017; Garreaud et al., 2009), and yields zones with $d_z \in [50, 90]$ stations and at least 200 training days per state, providing sufficient data for robust copula estimation.

Table 1. Top-ranked DTW zoning candidates (consensus CVI rank). Sil: Silhouette; CH: Calinski–Harabasz ($\times 10^{-1}$); DB: Davies–Bouldin; COP: COP index. $\uparrow$: larger is better; $\downarrow$: smaller is better. **Bold**: selected configuration.

Config.	Sil $\uparrow$	CH $\uparrow$	DB $\downarrow$	Dunn $\uparrow$	COP $\downarrow$	Rank
ward.D2, $G = 7$	0.390	186.4	0.743	0.012	0.241	1
ward.D2, $G = 6$	0.393	201.6	0.810	0.012	0.271	2
ward.D, $G = 6$	0.401	212.9	0.787	0.016	0.259	3
ward.D2, $G = 5$	0.390	235.4	0.816	0.011	0.291	4
ward.D, $G = 5$	0.376	224.3	0.843	0.016	0.306	5

2.3.2 State-conditional Step 2 – State-conditional marginal specification

215 Fix a zone $z \in \mathcal{Z}$ and a regime $s \in \{1, \dots, K\}$. For station $m \in z$, let

$$F_{m,s}(x) = \mathbb{P}(R_t^m \leq x \mid S_t = s), \quad f_{m,s}(x) = \frac{d}{dx} F_{m,s}(x),$$

be For station m in zone z and regime s , the state-conditioned CDF and density, $F_{m,s}(x) = \mathbb{P}(R_t^m \leq x \mid S_t = s)$ is estimated from the subset $\{R_t^m : S_t = s\}$.¹ To respect the nHMM emissions and explicitly address extremes, we use a Viterbi-decoded training subset using a three-part hybrid marginal: A (1) a point mass at zero with probability $p_{m,s}^{(0)} = \mathbb{P}(R_t^m = 0 \mid S_t = s)$. A $p_{m,s}^{(0)} = \mathbb{P}(R_t^m = 0 \mid S_t = s)$;

220 (2) a $(C - 1)$ -component Gamma mixture for positive amounts, with weights $w_{m,s,c}$, shapes $\omega_{m,s,c}$ and rates $\phi_{m,s,c}$, $c = 1, \dots, C - 1$:

$$g_{m,s}(x) = \sum_{c=1}^{C-1} w_{m,s,c} \frac{\phi_{m,s,c}^{\omega_{m,s,c}} x^{\omega_{m,s,c}-1} e^{-\phi_{m,s,c}x}}{\Gamma(\omega_{m,s,c})}, \quad x > 0.$$

A; and (3) a Generalized Pareto Distribution (GPD) for the upper tail, fitted to exceedances $Y = x - u_{m,s}$ above the above threshold $u_{m,s}$ -quantile ($u_{m,s}$ chosen at the empirical (0.90) level), with shape $\xi_{m,s}$ and scale $\sigma_{m,s} > 0$: empirical quantile of positive precipitation):

$$\mathbb{P}(Y \leq y) = 1 - 1 - \left(1 + \frac{\xi_{m,s} y}{\sigma_{m,s}} \frac{\xi_{m,s} y}{\sigma_{m,s}} \right)^{-1/\xi_{m,s} - 1/\xi_{m,s}}, \quad y = x - u_{m,s} \geq 0.$$

225 We enforce continuity Continuity of $F_{m,s}$ at the threshold $u_{m,s}$ by stitching the body and tail with a is enforced by probability-splicing construction:

$$F_{m,s}(x) = \begin{cases} p_{m,s}^{(0)} + (1 - p_{m,s}^{(0)}) \int_0^x g_{m,s}(y) dy, & 0 \leq x \leq u_{m,s}, \\ F_{m,s}(u_{m,s}) + (1 - F_{m,s}(u_{m,s})) G_{m,s}^{\text{GPD}}(x - u_{m,s}), & x > u_{m,s}, \end{cases}$$

where $G_{m,s}^{\text{GPD}}$ is the GPD CDF. This hybrid specification (Hughes et al., 1999; Khalil et al., 2010; Zhang et al., 2021; Bárdossy and Pegram, 2010; Vrac et al., 2012) preserves the nHMM emission structure for the bulk while representing state-specific heavy tails relevant to hydroclimatic risk in central Chile (Garreaud et al., 2020; Salazar et al., 2024). Tail parameter stability is checked via threshold sensitivity (0.85–0.95 quantiles) and (Hughes et al., 1999; Khalil et al., 2010; Zhang et al., 2021; Bárdossy and Pegram, 2010; Vrac et al., 2012).

¹ In practice we use the Viterbi-decoded states for simulation; Section 2.4 reports robustness under smoothed state probabilities.

Table 2 summarizes GPD estimates for the Zone 5 medoid (southern Chile, storm-dominated) across all five states. GPD shape parameters exhibit a physically consistent pattern: states 0, 3, and 4 (anticyclonic to transitional) show $\hat{\xi} \leq 0$ (bounded or near-exponential tails), while states 1 and 2 (episodic frontal intrusions) yield $\hat{\xi} > 0$ (heavy, unbounded tails). This contrast reflects the distinct precipitation generation mechanisms: light, intermittent precipitation associated with cold-air incursions under blocking conditions (bounded tails) versus intense, organized rainfall from deep frontal systems (heavy tails). Bootstrap 95 % confidence intervals ($B = 500$ parametric resamples of the medoid's own exceedances) confirm that this sign contrast is statistically robust across the range. Threshold sensitivity at $u_{m,s} \in [0.85, 0.95]$ was verified via mean-excess diagnostics (see Sect. 2.4) Supplement S3), confirming adequate GPD stability at the adopted 0.90 level.

Table 2. GPD parameter estimates at the Zone 5 medoid by nHMM state (southern Chile, storm-dominated). u^* : 0.90-quantile threshold (mm); $\hat{\sigma}$: scale; $\hat{\xi}$: shape; 95 % bootstrap CI from $B = 500$ parametric resamples of medoid exceedances; n_{exc} : number of exceedances above u^* .

State	u^*	$\hat{\sigma}$	$\hat{\xi}$	CI_L	CI_U	n_{exc}	Tail
0	7.72	3.92	-0.018	-0.30	+0.17	113	bounded
1	0.12	0.11	+0.458	-0.20	+0.80	60	heavy
2	7.37	3.72	+0.121	-0.19	+0.40	77	heavy
3	43.80	19.09	-0.161	-0.37	+0.00	102	bounded
4	41.68	12.53	-0.112	-0.53	+0.06	86	bounded

2.3.3 Step 3 – Transformation to copula space

Define State-conditional PIT pseudo-observations (state-conditional PITs) for zone z and regime s :

$$U_t^{(m)} = F_{m,s}(R_t^m), \quad m = 1, \dots, d_z,$$

and $\mathbf{U}_t^z = (U_t^{(1)}, \dots, U_t^{(d_z)}) \in [0, 1]^{d_z}$ are defined as $U_t^{(m)} = F_{m,s}(R_t^m)$, $m = 1, \dots, d_z$. By Sklar's theorem (Nelsen, 2006; Genest and MacKay, 1986)'s theorem (Nelsen, 2006; Genest and MacKay, 1986), the conditional joint law of $\mathbf{R}_t^z$ factorizes as

$$F_{z,s}(\mathbf{r}) = C_{z,s}(F_{1,s}(r_1), \dots, F_{d_z,s}(r_{d_z})), f_{z,s}(\mathbf{r}) = c_{z,s}(\mathbf{u}) \prod_{m=1}^{d_z} f_{m,s}(r_m),$$

where $C_{z,s}$ is the copula of $(R_t^{(1)}, \dots, R_t^{(d_z)})$ given $S_t = s$ and $c_{z,s}$ its density. Thus, marginal behavior so marginal behaviour is inherited from Sect. 2.3.2 Step 2 and residual spatial dependence is entirely encoded by $c_{z,s}$ encoded entirely by the copula $C_{z,s}$.

2.3.4 Step 4 – Regular vine construction ,simplifying assumption, and model selection

For d_z up to ~ 90 , we represent $c_{z,s} \sim 90$ we represent $C_{z,s}$ by a Regular Vine (R-vine) (Aas et al., 2009), decomposing the d_z -
250 variate density into bivariate pair-copulas along a sequence of linked trees: :

$$c_{z,s}(u_1, \dots, u_{d_z}) = \prod_{m=1}^{d_z-1} \prod_{e \in E_m} c_e(u_{i|D_i|D}, u_{j|D_j|D}; \theta_e),$$

where E_m denotes edges at tree level m , $c_e(\cdot, \cdot; \theta_e)$ is a bivariate copula family with parameter θ_e , and $u_{i|D}$ denotes conditional distribution functions given the conditioning
set D determined by the vine structure. We adopt the simplifying assumption (pair-copula parameters do not depend on the values of conditioning
variables), standard in high-dimensional applications and empirically adequate in our context (independent of conditioning variable values); diagnostics
255 in Sect. 2.4).

We allow a 2.4.2 confirm its empirical adequacy. The candidate family set $\mathcal{F} = \{\text{Gaussian, Student-}t, \text{Clayton, Gumbel, Frank}\}$ is $\mathcal{F} = \{\text{Gaussian, Student-}t, \text{Clayton, Gumbel, Frank}\}$
covering elliptical and Archimedean types families with different tail properties. Structure selection and family assignment are
performed use AIC edge-wise using likelihood-based criteria (AIC/BIC) as implemented by `RVineStructureSelect` in the `VineCopula` R package (Schepsmeier
et al., 2015). Parameters via `RVineStructureSelect` (Schepsmeier et al., 2015); parameters are estimated by maximum like-
260 lihood . Multiple random initializations and, when beneficial, truncated vines are considered to balance fit and parsimony. In practice, lower-level trees capture the strongest
with multiple random initialisations. Vines are truncated at Tree 2, which captures the dominant spatial links (e.g., typically
between nearby stations) , where Student- t or asymmetric Archimedean families may be selected; higher trees often favor Gaussian pairs due to attenuation of conditional
dependence beyond $\sim$ while limiting over-fitting in higher-order trees where conditional dependence attenuates.

Table 3 shows the selected structures for Zone 2 (semi-arid north) across all states. All are R-vines truncated at
265 Tree 2, with Archimedean families predominating (7–23 % of edges) and Gaussian pairs rare (0–3 %), reflecting the non-
Gaussian, asymmetric spatial dependence structure of daily precipitation in semi-arid terrain. As expected, the proportion
of Gaussian pairs increases southward (reaching $\sim 22\%$ in wet states of Zone 5–10 variables; Supplement S4), consistent with
Bárdossy and Pegram (2010) more elliptically distributed spatial fields under persistent frontal forcing.

Table 3. Selected vine structures for Zone 2, all five states. %Arch: percentage of Archimedean pair-copula families selected by AIC; n :
training days.

State	Type	Trunc.	%Gauss	%Arch	n
0	R-vine	2	0.0	11.7	1179
1	R-vine	2	0.0	7.2	1269
2	R-vine	2	0.9	23.1	845
3	R-vine	2	0.0	16.2	1013
4	R-vine	2	2.7	16.2	860

2.3.5 State-aware simulation and back-transformation

270 Given a Viterbi-decoded regime $\hat{S}_t$ for day t and zone z , we generate an ensemble member as follows: Draw $U_t^z \sim C_{z, \hat{S}_t}$ using RVineSim. Map :
 (1) draw $U_t^z \sim C_{z, \hat{S}_t}$ via RVineSim; (2) map back to precipitation using the stitched marginal inverse:

$$R_t^{(m)} = \begin{cases} F_{m, \hat{S}_t}^{-1}(U_t^{(m)}), & U_t^{(m)} \leq u_{m, \hat{S}_t}^*, \\ F_{m, \hat{S}_t}^{-1}(u_{m, \hat{S}_t}^*) + Q_{\text{GPD}, m}^{-1}\left(\frac{U_t^{(m)} - u_{m, \hat{S}_t}^*}{1 - u_{m, \hat{S}_t}^*}\right), & U_t^{(m)} > u_{m, \hat{S}_t}^*, \end{cases}$$

where $u_{m, \hat{S}_t}^*$ is, directing draws above the tail threshold quantile and $Q_{\text{GPD}, m}^{-1}$ through the GPD quantile function for station m . This preserves both central and tail behavior (Trenberth, 2014; Franzke and Osprey, 2014).

275 2.3.6 Computational aspects and reproducibility

(Trenberth, 2014; Franzke and Osprey, 2014). The full system entails 7 zones $\times$ $K=5$ regimes = 35 copula models $\{C_{z, s}\}$, each with state-conditioned marginals $\{F_{m, s}\}$ for d_z stations. Estimation and simulation are comprises $7 \times 5 = 35$ copula models, parallelized across (z, s) using parallel and doParallel in R. Vines are fitted with VineCopula (Schepsmeier et al., 2015); clustering uses TSclust (Montero and Vilar, 2015); GPD fitting uses fExtremes and evd. The training set comprises $n = 5166$ MJJA days (1980–2021) and the 2022 season is reserved for out-of-sample evaluation.

280 VineCopula (Schepsmeier et al., 2015) and TSclust (Montero and Vilar, 2015).

2.3.6 Uncertainty quantification and summary metrics

From M ensemble members, define for each day t the central 90% prediction interval $[\hat{L}_t, \hat{U}_t]$ (empirical 5th and 95th percentiles) and the ensemble mean $\hat{X}_t$ per day, define the central 90 % prediction interval $[\hat{L}_t, \hat{U}_t]$ (5th–95th percentiles). Let X_t^{obs} be the observation. We report:

$$\begin{aligned} \text{AW} &:= \frac{1}{T} \sum_{t=1}^T (\hat{U}_t - \hat{L}_t), \text{ (Average Width; sharpness)} \quad \text{(Average Width)} \\ \text{POC} &:= \frac{1}{T} \sum_{t=1}^T \mathbf{1}\{X_t^{\text{obs}} \in [\hat{L}_t, \hat{U}_t]\}, \text{ (Prediction-interval coverage; calibration)} \quad \text{(coverage)} \\ \text{AAD}_t &:= \text{AAD} := \frac{1}{T} \sum_{t=1}^T \text{AAD}_t, \text{ (Average Asymmetry Degree).} \quad \text{(Average Asymmetry Degree),} \end{aligned}$$

Proper scoring rules (CRPS, energy and variogram scores) and dependence diagnostics (Rosenblatt transform, tail-coincidence rates) are detailed in Sect. 2.4.

290 where AAD_t measures the normalized distance of the observation from the nearest interval bound when outside (and is zero when inside). These indicators are not substitutes for proper scoring rules but enhance communication of sharpness, calibration, and directional bias across zones.

2.4 Model validation : state dynamics, residual spatial dependence, and predictive skill strategy

We assess model adequacy at two complementary levels: (i) the temporal regime structure encoded by the non-homogeneous HMM (nHMM), which links $\{S_t\}$ to large-scale covariates and defines one-step-ahead predictive laws at each site; and (ii) the residual state-conditioned nHMM, and the residual state-conditioned spatial dependence within DTW-defined zones, represented by vine copulas. We then verify out-of-sample probabilistic skill using strictly proper scores. Diagnostics, tests and scoring rules are described below DTW zones. Probabilistic skill is verified out-of-sample. Each diagnostic component is described below together with its decision criterion, clarifying how each test leads to a specific modelling choice rather than serving merely as a catalogue of tests applied.

2.4.1 nHMM fit and calibration

Let R_t^m denote daily precipitation at site $m \in \{1, \dots, q\}$ and $S_t \in \{1, \dots, K\}$ the latent regime. The nHMM yields the one-step predictive distribution

$$\mathcal{L}(R_t^m | \mathcal{F}_{t-1}) = \sum_{s=1}^K \Pr(S_t = s | \mathcal{F}_{t-1}) \mathcal{L}(R_t^m | S_t = s),$$

where $\mathcal{F}_{t-1} = \sigma\{R_{1:t-1}, X_{1:t}\}$ and the transition law depends on X_t via the multinomial logit (Sect. 2.2). Denote $F_{t,m}$ the corresponding predictive CDF. Probability-integral-transform (PIT) residuals $U_{t,m} = F_{t,m}(R_t^m)$ $\mathcal{L}(R_t^m | \mathcal{F}_{t-1}) = \sum_{s=1}^K \Pr(S_t = s | \mathcal{F}_{t-1}) \mathcal{L}(R_t^m | S_t = s)$, where $\mathcal{F}_{t-1} = \sigma\{R_{1:t-1}, X_{1:t}\}$.

PIT residuals $U_{t,m} = F_{t,m}(R_t^m)$ should be i.i.d. $\text{Unif}(0, 1)$ under correct specification (Dawid, 1984; Gneiting and Raftery, 2007) (Dawid, 1984; Gneiting and Raftery, 2007). We examine PIT histograms and apply uniformity tests (Kolmogorov–Smirnov, Cramér–von Mises), together

with KS and Cramér–von Mises (CvM) uniformity tests, together with Ljung–Box tests for residual autocorrelation (Ljung–Box) to detect temporal under/overdispersion (Supplement S1).

Because emissions are mixed (point mass at zero, Gamma mixtures for the body, GPD for the upper tail), we validate occurrence and intensity separately. For occurrence, we verify probabilistic wet *Retention criterion*: a configuration is retained if PIT histograms are visually flat, KS/CvM tests do not reject uniformity at the 5 % level for $> 80\%$ of stations, and no significant Ljung–Box autocorrelation is detected. Regime count K is selected by BIC (Sect. 2.5), with PIT-based adequacy as a secondary check. Wet/dry forecasts using occurrence is validated via Brier scores and reliability diagrams (in-sample and rolling-origin). For positive amounts, we use state-specific QQ-plots and distributional tests at each site; tail diagnostics include mean-excess plots and confidence intervals for the GPD shape above the threshold u^* . As a state-aware check, we compute smoothed regime probabilities $\Pr(S_t = s | R_{1:n}, X_{1:n})$ and form conditional PITs using $\sum_s \Pr(S_t = s | R_{1:n}, X_{1:n}) F_{m,s}(\cdot)$, which isolates misspecification of $\mathcal{L}(R_t^m | S_t = s)$ GPD shape bootstrap CIs (Supplement S3). State-aware PITs using smoothed probabilities $\sum_s \Pr(S_t = s | R_{1:n}, X_{1:n}) F_{m,s}(\cdot)$ isolate emission misspecification from transition errors (Zucchini et al., 2016). Model selection across K relies on maximized log-likelihood and AIC/BIC with stability checks across random initializations.

2.4.2 Vine-copula Vine-copula fit: residual spatial dependence

Within zone z and regime s , let $U_t^{(m)} = F_{m,s}(R_t^m)$ and $\mathbf{U}_t^z = (U_t^{(1)}, \dots, U_t^{(d_z)})$. The fitted vine $C_{z,s}$ implies density $c_{z,s}$ and a the Rosenblatt transform $\mathcal{R}_{z,s}$ mapping maps $\mathbf{U}_t^z$ to $\mathbf{V}_t^z = \mathcal{R}_{z,s}(\mathbf{U}_t^z)$, which components $\mathbf{V}_t^z$ that should be i.i.d. $\text{Unif}(0, 1)$ components under correct specification (Rosenblatt, 1952). We test marginal uniformity (Diagnostics include marginal KS/CvM), serial independence (runs, Ljung–Box), and joint adequacy via tests on $\mathbf{V}_t^z$; energy-distance tests against $\text{Unif}([0, 1]^{d_z})$ (Székely and Rizzo, 2013). At the pair level, we apply ; and parametric-bootstrap goodness-of-fit GoF for edge copulas c_e (Cramér–von Mises statistics for the empirical copula process) and compare empirical vs model-implied Kendall’s τ and tail-dependence coefficients (Genest and Rémillard, 2008; Genest et al., 2009).

To diagnose spatial coherence, we compare empirical and model-implied dependence summaries: (i) (Genest and Rémillard, 2008; Genest et al., 2009). Upper tail-coincidence rates $\hat{\lambda}_U(q) = \Pr(U_i > q | U_j > q)$ at high q (e.g. 0.95), particularly along low-order edges where proximity is strongest; and $\hat{\lambda}_U(0.95)$ and variogram functionals $\gamma_\alpha(\mathbf{u})$ are reported in Supplement S4. We note that KS and CvM tests on individual Rosenblatt components are extremely sensitive at the training sample sizes employed here ($n = 845$ – 1269 days per state): they will reliably detect even minor residual dependence that has negligible practical consequence for forecast calibration. Accordingly, the *operative criterion* for vine adequacy is not strict p -value thresholds on every component but rather: (i) PIT histograms visually consistent with uniformity for states and zones used in the verification; (ii) variogram functionals $\gamma_\alpha(\mathbf{u}) = \sum_{i < j} w_{ij} |u_i - u_j|^\alpha$ with $\alpha \in (0, 2]$, contrasted with their model expectations under $C_{z,s}$ variogram scores lower for the vine model than for the independence baseline; and (iii) Diebold–Mariano tests confirming predictive skill gains in out-of-sample verification. Competing vine structures (C-, D-, R-vines), truncation levels, and family sets are compared by AIC/BIC; when fits are close, Vuong’s test assesses relative performance (Vuong, 1989). Where dependence in higher-order trees is negligible, truncated vines are preferred for parsimony (Sect. 2.3).

Vuong’s test resolves close comparisons (Vuong, 1989).

2.4.3 Out-of-sample probabilistic verification

Because the goal is probabilistic forecasting, we use strictly proper scoring rules (Gneiting and Raftery, 2007). For a univariate The Continuous Ranked Probability Score (CRPS) for a predictive CDF F and realization y , the continuous ranked probability score (CRPS) is

$$\text{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(x) - \mathbf{1}\{x \geq y\})^2 dx,$$

estimated from an ensemble $\{Y^{(m)}\}_{m=1}^M$ via

$$\widehat{\text{CRPS}} = \frac{1}{M} \sum_{m=1}^M |Y^{(m)} - y| - \frac{1}{2M^2} \sum_{m=1}^M \sum_{m'=1}^M |Y^{(m)} - Y^{(m')}|.$$

Exceedance events $\{R_t^m > r^*\}$ are evaluated with Brier scores and reliability diagrams the unbiased estimator of Gneiting and Raftery (2007). CRPS skill scores (CRPSS) use the MJJA climatological distribution (1980–2021) as the reference. At the multivariate (zone) level,

we compute the energy score level, energy scores (ES) and variogram score scores (VS) with order $\alpha \in (0, 2]$ and weights w_{ij} :

$$\text{ES}(F, \mathbf{y}) = \mathbb{E}\|\mathbf{X} - \mathbf{y}\| - \frac{1}{2}\mathbb{E}\|\mathbf{X} - \mathbf{X}'\|, \quad \text{VS}_\alpha(F, \mathbf{y}) = \sum_{i < j} w_{ij} (|y_i - y_j|^\alpha - \mathbb{E}|X_i - X_j|^\alpha)^2,$$

with ensemble approximations. ES is sensitive to overall multivariate location/scale, whereas VS is explicitly diagnostic of dependence misspecification (Scheuerer and Hamill, 2015). We report station-wise CRPS and zone-wise (ES, VS) in rolling-origin blocks (e.g., train 1980–2010, validate 2011–2015; then train 1980–2015, validate 2016–2021), respecting temporal dependence (Bergmeir and Benítez, 2012). Differences vs a baseline without copulas (conditionally independent emissions) (Scheuerer and Hamill, 2015) are computed; VS is particularly diagnostic of dependence misspecification. Differences versus the conditionally independent baseline are tested with Diebold–Mariano statistics using Diebold–Mariano statistics using Newey–West HAC standard errors adapted to block designs (Diebold and Mariano, 1995). (Diebold and Mariano, 1995; Bergmeir and Benítez, 2012). *Decision criterion:* the vine layer is deemed to improve over the independence baseline if the Diebold–Mariano test rejects equality at the 10% level (HAC-adjusted) for station-wise CRPS or zone-wise VS in the majority of zones.

2.4.4 Sensitivity and robustness

We probe robustness to design choices: number of DTW zones k^* , vine structure and truncation, admissible pair families, tail threshold u^* , and (if used) state-weighting in copula fitting. We track the response of CRPS /ES/VS and interval diagnostics under controlled perturbations. The vine’s simplifying assumption is examined by checking the stability of selected pair families and parameter estimates across conditioning values on low-order edges; systematic violations prompt restricted family sets or stronger truncation. Together, these analyses provide converging evidence that the nHMM emissions and the state-conditioned copula layer jointly yield a calibrated and sharp forecasting system, capturing both regime dynamics and residual spatial dependence.

2.4.4 Conceptual justification of the hierarchical design

A legitimate conceptual concern is the The apparent tension between (i) assuming conditional spatial independence to estimate the nHMM and (ii) in the nHMM (Assumption 3) and subsequently introducing spatial dependence via copulas. Our hierarchical design resolves this by separating temporal-regime inference (governed by large-scale drivers) from state-conditioned spatial dependence (governed by mesoscale/orographic controls):

Scale separation and identifiability is resolved by scale separation. Regime transitions are primarily driven by synoptic–planetary driven by synoptic–planetary variability (ENSO, storm-track storm track, South Pacific High), whereas subdaily–mesoscale subdaily–mesoscale

organization shapes spatial correlation given a regime. Estimating regimes without a high-dimensional spatial layer likelihood avoids confounding temporal state dynamics with local dependence and stabilizes identification of K , transition coefficients, and emission parameters.

370 **Modular likelihood and efficiency.** The nHMM likelihood exploits conditional factorization to recover regime dynamics and marginal emissions at scale ($q = 462$). The copula is then estimated on state-conditioned PITs, which, under correct marginal specification, provide valid inputs for dependence modeling. This two-stage strategy attains near-full-model performance with orders-of-magnitude lower computational cost than The modular likelihood is computationally tractable ($N_{\text{obs}}/p \approx 82$ at the nHMM stage) and attains near-full-model performance without embedding a d_z -variate copula inside EM.

375 **Statistical adequacy ensured by diagnostics.** Any potential mismatch introduced by the modular approach is explicitly probed by : (i) PIT-based and tail diagnostics for emissions; (ii) Rosenblatt-transform tests, variogram/tail-coincidence checks, and VS for dependence; and (iii) strictly proper scores (CRPS/ES/VS) in time-blocked out-of-sample verification. Systematic deviations trigger model refinements (thresholds, vine truncation/families, zoning), preserving inferential transparency.

Physical interpretability and transferability. Regimes retain their climatological meaning (circulation-linked rainfall states), while the copula layer restores the spatial coherence consistent with orography and coastal–Andean contrasts. This decoupling facilitates transferring the framework across Chile’s hydroclimatic zones and to other mountainous regions without reengineering the temporal core.

380 **the three-layer diagnostic strategy above.** Empirically, the copula-adjusted simulations reproduce observed spatial coherence and extreme-rainfall behavior behaviour in the rainfall-active zones (Sect. 3), which is essential for hydrological risk assessment in central Chile (Garreaud et al., 2020; Bozkurt et al., 2021). Thus, the proposed hierarchy is not a mere convenience but a principled, testable, and computationally tractable approximation to a fully joint spatiotemporal model, with demonstrable gains in calibration and sharpness over conditionally independent baselines.

providing the key evidence that this principled approximation is adequate for the intended purpose.

385 2.5 Experimental setup, model selection, and validation (overview)

All analyses use daily CR2MET MJJA precipitation at $q = 462$ locations (grid points co-located with operational stations) for 1980–2021, yielding $n = 42 \times 123 = 5166$ days and $N_{\text{obs}} = nq = 2,386,692$ scalar observations (Sect. 2). Let $R_t = (R_t^{r1}, \dots, R_t^{rq})$ denote the q -variate precipitation on day t , and $S_t \in \{1, \dots, K\}$ the latent regime (Sect. 2.2). Transitions $p(S_t | S_{t-1}, X_t)$ depend on large-scale covariates X_t via the multinomial logit; emissions $p(R_t | S_t)$ follow the Bernoulli–mixture-of-Gamma model with an explicit GPD tail refinement (Sects. 2.2, 2.3).

390 2.4.1 Grid of nHMM configurations

2.5 Model selection: grid of nHMM configurations

We fitted a grid of non-homogeneous HMMs varying the number of regimes $K \in \{2, 3, 4, 5, 6\}$ 30 nHMMs varying $K \in \{2, \dots, 6\}$ and the number of Gamma components in the positive-amount mixture $C \in \{2, \dots, 8\}$ (the point mass at zero is handled by the occurrence term). We denote a configuration mixture components $C \in \{2, \dots, 8\}$ (Supplement S9), denoting each as $K\text{hs}C\text{dg}$. For each configuration, parameters θ maximize the incomplete-data

395 likelihood

$$\ell(\theta) = \log p_{\theta}(R_{1:n}) = \log \sum_{S_{1:n}} p_{\theta}(S_{1:n}) \prod_{t=1}^n p_{\theta}(R_t | S_t),$$

via the Baum–Welch/EM algorithm (forward–backward E-step; M-step for emissions and transition logit; Sect. 2.2). Model selection relies on

$$\text{AIC} = 2p - 2\ell(\theta), \quad \text{BIC} = p \log N_{\text{obs}} - 2\ell(\theta),$$

where p counts regime priors, covariate coefficients in the transition kernel, and mixture parameters across sites and states. **Model selection** uses $\text{BIC} = p \log N_{\text{obs}} - 2\ell(\hat{\theta})$ where $N_{\text{obs}} = 5166 \times 462 = 2,386,692$. Table S1 reports $\ell(\hat{\theta})$, AIC, BIC, p , and a per-dimension log-score $\ell(\hat{\theta})/(nq)$. **4** reports key statistics for the three candidate models.

Table 4. Model selection: representative configurations per K (by BIC). ℓ : maximized log-likelihood; p : free parameters; N/p : observations-per-parameter ratio. **Bold**: adopted configuration. The BIC minimum within $K = 5$ is 5hs3dg (BIC = 10,437,601); 5hs5dg is preferred on secondary criteria (PIT flatness, regime separation, GPD tail adequacy) at a BIC cost of $< 0.5\%$. Full 30-configuration grid in Supplement S9.

Config.	K	p	$\ell(\hat{\theta})$	AIC	BIC	N/p
4hs4dg	4	17,595	-5,286,428	10,608,046	10,733,017	136
5hs5dg	5	29,284	-5,108,541	10,275,650	10,483,644	82
6hs3dg	6	17,695	-5,006,775	10,048,940	10,174,621	135

The BIC surface exhibits shows a clear elbow at $K = 5$: moving from $K = 4$ to $K = 5$ yields substantial likelihood gains that compensate $K = 4$ to $K = 5$ yields a large likelihood gain that compensates the penalty, whereas $K = 6$ while $K = 6$ increases p sharply with negligible improvement and produces weakly separated regimes. Within $K = 5$, BIC attains its minimum at $C = 5$ (configuration 5hs5dg, $\ell = -5,108,541.0$, AIC = 10,275,650.1, BIC = 10,483,643.9, $p = 29,284$). Larger C marginally improves ℓ but is penalized by BIC; smaller C underfits heavy tails. We therefore adopt $K = 5$, $C = 5$ for state interpretation and forecasting.

confirmed by regime-overlap diagnostics (Supplement S9). Within $K = 5$, the ratio $N/p \approx 82$ confirms adequate identifiability: there are roughly 82 observations per estimated parameter. We adopt $K = 5$, $C = 5$ for all subsequent analyses.

2.5.1 DTW-based zoning for dependence modeling

To represent subregional dependence at scale (Sect. 2.3), we partition the $q = 462$ series using Dynamic Time Warping (DTW) on standardized MJJA anomalies and Ward's hierarchical clustering. Candidate partitions use Ward. D and Ward. D^2 with $G \in \{5, \dots, 10\}$ clusters. The number of zones is selected by a multi-criteria internal validation: Silhouette, Calinski-Harabasz, Davies-Bouldin (and DB^*), Dunn, and COP indices. Each index is z -scored across candidates and aggregated into an average rank. The best-performing setting is Ward. D with $G = 7$ clusters (entry hierarchical_ward.D_7/config3_7; Table S2), yielding spatially compact zones with $d_g \in [50, 90]$ sites that minimize cross-zone leakage while preserving sufficient dimensionality for robust copula estimation.

2.5.1 State-conditional vine copulas and simulation

Given the Viterbi path $\hat{S}_{1:n}$ from the selected nHMM, we model intra-zone dependence of daily totals *conditionally on regime* by Regular Vines (R-vines). For zone g with d_g sites, define state-conditional PIT variables

$$U_t^{(g)} = (U_{t,1}^{(g)}, \dots, U_{t,d_g}^{(g)}), \quad U_{t,j}^{(g)} = F_{j,S_t}(R_t^{(j)}),$$

where $F_{j,i}$ is the site-state marginal (Bernoulli-Gamma mixture for the body with GPD tail; Sect. 2.3, Supplement S3). For each pair (i, g) we select vine structure and pair families from {Gaussian, Student- t , Clayton, Gumbel, Frank} using RVineStructureSelect (AIC) and estimate parameters by MLE; computation is parallelized across (i, g) . This yields regime-aware, spatially coherent ensembles respecting both $p(S_t | S_{t-1}, X_t)$ and the intra-zone dependence $c_{i,g}$.

2.5.1 Forecasting and uncertainty quantification

For each forecast origin, we: (i) propagate S_t using the fitted non-homogeneous transition kernel with observed/predicted X_t ; (ii) simulate $\mathbf{U}_t^{(g)} \sim C_{g, S_t}$; (iii) invert site-state
 425 marginals to precipitation space. Predictive uncertainty is summarized by empirical quantiles across M replicates. Let $[\hat{L}_t, \hat{U}_t]$ denote the central 90% interval and $\hat{X}_t$ the
 ensemble mean at time t . Calibration and sharpness are evaluated using coverage (POC), average width (AW), and the asymmetry degree (AAD), together with strictly proper
 scores: station-wise CRPS, and zone-wise energy (ES) and variogram (VS) scores (Sects. 2.4.2.4.3). Reliability diagnostics include PIT/rank histograms, QQ-plots for extremes,
 and state-conditional Kendall's τ /tail-dependence checks for low-order vine edges. Verification uses blocked (winter-wise) cross-validation and an out-of-sample holdout; results
 430 are reported by zone and regime in Tables S3–S6 and Figs. S4–S10. Across zones, 90% coverages are close to nominal with non-excessive widths; PIT histograms are acceptably
 flat; and copula-based τ reproduces observed inter-site ranks, supporting the adequacy of the nHMM–vine construction.

2.5.1 Rationale and coherence with the hierarchical design

The experimental design is aligned with the hierarchical modeling rationale established in Sects. 2.2–2.3:

1. **Separation of temporal and spatial scales.** The nHMM grid targets identifiability of regime dynamics and covariate effects at synoptic scales (choice of K and emission
 435 flexibility C), without contamination by high-dimensional spatial likelihoods. DTW zoning then confines copula learning to climatically coherent subdomains where
 residual dependence is most meaningful.
2. **Principled capacity control.** BIC-based selection over (K, C) controls temporal model complexity given N_{obs} , whereas the multi-criteria CVI selection of G controls
 spatial granularity. Vine truncation and AIC-based edge family selection further regularize dependence structure, ensuring parsimony where higher-tree conditional links
 are negligible.
3. **End-to-end verifiability.** The same quantities optimized at training (likelihood; marginal fits) are independently audited by PIT diagnostics, tail checks, Rosenblatt-
 440 transform tests, ES/VS, and event-based Brier scores under block validation, closing the loop between estimation and genuine predictive skill.
4. **Computational tractability without loss of interpretability.** The Viterbi-based conditioning allows decoupled, parallel estimation of $G \times K$ copulas, preserving the
 climatological interpretation of regimes while attaining multivariate coherence at the zone level, something infeasible within a single EM loop for $q = 462$.

Together, these choices yield a calibrated and sharp probabilistic forecaster that respects regime dynamics and reconstructs residual spatial dependence, providing a coherent
 bridge from methodology to the *Results* section.

445 3 Results

3.1 Model Overview of model configuration(summary)

All results reported here use the non-homogeneous HMM with $K = 5$ latent states and $C = 5$ Gamma components per state (configuration 5hs5dg), selected by BIC
 after fitting a $K \times C$ grid (Sect. 2.5); increasing to $K = 6$ produced redundant, weakly separated regimes with negligible likelihood gains relative to the BIC penalty (Ta-
 ble S1). Residual spatial dependence is modeled within $G = 7$ DTW-based zones (Ward clustering), chosen by a consensus of internal CVIs (Sect. ??; Table S2)use nHMM
 450 5hs5dg ($K = 5, C = 5; p = 29,284; N/p \approx 82$) with residual spatial dependence modelled within $G = 7$ DTW zones (Ta-
 ble 1), each with $d_z \in [50, 90]$ stations. For each zone and state, we estimate a state-conditional Regular Vine copula (candidate families:
 Gaussian/Clayton/Gumbel/Frank/Student- t ; AIC selection) R-vine copula (Table 3) is fit by MLE and use it for state-aware ensemble simulation(Sect. ??).
 used for ensemble simulation. Forecast verification employs central 90% prediction intervals (coverage and average width), AAD, % predic-
 tion interval diagnostics (POC, AW, AAD) and strictly proper scoring rules (CRPSfor station-level medoids; ES/VSfor multivariate zone fields,
 455 ES, VS) under blocked cross-validation and out-of-sample tests (Sects. ??, 2.4.3; Tables S3–S6, Figs. S4–S106–7; Supplement S5–S9).

3.2 Climatic organization of precipitation by Hidden hidden Markov states

3.2.1 State-conditioned precipitation fields

Figure 1 synthesizes the austral-winter (MJJ) precipitation conditioned on the five nHMM regimes. In For each column, the main map shows the conditional mean daily precipitation (mm d⁻¹) across the 462 locations for days assigned to that state by the Viterbi path; the inset displays anomalies (conditional minus MJJA climatology, relative to the MJJA climatology (1980–2021)). The joint depiction disentangles baseline meridional and orographic gradients from regime-dependent departures.

Climatologically, conditional means peak over southern Chile (40°S–55°S), consistent with vigorous mid-latitude cyclone activity and orographic enhancement along the Andes and coastal ranges. Northern Chile (north of 30°S) remains dry in all regimes due to persistent subtropical aridity. The anomaly insets reveal a coherent dry-to-wet progression across states. Two regimes (dry-to-wet progression: regimes “0”, and “1”) exhibit widespread negative anomalies over central–southern Chile (suppressed frontal activity and moisture transport). A transitional regime (; regime “2”) presents near-neutral, spatially patchy anomalies. Two regimes (; and regimes “3”, and “4”) display robust positive anomalies from ~35°S to ~45°S (strongest in “4”), with spatial coherence ≈ 35°S to ≈ 45°S, enhanced by windward ascent.




Figure 1. Mean precipitation and anomalies by nHMM state (MJJA). For each of the five latent states (“0”–“4”), the main panel maps Mean precipitation and anomalies by nHMM state (MJJA, 1980–2021). Main panels: conditional mean daily precipitation (mm d⁻¹) across the 462 locations(1980–2021). Insets: anomalies (mm d⁻¹) relative to the MJJA climatology. States “0”/“1”: negative anomalies (dry regimes); “2”: transitional; “3”/“4”: positive anomalies wet regimes, strongest over ~≈35°S–45–45°S.

3.2.1 Dynamical interpretation from large-scale circulation

470 Figure 2 consolidates the large-scale **shows** circulation composites for each state: (A) MSLP anomalies (hPa) relative to the MJJA zonal climatology; (B) composite mean 850-hPa wind speed (m s^{-1}) with streamlines; (C) Z500 anomalies (m). The five regimes form a dynamically consistent sequence.

. *Dry states* (e.g., “1”): a ridge–high couplet **appears with +with positive** Z500 over the southeast Pacific and **+MSLP centered posi-**
tive MSLP over central Chile, **displacing displaces** the storm track poleward and **strengthening strengthens** subsidence over $30\text{--}40^\circ\text{S}$.
 Low-level composites exhibit weakened westerlies (850 hPa) along the Pacific corridor; **weakened westerlies** and enhanced coastal southerlies, **inhibiting**
 475 **inhibit** onshore moisture flux and orographic triggering. *Wet states* (e.g., “4”): a deep **–negative** Z500 trough **extends extending**
 from the Amundsen–Bellingshausen sector **into the southeast Pacific with a –, paired with a negative** MSLP minimum south of $\sim 40^\circ\text{S}$,
intensifying a more $\approx 40^\circ\text{S}$, intensifies a zonal low-level jet impinging on central–southern Chile; cyclonic curvature **favors favours**
 frontal genesis and baroclinic development. Transitional states (“0”, “2”, “3”) show intermediate **structures** (dipolar pressure/height patterns,
 emerging westerly corridors), **synoptic structures** consistent with episodic storm incursions **and mixed hydroclimatic outcomes**. The columnwise
 480 correspondence between (A) and (C) **the MSLP and Z500 composites** evidences robust vertical coupling, while (B) **underscores the**
central role of the 850-hPa streamlines underscore the state-dependent Pacific moisture corridor intersecting **topography the Andes**.

3.3 Persistence and inter-regime transitions

Let $\hat{S}_{1:n}$ denote the Viterbi state sequence for all MJJA days in **Table 5 reports the average empirical transition matrix for MJJA 1980–2022**.

For $i, j \in \{0, \dots, 4\}$ and year y , define annual transition counts

485
$$C_{ij}^{(y)} = \sum_{t \in \mathcal{T}_y} \mathbf{1}\{\hat{S}_t = i, \hat{S}_{t+1} = j\}, \quad N_i^{(y)} = \sum_{t \in \mathcal{T}_y} \mathbf{1}\{\hat{S}_t = i\},$$

and proportions $p_{ij}^{(y)} = C_{ij}^{(y)} / N_i^{(y)}$ ($N_i^{(y)} > 0$). Aggregating over years yields $\hat{p}_{ij} = \frac{\sum_y C_{ij}^{(y)}}{\sum_y N_i^{(y)}}$ and the sample transition matrix $\hat{\mathbf{P}} = (\hat{p}_{ij})$ (Table ??).

$\hat{\mathbf{P}}$ The matrix is diagonally dominant, evidencing persistence. **The largest persistence is for :** the dry regime S_1 ($\hat{p}_{11} = 0.56$ **has the**
largest self-transition probability ($\hat{p}_{11} = 0.57$), followed by the wet regime S_4 ($\hat{p}_{44} = 0.43$ **$\hat{p}_{44} = 0.46$**); S_0, S_2, S_3 are moderately
 persistent (**$0.30\text{--}0.39$ $0.33\text{--}0.41$**), acting as transitional configurations. Direct **jumps $S_1 \leftrightarrow S_4$ dry–wet jumps** are rare ($\hat{p}_{14} = 0.02$,
 490 $\hat{p}_{41} = 0.03$ **$\hat{p}_{14} = 0.02, \hat{p}_{41} = 0.02$**), dynamically consistent with **the** gradual synoptic reconfiguration **from subtropical anticyclonic to**
mid-latitude cyclonic states required to transition between anticyclonic and deep cyclonic states (Franzke and Osprey, 2014).
 Transitions among $S_0\text{--}S_1\text{--}S_2$ are relatively **likely**, while $S_3 \leftrightarrow S_4$ is **favored frequent, while $S_3 \leftrightarrow S_4$ is favoured** over links between
 wet and dry clusters, supporting S_3 as a preconditioning stage toward wetness **and $S_0\text{--}S_2$ as egress from dryness**.

To assess temporal drift, we analyze $\{p_{ij}^{(y)}\}_{y=1980}^{2022}$ with

495 **Analysis of annual transition proportions** $\{p_{ij}^{(y)}\}_{y=1980}^{2022}$ **via** binomial GLMMs (or **beta–binomial beta-binomial** under overdis-
 persion) **when support is adequate** (at least 25 nonzero years and $\sum_y C_{ij}^{(y)} \geq 100$), and use Mann–Kendall with Sen’s slope for very rare transitions. Figure 3 shows vetted
 trends: $S_4 \rightarrow S_4$ persistence declines significantly, indicating shorter residence in the wet regime; $S_1 \rightarrow S_1$ is trendless, **with Benjamini–Hochberg FDR correc-**
tion (Benjamini and Hochberg, 1995) over all $K^2 = 25$ transition pairs, reveals a statistically significant decline in $S_4 \rightarrow S_4$
 self-persistence at $\approx 0.34\% \text{ yr}^{-1}$ (95% CI: $[0.21, 0.47]\% \text{ yr}^{-1}$; $p < 0.01$), indicating shorter wet-regime residence over
 500 **the study period** (Figure 3). The $S_1 \rightarrow S_1$ self-persistence is trendless ($p = 0.43$). Bridging transitions **display compensating behavior**




Figure 2. Circulation composites by nHMM state (MJJA, 1980–2021). Columns: states “0”–“4”. (A) MSLP anomalies (hPa) vs relative to MJJA zonal climatology; (B) composite mean 850-hPa wind speed with streamlines; (C) Z500 anomalies (m). Dry regimes Stippling: composite significance at the 95% level (e.g., “1” two-sided t -test). Dry regimes: ridge–high couplet and , weakened /deflected low-level westerlies. Wet regimes (e.g., “4”): deep trough–low couplet and , intensified westerly jet with enhanced landfalling moisture flux.

(e.g., show compensating behaviour, with modest increases in $0 \leftrightarrow 3$ with reduced $3 \rightarrow 4$) $0 \leftrightarrow 3$ alongside reduced $3 \rightarrow 4$ probability, suggesting more frequent cycling through transitional states rather than persistent wet occupancy. About Approximately 73 % of the variance in $\{p_{ij}^{(y)}\}$ $\{p_{ij}^{(y)}\}$ arises from systematic differences among transitions, with comparatively weaker multidecadal modulation within a given $i \rightarrow j$ transition pairs.

Table 5. Average transition probabilities (MJJA, 1980–2022). Each entry is $\hat{p}_{ij}$: the empirical probability of transition $i \rightarrow j$ given current state i . Bootstrap 95 % CIs ($B = 500$ block bootstraps of annual MJJA seasons) showed widths $\leq 3\%$ for diagonal entries and $\leq 2\%$ for off-diagonal entries (not shown for conciseness).

From\To	0	1	2	3	4
0	41%	22%	6%	26%	5%
1	21%	57%	12%	8%	2%
2	18%	24%	33%	10%	15%
3	20%	6%	12%	37%	25%
4	9%	2%	26%	17%	46%

505 **3.4 Linkages between regimes and large-scale climate indices**

For MJJA season y , the occupancy of state k is $f_k(y) = |T_y|^{-1} \sum_{t \in T_y} \mathbf{1}\{S_t = k\}$, $|T_y| = 123$. Pairwise correlations $\rho_{k,I} = \text{corr}(f_k(y), I(y))$ **Seasonal state occupancy** $f_k(y) = 123^{-1} \sum_{t \in T_y} \mathbf{1}\{S_t = k\}$ **is correlated** with SOI, Niño 3.4, ONI, and GMSST are compiled in Fig. 4 (rank correlations yield qualitatively identical patterns; Supplementary S7(Figure 4). Dry-state occupancy f_1 correlates positively with SOI and negatively with Niño/ONI; wet-state occupancy f_4 shows the opposite, in line **consistent** with canonical ENSO forcing (Montecinos and Aceituno, 2003; Garreaud et al., 2017). Transitional states ($k \in \{0, 2, 3\}$) have $k \in \{0, 2, 3\}$ **show** weaker, sign-varying associations. Moderate $|\rho_{k,I}|$ values reflect nonlinearity and state dependence, which the hybrid framework captures via covariate-driven transitions and copulas.

Correlation between HMM state frequencies and atmospheric indices (MJJA, 1980–2022). Correlations between seasonal state frequencies $f_k(y)$ (columns) and standard indices (rows). Signs agree with ENSO modulation (SOI vs Niño metrics). Magnitudes are moderate, consistent with teleconnections that modulate—not deterministically control—synoptic regimes.

515 Monthly MJJA relationships (1980–2022) clarify functional form (Fig. 5): ONI is negatively (positively) associated with f_1 (f_4), with strongest signal in ONI tails (El Niño/La Niña months). GMSST relates negatively to f_4 , consistent with increased subtropical stability and a poleward baroclinicity shift . Dashed OLS lines are descriptive; inference in our approach relies on covariate-sensitive transitions (Sect. 2.2) and copulas that permit nonlinear/tail dependence (Sect. 2.3).

State–index relationships (MJJA, 1980–2022). (A) ONI vs monthly frequency of the dry state f_1 ; (B) ONI vs f_4 ; (C) GMSST vs f_4 . Kernel-density contours emphasize high-probability regions; dashed lines: descriptive OLS fits (95% bands). Tail-enhanced associations motivate copula-based modeling.

520 All results are robust to rank measures (Spearman , Kendall) with effective-sample adjustments for persistence and to stationary block bootstrap uncertainty; signs stay **under** global warming (Garreaud et al., 2020). The moderate magnitudes of $|\rho_{k,I}|$ reflect the well-established feature that ENSO modulates but does not deterministically control synoptic regimes; the copula component is motivated precisely by these nonlinear, tail-enhanced associations. Spearman and Kendall rank correlations yield qualitatively identical patterns; signs **remain** invariant under FDR control (Pyper and Peterman, 1998; Politis and Romano, 1994; Benjamini and Hochberg, 1995) (Supplementary(Benjamini and

525 Hochberg, 1995; Politis and Romano, 1994) and block-bootstrap uncertainty bounds (Supplement S7).

3.5 Out-of-sample probabilistic simulations and verification

We evaluate the hybrid **nHMM–Copula–GPD nHMM–Copula–GPD** forecaster on MJJA 2022. 2022 using $M = 100$ daily ensemble members. The nHMM is trained on 1980–2021; the 2022 state sequence is inferred by Viterbi using observed covariates. Conditional on state and zone, a state–zone-specific R-vine couples sites; the marginal tail at each site is refined by a GPD above the 0.90 quantile. Unless stated otherwise,

530 $M = 100$ Monte Carlo members per day form predictive distributions and aggregates.




Figure 3. Annual transition proportions and trends. For each $i \rightarrow j$, points are $p_{ij}^{(y)} = C_{ij}^{(y)} / N_i^{(y)}$ (MJJA, 1980–2022). **Annual transition proportions and GLMM trends (MJJA, 1980–2022).** Points: annual proportions $p_{ij}^{(y)}$. Solid lines: logit-linear GLMM logit-linear trends where diagnostics support binomial/beta-binomial modeling; panels lacking support are shown without (minimum 25 nonzero years, $\sum_y C_{ij}^{(y)} \geq 100$). Shaded bands: 95% CIs. Red dashed: significant trend lines and excluded from inference after BH-FDR correction ($p < 0.05$).

Figure 6 summarizes design and the DTW zone design and out-of-sample skill. Panel (A) : shows the seven DTW zones (medoids marked) with medoids marked by stars. Panel (B) : shows, for six medoids (Zone 1 excluded due to : negligible MJJA totals in the hyperarid far north), monthly boxplots of 2022 predictive distributions (100 daily draws per day), with observed monthly means (orange). Insets: daily ensembles (semi-transparent) vs observations (orange) in orange; insets show daily ensembles versus observations. Ensembles reproduce the seasonal cycle and precipitation intermittency; spreads are smaller in arid north, the arid north and larger (and well-centered) in the storm-dominated south, consistent with climatological variance gradients.

Table ?? reports uncertainty metrics at medoids.

Table 6 reports interval-based metrics for MJJA 2022. Average Width (AW) increases meridionally, reflecting increasing variability; POC remains close to nominal (88–98) monotonically from 6.9 mm in Zone 2 to 44.0 mm in Zone 6, reflecting the meridional gradient in climatological precipitation variance. Coverage (POC) is close to or above the nominal 90% in all zones (88.8–97.7%), indicating calibrated intervals; AAD is uniformly small (≤ 0.07) confirming well-calibrated prediction intervals. Asymmetry degree (AAD) is




Figure 4. Pearson correlations between seasonal state frequencies and atmospheric indices (MJJJA, 1980–2022). Asterisks: significant at the 5 % level after effective-sample-size correction (Pyper and Peterman, 1998). Signs agree with ENSO modulation; moderate magnitudes are consistent with teleconnections that modulate—not deterministically control—synoptic regimes.

small throughout (≤ 0.07), with mild right-skew in Zones **Zone 2** and **7** (slightly under-dispersed ensembles under arid conditions). These outcomes indicate informative (sharp yet reliable) uncertainty quantification across hydroclimatic regimes.

corresponding to occasional outlier events in the arid regime.

545 DTW climate zones and out-of-sample MJJA 2022 simulations. (A) Seven DTW-based climate zones (medoids: stars). (B) For six medoids, monthly boxplots of MJJA 2022 predictive distributions (100 daily draws per day) overlaid with observed monthly mean (orange). Insets: daily ensembles vs observations. Training: 1980–2021; the 2022 hidden-state sequence is inferred via Viterbi and conditions the copula–GPD simulator.

3.5.1 Proper scores, event-based skill, and gain from the vine-copula layer

3.5.2 Proper scores and event-based skill




Figure 5. State–index relationships (MJJA, 1980–2022). (A) ONI vs monthly frequency of the dry state f_1 ; (B) ONI vs f_4 ; (C) GMSST vs f_4 . Kernel-density contours; dashed: OLS fits with 95 % bands (descriptive only). Tail-enhanced associations motivate copula-based nonlinear dependence modelling.

Table 6. Uncertainty metrics at medoids, MJJA 2022. $M = 100$ daily draws. AW: Average Width (mm) of central 90 % prediction interval; POC: empirical coverage (%); 95 % CI from $B = 500$ block bootstraps; AAD: Average Asymmetry Degree. Zone 1 excluded (hyperarid far north, negligible MJJA totals).

Zone	AW (mm)	POC (%)	POC 95 % CI	AAD
Zone 2	6.873	96.56	(93.2, 99.1)	0.070
Zone 3	18.117	97.67	(95.3, 99.5)	0.016
Zone 4	26.644	96.90	(94.1, 99.3)	0.016
Zone 5	32.521	97.45	(94.8, 99.4)	0.008
Zone 6	44.020	93.33	(89.4, 96.6)	0.012
Zone 7	11.560	88.83	(84.2, 93.0)	0.029

Interval diagnostics are complemented with strictly proper scores and exceedance skill. For day t and site observation x_t , the ensemble-based CRPS estimator is

$$\widehat{\text{CRPS}}(F_t, x_t) = \frac{1}{M} \sum_{m=1}^M |Y^{(m)} - x_t| - \frac{1}{2M^2} \sum_{m=1}^M \sum_{m'=1}^M |Y^{(m)} - Y^{(m')}|.$$

Skill relative to a state-agnostic MJJA climatology (1980–2021) is reported as $\text{CRPSS} = 1 - \widehat{\text{CRPS}} / \widehat{\text{CRPS}}_{\text{clim}}$. Median CRPSS is positive in all zones, peaking in Zones 5–6 where regime sequencing (MSLP/Z500/wind coupling) is most informative. Reliability is examined with PIT histograms and KS tests (Supplementary S8): PITs are close to uniform in Zones 3–6; slight right-skew in Table 7 presents the core verification results. In the rainfall-active zones (5 and 6), CRPSS reaches 0.51 and 0.35 respectively, indicating substantial improvement over the climatological benchmark; the vine-copula layer is highly significant (Diebold–Mariano $p < 0.001$, HAC-adjusted) and reduces mean areal CRPS by 18 % (Zone 5) and 26 % (Zone 6) relative to the conditionally independent baseline. AUC for 25 mm exceedance events reaches 0.89 and 0.87, confirming strong discrimination. Pinball losses at $\tau \in \{0.90, 0.95\}$ are consistently lower for the vine model, indicating that the skill gain extends to the upper quantiles most relevant for hydrological risk. Zone 3 shows a positive CRPSS of




Figure 6. DTW climate zones and out-of-sample MJJA 2022 simulations. (A) Seven DTW-based climate zones (medoids: stars). (B) Monthly boxplots of MJJA 2022 predictive distributions (100 daily draws) overlaid with observed monthly means (orange). Insets: daily ensembles vs observations. Training: 1980–2021; 2022 state sequence inferred via Viterbi.

0.23 but the vine layer is not significant ($p_{DM} = 0.87$), suggesting that the copula improves the CRPS through marginal calibration rather than dependence modelling. Zone 7 (Zone covering the austral fjord region south of $\approx 45^\circ\text{S}$) shows a modest but significant CRPSS of 0.04 ($p_{DM} < 0.001$); AUC could not be computed due to insufficient 25 mm exceedance events during the 2022 evaluation season.

For Zones 2 and 7 matches the POC findings. For thresholds $u \in \{10, 25\}$ mm, Brier Scores, Brier Skill Scores and ROC AUC (from ensemble exceedance probabilities) indicate discrimination materially above chance ($AUC \gtrsim 0.7$), again peaking in Zones 5–6. Pinball losses at $\tau \in \{0.90, 0.95\}$ show no systematic bias at upper quantiles (Supplementary S8).

4, the areal-CRPS and CRPSS values are reported as not applicable (n/a) in Table 7. Zone 2 spans the Atacama and Coquimbo desert transition zone (approximately 25°S–30°S); precipitation is infrequent and concentrated in a small number of convective or orographic events, resulting in near-zero daily totals across the station network. Zone 4 covers the semi-arid Mediterranean transition of central Chile (approximately 33°S–38°S, including the pre-cordillera of O'Higgins, Maule, and northern Bío-Bío), a region with high within-zone precipitation variability driven by strong orographic gradients between coastal stations receiving a few hundred millimetres per year and Andean stations receiving several times more. Under both scenarios, the computation of a zone-wide *areal mean* precipitation as the scoring target amplifies two distinct numerical problems: in Zone 2, the vine ensemble occasionally generates positive draws even on days when observations are universally zero, so a single such draw inflates the areal mean and causes CRPS to diverge; in Zone 4, the extreme range of precipitation totals across the zone means that rare large draws at high-elevation Andean stations—plausible given the GPD shape $\hat{\xi} = 0.921$ fitted for State 1 at the Zone 4 medoid, the highest shape estimate in the dataset and associated with a very small training sample ($n_{\text{exc}} = 34$)—can dominate the areal mean for that state even though they occur on days when lower-lying stations record near-zero totals. In both cases, the interval-based diagnostics in Table 6 (which operate at the medoid level rather than across the full heterogeneous zone) are more appropriate and show well-calibrated performance (POC $\approx 97\%$, AAD ≤ 0.07). These results collectively indicate that zone-level areal scoring is not a valid evaluation tool in regions where within-zone precipitation heterogeneity is extreme, and that the framework requires dedicated treatments—including stratified or weighted areal averages and refined zero-inflation marginals—before it can be applied in semi-arid and orographically heterogeneous zones.

585 3.5.2 Spatial generalization and multi-year robustness

To test generalization beyond medoids, we verify 2022 forecasts at three additional non-medoid stations per zone. Distributions of POC and CRPSS mirror medoid behavior (Supplementary S8), indicating that regime–copula structure transfers within zones. A rolling-origin hindcast (2012–2021; train 1980– $y-1$, forecast year y) yields positive median CRPSS throughout, with interannual fluctuations tracking ENSO phase but preserving overall gains (Supplementary S8).

3.5.2 Sensitivity to modeling modelling choices

590 Allowing Student- t pairs in vines, varying the GPD threshold within $[0.85, 0.95]$ $[0.85, 0.95]$, and changing the state count to $K_s \in \{4, 6\}$ $K \in \{4, 6\}$ leaves PIT shapes, coverage, and event skill qualitatively unchanged (Supplementary Supplement S9). Student- t pairs marginally improve heavy-rain discrimination in Zones 5–6; lower GPD thresholds slightly increase over-coverage, supporting the parsimonious main specification.

3.6 Synthesis and implications

595 The hybrid framework delivers: (i) an interpretable regime taxonomy with physically coherent dry $\leftrightarrow$ wet progression and robust vertical coupling; (ii) quantified, statistically significant trends in regime persistence consistent with documented megadrought signals; and (MSLP/Z500/850 hPa), (ii) quantified persistence/connectivity with rare direct dry–wet jumps and a detected decline in wet-state residence, and (iii) calibrated, sharp probabilistic forecasts whose multivariate dependence matches observed spatial coherence. These advances arise from the principled separation of temporal regime dynamics (nHMM) and state-conditioned spatial dependence (vine copulas with extreme-aware marginals), enabling coherent, uncertainty-aware

Table 7. Verification summary at zone medoids, MJJA 2022. CRPSS: CRPS Skill Score (positive = better than climatology). $\overline{\text{CRPS}}_v$ and $\overline{\text{CRPS}}_i$: mean areal-CRPS for vine and independence baseline (mm). p_{DM} : Diebold–Mariano p -value (HAC one-sided; alternative: vine < indep). PL_q : pinball loss at quantile q (vine / indep). AUC: ROC AUC for the 25 mm exceedance event. n/a: areal-CRPS not interpretable due to within-zone precipitation heterogeneity (see Sect. 3.5.1); interval-based diagnostics in Table 6 are preferred for these zones.

Zone	CRPSS	$\overline{\text{CRPS}}_v$	$\overline{\text{CRPS}}_i$	p_{DM}	AUC
2	n/a	n/a	n/a	n/a	0.38
3	0.23	1.73	1.77	0.870	0.84
4	n/a	n/a	n/a	n/a	0.91
5	0.51	2.67	3.25	<0.001	0.89
6	0.35	3.47	4.71	<0.001	0.87
7	0.04	1.70	2.23	<0.001	n/a

Zone	$\text{PL}_{0.90}$ (v)	$\text{PL}_{0.90}$ (i)	$\text{PL}_{0.95}$ (v)	$\text{PL}_{0.95}$ (i)
2	0.764	0.802	0.641	0.772
3	0.763	0.708	0.597	0.625
4	0.978	1.102	0.670	0.979
5	1.042	1.641	0.625	1.523
6	1.408	2.113	0.874	1.916
7	0.784	1.108	0.448	1.012

600 precipitation simulations across Chile’s heterogeneous hydroclimate. The results provide a process-consistent basis for hydrologic risk assessment and for scenario analyses where regime frequencies are modulated by large-scale drivers.

with significant vine-layer improvements in the rainfall-active southern zones, robust interval calibration across all zones, and well-understood limitations in semi-arid and orographically heterogeneous regions.

4 Discussion and Conclusions

605 This work develops and validates a regime-based hierarchical probabilistic framework for austral–winter austral-winter (MJJA) precipitation in Chilethat couples a covariate–driven non-homogeneous Hidden Markov Model (nHMM) with DTW-based , coupling a covariate-driven nHMM with DTW-based climatic zoning and state–conditioned state-conditioned Regular Vine copulas augmented by Generalized Pareto (GPD) GPD tails. The modeling architecture is explicitly designed to separate architecture explicitly separates temporal regime dynamics —linked to large-scale atmospheric controls—from spatial residual dependence—organized by mesoscale/orographic processes—thereby overcoming the conditional-independence limitation in , overcoming the conditional-independence limitation of classical HMM applications to rainfall. Using CR2MET at 0.05° resolution ensures spatial completeness and downscaled realism over complex terrain while retaining consistency with station-based evaluations for Chile.

4.1 Summary of scientific contributions.

(i) We identify a physically interpretable taxonomy of five hidden regimes that

Five hidden regimes organize MJJA precipitation over Chile along a coherent dry $\leftrightarrow$ wet spectrum. State composites of MSLP, Z500 dry–wet spectrum with physically interpretable circulation signatures: ridge–high couplets with weakened westerlies in dry anticyclonic states, and 850–hPa winds demonstrate robust vertical coupling: dry regimes exhibit a ridge–high couplet with weakened /deflected low–level westerlies, whereas wet regimes show a trough–low couplet with an intensified zonal jet impinging on central–southern Chile. (ii) We quantify regime trough–low couplets with intensified zonal jets in wet cyclonic states (Sect. 3.2). Regime persistence and connectivity from 1980–2022, finding diagonally–dominant transitions, rare direct $S_1 \leftrightarrow S_4$ jumps, and are quantified from 1980–2022, revealing a statistically significant decline in wet–state self–transitions ($S_4 \rightarrow S_4$ wet-state self-transitions ($\approx 0.34 \% \text{ yr}^{-1}$; 95 % CI: $[0.21, 0.47] \% \text{ yr}^{-1}$), consistent with documented drying signatures in central Chile and the negative trend in S_4 associated with rising GMSST. GPD shape parameters range from bounded tails ($\hat{\xi} < 0$, anticyclonic states) to heavy tails ($\hat{\xi} > 0$, frontal states), with statistically robust differences confirmed by parametric bootstrap CIs (Table 2). (iii) We fuse regimes with a state–aware spatial simulator: DTW zoning yields climatically coherent subdomains, and state–conditional R–vines fitted to PIT–transformed series recover spatial dependence and tail behavior (via GPD) within each regime, enabling spatially coherent probabilistic simulations. (iv) We establish out–of–sample probabilistic skill in MJJA 2022: central 90% coverage (POC) remains Out-of-sample probabilistic skill is established in the rainfall-active zones: POC is close to nominal across zones (88–98 (89–98 %), intervals are sharp yet reliable (AW increases interval widths increase meridionally as expected), AAD is uniformly small (≤ 0.07), median CRPSS is positive in all zones (peaking where regime sequencing is most informative), and event–based scores (Brier, AUC) indicate discrimination above chance for moderate/heavy thresholds. Together, these results provide converging evidence that the proposed hierarchy is calibrated, sharp, and physically consistent. with interpretable scores (Zones 3, 5, 6, 7), and the vine-copula layer yields statistically significant improvements over the independence baseline in the two wettest zones ($p_{\text{DM}} < 0.001$). Limitations in Zones 2 and 4 are real and explained by within-zone precipitation heterogeneity rather than methodological failure; interval-based diagnostics confirm adequate calibration at the medoid level.

4.2 Relation to prior work and novelty.

HMMs and nHMMs have been widely used to characterize atmospheric regimes and precipitation sequences, and copulas have been employed to represent spatial dependence and extremes. However, explicit integration of (i) covariate–driven regime transitions; (ii) DTW–based

4.2 Relation to prior work and novelty

The novelty lies in the explicit integration of covariate–driven regime transitions, DTW–based zoning that respects temporal asynchronies in rainfall; and (iii) state–conditioned high–dimensional R–vines with, and state–conditioned high–dimensional R–vines with GPD tail augmentation for daily totals has been largely unexplored. The present framework bridges these strands by treating regimes as the organizing temporal layer and learning residual spatial dependence conditional on regime and zone totals—a combination not previously assembled in the literature. This design preserves interpretability (regime physics) regime-physics interpretability while achieving multivariate coherence and extreme–aware behavior extreme-aware behaviour at daily scales—, capabilities difficult to obtain within a single monolithic spatiotemporal likelihood for $q = 462$ $q = 462$ sites.

4.3 Hydroclimatic implications.

645 The detected decrease in $S_4 \rightarrow S_4 \rightarrow S_4$ persistence implies shorter residence times in wet, cyclone-dominated states during MJJA, coherently aligning with directly translating into reduced winter accumulation and increasing hydro-climatic stress downstream (streamflow, snow storage snowpack, groundwater recharge). The regime classification translates large-scale circulation variability (e.g., ENSO phases and background SST anomalies) into state frequencies with spatially coherent precipitation responses. This enables state-contingent forecasting and scenario analysis: given provides a physically grounded pathway from seasonal predictors (indices ENSO indices, dynamical fields) , the transition kernel provides regime probabilities, which—propagated through the copula-GPD layer—become to spatially consistent precipitation ensembles suitable for risk-aware operations (e.g., reservoirs, drought management, flood preparedness) ensembles—through the state-transition kernel and the copula-GPD simulator—suitable for risk-aware hydroclimate operations including reservoir rule-curve testing, drought early-warning, and flood preparedness.

4.4 Robustness Limitations and validation pathways for improvement

655 Multiple, complementary diagnostics support the adequacy of each model layer. For emissions, PIT histograms, wet/dry reliability, QQ plots and GPD threshold stability confirm calibration across regimes, with expected challenges in the hyper-arid north (mass at zero, data scarcity in the tail). For spatial dependence, Rosenblatt-transform tests, variogram/tail-coincidence diagnostics, and variogram scores validate the selected vines; truncated structures are preferred when higher-order conditional links attenuate, limiting over-parameterization. Proper scoring rules (CRPS/ES/VS) under block cross-validation and out-of-sample evaluation document genuine predictive skill beyond in-sample fit. Sensitivity checks (vine families, thresholds, $K \in \{4, 6\}$) leave conclusions qualitatively unchanged, with Student- t pairs offering marginal gains for heavy-rain discrimination in southern The most pressing limitation is the inadequacy of zone-level areal-CRPS in climatically heterogeneous zones (2 and 4). The solution is threefold: replace the uniform areal average by an orographically weighted average; implement dedicated zero-inflated marginals for stations with very high dry-day frequencies; and apply stratified validation (separately for coastal and Andean sub-domains) before summarizing at the zone level. These modifications would not change the fundamental framework architecture but would extend its applicability to Chile's most climatically complex zones.

665 4.5 Limitations and scope.

First, Viterbi decoding induces state-assignment error ; residual misclassification Other limitations include Viterbi-induced state-assignment error that can propagate to the copula layer. In practice, effects are modest in our diagnostics, but a smoothed-probability simulation scheme could further attenuate classification noise. Second, within-state nonstationarity (e.g., slow structural changes in spatial dependence) ; within-state nonstationarity that is not explicitly modeled; regime-varying or covariate-varying vines, or mixture copulas, could capture such evolution. Third, tail modeling relies on a fixed exceedance threshold (here 0.90), which trades bias and variance; threshold ensembles or Bayesian threshold selection could formalize this choice. Fourth, the nHMM assumes conditional independence during estimation; although our hierarchy repairs dependence post hoc captured; and the modular estimation strategy that assumes conditional spatial independence during EM, though the hierarchical design repairs this in a principled manner, a fully joint spatiotemporal likelihood would be ideal but is computationally prohibitive at this dimensionality.

4.5 Pathways for future work.

675 Methodological extensions include (i) zero-inflated/hurdle emissions with regime-specific tails to improve arid-zone behavior; (ii) and diagnostically verified manner. Future work should explore hierarchical Bayesian formulations that integrate regime, vine, and tail components with shared priors, enabling joint uncertainty quantification and model selection; (iii) regime-aware multivariate fields (precipitation-temperature-snow) for compound-risk ; regime-aware multivariate fields for compound-risk analysis; and (iv) assimilation of sub-seasonal/seasonal assimilation of sub-seasonal dynamical predictors (e.g., MJO /jet MJO diagnostics, RCM output) as covariates in the transition kernel to extend lead times. From an applications perspective, state-contingent ensembles support reservoir rule-curve testing, drought early-warning, and impact-model forcing with coherent spatial fields.

680 , VIMF forecasts) as transition-model covariates.

4.5 Concluding remarks.

By disentangling large-scale large-scale regime dynamics from residual spatial dependence and explicitly accounting for extremes, the nHMM-DTW-vine-GPD framework delivers (i) nHMM-DTW-vine-GPD framework delivers meteorologically coherent regimes,

685 (ii) quantified and interpretable quantified changes in regime persistence , and (iii) consistent with observed megadrought trends, and calibrated, sharp , spatially consistent probabilistic forecasts probabilistic forecasts in the rainfall-active zones of Chile. The approach reconciles physical insight with statistical parsimony and computational tractability at continental scale, offering a transferable template for probabilistic hydroclimate modeling modelling in Chile and other mountainous, hydroclimatically heterogeneous regions, with clearly specified conditions for its applicability in arid and orographically complex environments.

690 *Code and data availability.* All code and data required to reproduce the analyses and figures are openly available at <https://github.com/mauricio-herrera/hybrid-nhmm-copula-precipitation-framework>. The repository contains the R script `Hybrid_nHMM_Copula_Simulation_Valid`, the compressed dataset `mapa_estaciones_precipitaciones_clean.csv.zip`, and documentation files (`README.md`, `LICENSE`, `CITATION.cff`). All analyses used R 4.3 or later with the packages `VineCopula`, `evd`, and `scoringRules`.

Author contributions. Mauricio Herrera-Marín: Conceptualization, Methodology, Formal Analysis, Investigation, Writing – Original Draft,

695 Visualization. Francisca Kleisinger: Data Curation, Methodology, Software, Visualization. Diego Rivera: Conceptualization, Writing – Review & Editing, Supervision. Andrés Wilson: Data Curation, Methodology, Visualization. Alex Godoy-Faúndez: Conceptualization, Writing – Review & Editing, Supervision.

Competing interests. The authors declare no competing interests.

Acknowledgements. D. Rivera acknowledges financial support from FONDECYT 1230520 and ANID/FONDAP/1523A0001.

- Aas, K., Czado, C., Frigessi, A., and Bakken, H.: Pair-copula constructions of multiple dependence, *Insurance: Mathematics and Economics*, 44, 182–198, 2009.
- Ailliot, P., Thompson, C., and Thomson, P.: Space–Time Modeling of Precipitation Using a Hidden Markov Model and Censored Gaussian Distributions, *Journal of Hydrology*, 365, 56–73, <https://doi.org/10.1016/j.jhydrol.2008.11.027>, 2009.
- 705 Arias, P. A., Garreaud, R., Poveda, G., Espinoza, J. C., Molina-Carpio, J., Masiokas, M., Viale, M., Scaff, L., and Van Oevelen, P. J.: Hydroclimate of the Andes part II: Hydroclimate variability and sub-continental patterns, *Frontiers in Earth Science*, 8, 505 467, 2021.
- Barrett, B. S. and Hameed, S.: Seasonal variability in precipitation in central and southern Chile: Modulation by the South Pacific high, *Journal of Climate*, 30, 55–69, 2017.
- Barría, P., Rojas, M., Moraga, P., Muñoz, A., Bozkurt, D., and Alvarez-Garreton, C.: Anthropocene and streamflow: Long-term perspective
710 of streamflow variability and water rights, *Elem Sci Anth*, 7, 2, 2019.
- Becker, E. and Schmitz, G.: Climatological effects of orography and land–sea heating contrasts on the gravity wave–driven circulation of the mesosphere, *J. Atmos. Sci.*, 60, 103–118, [https://doi.org/10.1175/1520-0469\(2003\)060<0103:CEOOAL>2.0.CO;2](https://doi.org/10.1175/1520-0469(2003)060<0103:CEOOAL>2.0.CO;2), 2003.
- Bellone, E., Hughes, J. P., and Guttorp, P.: A Hidden Markov Model for Downscaling Synoptic Atmospheric Patterns to Precipitation Amounts, *Climate Research*, 15, 1–12, <https://doi.org/10.3354/cr015001>, 2000.
- 715 Benjamini, Y. and Hochberg, Y.: Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing, *Journal of the Royal Statistical Society: Series B*, 57, 289–300, 1995.
- Bergmeir, C. and Benítez, J. M.: On the use of cross-validation for time series predictor evaluation, *Information Sciences*, 191, 192–213, <https://doi.org/10.1016/j.ins.2011.12.028>, 2012.
- Berndt, D. J. and Clifford, J.: Using dynamic time warping to find patterns in time series, *KDD workshop*, 10, 359–370, 1994.
- 720 Boisier, J. P.: CR2MET v2.5: High-resolution precipitation and temperature dataset for 1960–2021 in continental Chile, Dataset, <https://doi.org/10.5281/zenodo.7529682>, data set, 2023.
- Boisier, J. P., Alvarez-Garretón, C., Cepeda, J., Osses, A., Vásquez, N., and Rondanelli, R.: CR2MET A high-resolution precipitation and temperature dataset for hydroclimatic research in Chile, 2018.
- Boisier, J. P., Alvarez-Garreton, C., Cordero, R. R., Damiani, A., Gallardo, L., Garreaud, R. D., Lambert, F., Ramallo, C., Rojas, M., and
725 Rondanelli, R.: Anthropogenic drying in central-southern Chile evidenced by long-term observations and climate model simulations, *Elementa*, 6, <https://doi.org/10.1525/elementa.328>, 2018.
- Bozkurt, D., Rondanelli, R., Gonzalez-Reyes, A., Garreaud, R. D., Muñoz, A., Flores-Aqueveque, V., Aguilera, D., Farias-Barahona, D., Lambert, F., and Ramallo, C.: Climate change impacts on water availability and agriculture in Central Chile, *International Journal of Climatology*, 41, 1–21, <https://doi.org/10.1002/joc.6789>, 2021.
- 730 Burger, F., Brock, B., and Montecinos, A.: Seasonal and elevational contrasts in temperature trends in Central Chile between 1979 and 2015, *Global and Planetary Change*, 162, 136–147, 2018.
- Bárdossy, A. and Pegram, G.: Copula-Based Multivariate Distribution Functions for Hydrological Modeling, *Water Resources Research*, 46, W07 536, <https://doi.org/10.1029/2009WR007857>, 2010.
- Cavalleri, F., Lussana, C., Viterbo, F., Brunetti, M., Bonanno, R., Manara, V., Lacavalla, M., Sperati, S., Raffa, M., Capecchi, V., Cesari,
735 D., Giordani, A., Cerenzia, I. M. L., and Maugeri, M.: Multi-scale assessment of high-resolution reanalysis precipitation fields over Italy, *Atmospheric Research*, 312, 107 734, <https://doi.org/https://doi.org/10.1016/j.atmosres.2024.107734>, 2024.

- Charles, S. P., Timbal, B., and Nguyen, K.: Statistical Downscaling for Assessing Simulated Precipitation over Western Victoria, Australia, *International Journal of Climatology*, 34, 719–732, <https://doi.org/10.1002/joc.3727>, 2014.
- Cioffi, F., Conticello, F. R., and Lall, U.: Stochastic Scenarios for 21st Century Rainfall Seasonality, Daily Frequency, and Intensity in South
740 Florida, *Journal of Water Resources Planning and Management*, 146, 04020 058, [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001250](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001250),
2020.
- CR2: Centro de Ciencia del Clima y la Resiliencia. Datos de Precipitación, <https://www.cr2.cl/datos-de-precipitacion/>, accessed: 2024-05-29,
2024.
- Dawid, A. P.: Present Position and Potential Developments: Some Personal Views: Statistical Theory: The Prequential Approach, *Journal of*
745 *the Royal Statistical Society. Series A*, 147, 278–292, <https://doi.org/10.2307/2981683>, 1984.
- Diebold, F. X. and Mariano, R. S.: Comparing Predictive Accuracy, *Journal of Business & Economic Statistics*, 13, 253–263,
<https://doi.org/10.1080/07350015.1995.10524599>, 1995.
- Franzke, C. L. E. and Osprey, S.: Climate variability and extremes: Statistical Erameworks and Applications, *Proceedings of the Royal*
Society A: Mathematical, Physical and Engineering Sciences, 470, 20130 267, <https://doi.org/10.1098/rspa.2013.0267>, 2014.
- 750 Garreaud, R., Alvarez-Garretón, C., Barichivich, J., Boisier, J., Christie, D., Galleguillos, M., LeQuesne, C., McPhee, J., and Zambrano-
Bigiarini, M.: The 2010–2015 megadrought in central Chile: impacts on regional hydroclimate and vegetation, *Hydrology and Earth*
System Sciences, 21, 6307–6327, <https://doi.org/10.5194/hess-21-6307-2017>, 2017.
- Garreaud, R., Alvarez-Garretón, C., Barichivich, J., and et al.: The 2010–2018 Megadrought in Central Chile: Causes, Impacts and Lessons,
Journal of Climate, 33, 5765–5780, 2020.
- 755 Garreaud, R. D.: The Andes climate and weather, *Advances in geosciences*, 22, 3–11, 2009.
- Garreaud, R. D., Vuille, M., Compagnucci, R., and Marengo, J.: Present-day South American climate, *Palaeogeography, Palaeoclimatology,*
Palaeoecology, 281, 180–195, 2009.
- Genest, C. and MacKay, J.: A primer on the copula approach to modeling dependence between random variables, *Statistical Science*, 7,
113–122, 1986.
- 760 Genest, C. and Rémillard, B.: Validity of the Parametric Bootstrap for Goodness-of-Fit Testing in Semiparametric Models, *Annales de*
l’Institut Henri Poincaré, Probabilités et Statistiques, 44, 1096–1127, <https://doi.org/10.1214/07-AIHP148>, 2008.
- Genest, C., Rémillard, B., and Beaudoin, D.: Goodness-of-Fit Tests for Copulas: A Review and a Power Study, *Insurance: Mathematics and*
Economics, 44, 199–213, <https://doi.org/10.1016/j.insmatheco.2007.10.005>, 2009.
- Giorgino, T.: Computing and Visualizing Dynamic Time Warping Alignments in R: The dtw Package, *Journal of Statistical Software*, 31,
765 1–24, <https://doi.org/10.18637/jss.v031.i07>, 2009.
- Gneiting, T. and Raftery, A. E.: Strictly Proper Scoring Rules, Prediction, and Estimation, *Journal of the American Statistical Association*,
102, 359–378, <https://doi.org/10.1198/016214506000001437>, 2007.
- Herrera, M., Kleisinger, F., and Wilsón, A.: Spatiotemporal Forecasting in Climate Data Using EOFs and Machine Learning Models: A Case
Study in Chile, <https://arxiv.org/abs/2502.17495>, 2025.
- 770 Hersbach, H., Bell, B., and et. al.: ERA5 hourly data on single levels from 1940 to present., Copernicus Climate Change Service (C3S)
Climate Data Store (CDS), <https://doi.org/10.24381/cds.adbb2d47>, 2023.
- Hoskins, B. J. and Hodges, K. I.: A new perspective on Southern Hemisphere storm tracks, *Journal of Climate*, 18, 4108–4129, 2005.
- Hughes, J. P., Guttorp, P., and Charles, S. P.: A non-homogeneous hidden Markov model for precipitation occurrence, *Journal of the Royal*
Statistical Society: Series C (Applied Statistics), 48, 15–30, 1999.

- 775 Khalil, A. F., Kwon, H.-H., Lall, U., and Kaheil, Y. H.: Predictive Downscaling Based on Non-Homogeneous Hidden Markov Models, *Hydrological Sciences Journal*, 55, 333–350, <https://doi.org/10.1080/02626661003780342>, 2010.
- Kirshner, S.: Modeling of Multivariate Time Series Using Hidden Markov Models, Ph.d. thesis, University of California, Information and Computer Sciences, 2005.
- Latoja, D., Lillo-Saavedra, M., Gonzalo-Martin, C., Godoy-Faúndez, A., Somos-Valenzuela, M., and Rivera, D.: Decadal Variability of Dry
780 Days in Central Chile, *Earth Systems and Environment*, pp. 1–18, 2024.
- Maharaj, E. A. and Alonso, A. M.: Discriminant analysis of multivariate time series: Application to diagnosis based on ECG signals, *Computational Statistics & Data Analysis*, 55, 973–988, <https://doi.org/10.1016/j.csda.2010.07.015>, 2011.
- Mayer, J., Mayer, M., and Haimberger, L.: Consistency and Homogeneity of Atmospheric Energy, Moisture, and Mass Budgets in ERA5, *Journal of Climate*, 34, 3955–3974, <https://doi.org/10.1175/JCLI-D-20-0676.1>, publisher Copyright: © 2021 American Meteorological
785 Society., 2021.
- McLachlan, G. and Krishnan, T.: The EM Algorithm and Extensions, vol. 382, John Wiley & Sons, Hoboken, NJ, 2nd edn., 2008.
- Montecinos, A. and Aceituno, P.: Seasonality of the ENSO-related rainfall variability in central Chile and associated circulation anomalies, *Journal of climate*, 16, 281–296, 2003.
- Montecinos, A., Kurgansky, M. V., Muñoz, C., and Takahashi, K.: Non-ENSO interannual rainfall variability in central Chile during austral
790 winter, *Theoretical and Applied Climatology*, 106, 557–568, <https://doi.org/10.1007/s00704-011-0457-1>, 2011.
- Montero, P. and Vilar, J. A.: TSclust: An R package for time series clustering, *Journal of Statistical Software*, 62, 1–43, <https://doi.org/10.18637/jss.v062.i01>, 2015.
- Nelsen, R. B.: An Introduction to Copulas, Springer Series in Statistics, Springer, New York, 2 edn., <https://doi.org/10.1007/0-387-28678-0>, 2006.
- 795 Pablo Sarricolea, M. H.-O. and Óliver Meseguer-Ruiz: Climatic regionalisation of continental Chile, *Journal of Maps*, 13, 66–73, <https://doi.org/10.1080/17445647.2016.1259592>, 2017.
- Pica-Téllez, A., Garreaud, R., Meza, F., Bustos, S., Falvey, M., Ibarra, M., Duarte, K., Ormazábal, R., Dittborn, R., and Silva, I.: Informe proyecto arclim: Atlas de riesgos climáticos para chile, Tech. rep., Centro de Ciencia del Clima y la Resiliencia, Centro de Cambio Global UC y Meteodata para el Ministerio del Medio Ambiente a través de la Deutsche Gesellschaft für Internationale Zusammenarbeit (giz),
800 Santiago, Chile, 2020.
- Politis, D. N. and Romano, J. P.: The Stationary Bootstrap, *Journal of the American Statistical Association*, 89, 1303–1313, <https://doi.org/10.1080/01621459.1994.10476870>, 1994.
- Pyper, B. J. and Peterman, R. M.: Comparison of methods to account for autocorrelation in correlation analyses of fish data, *Canadian Journal of Fisheries and Aquatic Sciences*, 55, 2127–2140, <https://doi.org/10.1139/f98-104>, 1998.
- 805 Robertson, A. W., Kirshner, S., and Smyth, P.: Downscaling of Daily Rainfall Using a Hidden Markov Model, *Journal of Climate*, 17, 4407–4429, <https://doi.org/10.1175/3190.1>, 2004a.
- Robertson, A. W., Kirshner, S., and Smyth, P.: Nonstationary hidden Markov models for modeling daily rainfall occurrence over Brazil, *Journal of Climate*, 17, 4407–4424, 2004b.
- Robertson, A. W., Kirshner, S., Smyth, P., Charles, S. P., and Bates, B. C.: Subseasonal-to-interdecadal variability of the Australian monsoon
810 over North Queensland, *Quarterly Journal of the Royal Meteorological Society*, 132, 519–542, 2006.
- Rosenblatt, M.: Remarks on a Multivariate Transformation, *Annals of Mathematical Statistics*, 23, 470–472, <https://doi.org/10.1214/aoms/1177729394>, 1952.

- Salazar, A., Thatcher, M., Goubanova, K., et al.: CMIP6 precipitation and temperature projections for Chile, *Climate Dynamics*, 62, 2475–2498, <https://doi.org/10.1007/s00382-023-07034-9>, 2024.
- 815 Sarricolea, P., Herrera-Ossandon, M., and Meseguer-Ruiz, Ó.: Climatic regionalisation of continental Chile, *Journal of Maps*, 13, 66–73, 2017.
- Schepsmeier, U., Stöber, J., Brechmann, E. C., Graeler, B., Nagler, T., Erhardt, T., Almeida, C., Min, A., and Ruckdeschel, P.: Package ‘VineCopula’, R Package Version 1.6.0, 2015.
- Scheuerer, M. and Hamill, T. M.: Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities, *Monthly Weather*
820 *Review*, 143, 1321–1334, <https://doi.org/10.1175/MWR-D-14-00269.1>, 2015.
- Székely, G. J. and Rizzo, M. L.: Energy Statistics: A Class of Statistics Based on Distances, *Statistical Science*, 28, 558–575, <https://doi.org/10.1214/13-STS437>, 2013.
- Trenberth, K. E.: Changes in Precipitation with Climate Change, *Climate Research*, 47, 123–138, 2014.
- Valdés-Pineda, R., Valdes, J. B., Diaz, H. F., and Pizarro-Tapia, R.: Analysis of spatio-temporal changes in annual and seasonal precipitation
825 variability in South America-Chile and related ocean–atmosphere circulation patterns, *International Journal of Climatology*, 36, 2979–3001, 2016.
- Vrac, M., Naveau, P., Li, Z., Noblet-Ducoudré, N., Drobinski, P., Somot, S., Cattiaux, J., and Yiou, P.: Bridging Climate Gaps: A Stochastic Downscaling Method to Link Large-Scale Variables with Local Climate, *Journal of Climate*, 25, 6447–6464, <https://doi.org/10.1175/JCLI-D-11-00628.1>, 2012.
- 830 Vuong, Q. H.: Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses, *Econometrica*, 57, 307–333, <https://doi.org/10.2307/1912557>, 1989.
- Ward, J. H.: Hierarchical Grouping to Optimize an Objective Function, *Journal of the American Statistical Association*, 58, 236–244, <https://doi.org/10.1080/01621459.1963.10500845>, 1963.
- Wilks, D. S.: *Statistical Methods in the Atmospheric Sciences*, Academic Press, 3rd edn., 2011.
- 835 Zhang, L., Chen, X., and Li, S.: Application of Non-Homogeneous Hidden Markov Model to Simulate Rainfall Sequences in a Changing Climate, *Hydrology and Earth System Sciences*, 25, 3925–3941, <https://doi.org/10.5194/hess-25-3925-2021>, 2021.
- Zucchini, W., MacDonald, I. L., and Langrock, R.: *Hidden Markov Models for Time Series: An Introduction Using R*, Chapman & Hall/CRC, 2nd edn., 2016.