POLITECNICO DI TORINO

Master's Degree in Mathematical Engineering

Master's Degree Thesis

Statistical modelling of non-stationary extreme winds under anthropic land consumption and climate change

Supervisors

Luca BRUNO

Gianluca MASTRANTONIO

Lorenzo RAFFAELE

Candidate

Elisabetta PAUTASSO

July 2026

Abstract

The probabilistic analysis of extreme wind speed grounds the wind hazard assessment and the design of safe structures under wind loads. In wind engineering, this usually translates into predicting the design wind speed having a specified return period, a task that traditionally relies on Extreme Value Analysis (EVA) under stationarity assumptions. However, this paradigm has been challenged within the last decades by land surface anthropogenic modifications and climate change.

To overcome the above modelling assumption, in this MSc Thesis a specific framework based on Bayesian inference, specifically a Markov Chain Monte Carlo algorithm for Non-stationary Extreme Value Analysis (NEVA), is adopted for the analysis of extreme wind speeds. Such an approach is tested on both synthetic and real-world time series.

First, the model is applied to synthetic time series to critically discuss its robustness under forced model misspecification, and its technical applicability to time series of different lengths.

Then, the model is applied to historical wind speed time series recorded at real world anemometric stations. Their neighbourhoods are affected by changes of the land surface aerodynamic roughness due to increasing urbanisation. The model proves encouraging ability to capture the non-stationary statistical features of the available wind speed data.

Finally, the model is implemented in the estimate of future climatic projections according to the Representative Concentration Pathways (RCP) adopted by the Intergovernmental Panel on Climate Change (IPCC). Both RCP 4.5 and 8.5 emission scenarios are considered, and related wind fields obtained from the convection-permitting Very High Resolution simulations made available by the Euro-Mediterranean Center on Climate Change (CMCC).

The MSc Thesis provides unprecedented findings in literature, namely by mapping future trends of the 50-year Effective Return Levels across 156 spatial points over and around Italy. The maps provide a spatially coarse but comprehensive overview of future wind hazards, and allow to identify the regions subjected to increase or decrease of extreme wind events.

Acknowledgements

I would like to express my sincere gratitude to all my supervisors for their precious time, patience, and valuable advice. Thank you for giving me the opportunity to work on this project, and for sharing your deep passion for this field of research.

Table of Contents

List of Tables VI

List of Figures VII

1 Introduction 1

2 State of the Art 4
- 2.1 Scope and Boundaries of the Review . . . 4
- 2.2 Stationary Extreme Value Analysis . . . 4
  - 2.2.1 The Block Maxima Approach and the GEV Distribution . . 5
  - 2.2.2 The Peaks-Over-Threshold Approach and the GP Distribution 8
  - 2.2.3 Rates of Convergence: Ultimate and Penultimate Approximations . . . 9
- 2.3 Non-Stationary Extreme Value Analysis . . . 12
  - 2.3.1 Maximum Likelihood Estimation (MLE) . . . 13
  - 2.3.2 Bayesian Inference . . . 13
- 2.4 NEVA Application to Climate Variables . . . 15
  - 2.4.1 Non stationary Wind Hazard . . . 15
  - 2.4.2 Other non stationary climate hazards . . . 16

3 Statistical model 18
- 3.1 Formulation of the Extreme Value Models . . . 20
  - 3.1.1 The Stationary Baseline . . . 20
  - 3.1.2 The Non-Stationary Extension . . . 20
- 3.2 Bayesian Inference and Hybrid MCMC Algorithm . . . 23
  - 3.2.1 Model Diagnostic and Convergence . . . 26
  - 3.2.2 Hazard prediction . . . 28

4 Application to synthetic wind speed time series 29
- 4.1 Generation of Synthetic Time Series . . . 29
- 4.2 Demonstrative Application of EVA to a Stationary Random Variable 31

4.3 Application of NEVA to stationary and non-stationary Random Variables . . . 34
4.3.1 Stationary time series . . . 34
4.3.2 Non-Stationary Time Series . . . 37

5 Application to real world wind speed time series 44
5.1 Time series measured at anemometric stations under changing environmental conditions . . . 44
5.2 Time series simulated over Italy under climate change . . . 56

6 Conclusions and Future Perspectives 65

Bibliography 67

List of Tables

4.1 Summary of the temporal evolution combinations for the mean and variance of the *parent* distribution. The cases analysed in this chapter are marked with an **X** 30
4.2 Comparison of parameter estimates and Goodness-of-Fit (GoF) metrics for the 2000-year stationary synthetic time series. The parameters are estimated via MLE and MCMC, alongside the analytical formulas. For the GP distribution, the location parameter refers to the fixed threshold u. 32
4.3 Summary of Non-Stationary GPD Parameters, and Goodness-of-Fit Metrics, specifically for the test case featuring a linear mean and quadratic variance. 39
4.4 Summary of Non-Stationary GPD Parameters, and Goodness-of-Fit Metrics, specifically for the test case featuring a constant mean and quadratic variance. 43

5.1 Mann-Kendall (MK) and White test results (p-value) and percentage variation of the built surface (Δr) for the analysed anemometric stations. The statistical tests refer to the 1981-2022 period, while Δr is calculated for the 1991-2011 period. 46
5.2 Estimated Non-Stationary GPD Parameters (MAP) and Diagnostic Metrics, for the time series measured. 48
5.3 Results of the Mann-Kendall (monotonic trend) and White (heteroscedasticity) non-stationarity tests for the RCP 4.5 and 8.5 projections. For each test, the p-value and test decision regarding the null hypothesis H_0 are reported at a significance level of $\alpha = 0.05$. . 57
5.4 Estimated GPD parameters (MAP) and performance metrics for each scenario. 57
5.5 Summary of CPU execution times for the non-stationary extreme value analysis under RCP4.5 and RCP8.5 scenarios across a 156-point grid over Italy. 61

List of Figures

3.1 Overall architecture and workflow of the ProNEVA software package. 19
3.2 Flowchart of the Hybrid MCMC algorithm 25

4.1 Synthetic time series generated from a stationary random variable on a 2000-year projection. Left: parent (Weibull) and GEV distributions. Middle: time series with annual maxima in red and peaks over threshold in blue. Right: GP distribution. 32
4.2 Quantile-Quantile (Q-Q) plots comparing the empirical quantiles against the theoretical estimates (MLE and MCMC). Left: GEV distribution. Right: GP distribution. 33
4.3 Non-stationary MCMC framework results for the GEV model (top row) and the GP model (bottom row). From left to right, the panels illustrate: the predictive probability density functions (pdf), the effective return levels (ERL_T) for T = 2, 50, 100, and 200 years. . . 35
4.4 Non-stationary MCMC framework results for the GEV model (top row) and the GP model (bottom row). From left to right, the panels illustrate: the temporal evolution of the normalised distribution parameters (location and scale for the GEV; constant threshold and scale for the GPD), and the diagnostic Q-Q plots. 36
4.5 Convergence analysis of the parameter and ERL_{50} slopes for the GEV and GP distributions against sampled years (top) and theoretical sample size (bottom). 38
4.6 Non-stationary MCMC framework results for the GPD model against analytical framework, specifically for the test case featuring a linear mean and quadratic variance. Top row: from left to right the evolution of the GPD from year 0 to year 100 and the effective return level. Bottom row: from left to right, the panels illustrate the temporal evolution of the normalised distribution parameters (dynamic threshold and scale), and the diagnostic Q-Q plot. 40

4.7 Relative errors of the estimated parameters and effective return levels against time, specifically for the test case featuring a linear mean and quadratic variance. 41
4.8 Non-stationary MCMC framework results for the GPD model against analytical framework, specifically for the test case featuring a constant mean and quadratic variance. Top row: from left to right the evolution of the GPD from year 0 to year 100 and the effective return level. Bottom row: from left to right, the panels illustrate the temporal evolution of the normalised distribution parameters (dynamic threshold and scale), and the diagnostic Q-Q plot. 42
4.9 Relative errors of the estimated parameters and effective return levels against time, specifically for the test case featuring a constant mean and quadratic variance. 43

5.1 View of the land use around Ciampino and Catania in 1984 and 2024 (satellite imagery Landsat / Copernicus). 47
5.2 Closed view (2 km radius fetch) of the land use around the anemometric stations at Roma Ciampino and Catania Fontanarossa in 1984 and 2024 (satellite imagery Landsat / Copernicus). Wind roses after anemometric records of U_{10} at the stations from Jan 1981 to Dec 2022. 49
5.3 Trend of the population and of the percentage of built surface over total one at Catania and Ciampino from 1970 to 2020 (data source: Italian National Institute of Statistics - ISTAT). 50
5.4 Results of the non-stationary Generalized Pareto Distribution (GPD) analysis for Catania Fontanarossa. The top panel displays the temporal evolution of the Probability Density Function (PDF) and Effective Return Level (ERL), while the bottom panel illustrates the parameter evolution and the corresponding Q-Q plot. 51
5.5 Results of the non-stationary Generalized Pareto Distribution (GPD) analysis for Rome Ciampino. The top panel displays the temporal evolution of the Probability Density Function (PDF) and Effective Return Level (ERL), while the bottom panel illustrates the parameter evolution and the corresponding Q-Q plot. 52
5.6 Trend of the ERL_{50} and built surface percentage (r) in the top panels (Catania Fontanarossa on the left, Roma Ciampino on the right). Scatter plot of ERL_{50} versus built surface percentage in the bottom panels. 53
5.7 Bubble map of the 50-year effective return level (ERL_{50}) trend for the selected anemometric stations. 54

5.8 Bubble map of the 50-year effective return level (ERL_{50}) percentage for the selected anemometric stations. 55
5.9 Comprehensive non-stationary extreme value analysis comparing RCP 4.5 (left column) and RCP 8.5 (right column) emission scenarios, for the location 39.5°N, 9.5°E (South Sardinia). 58
5.10 Comprehensive non-stationary extreme value analysis comparing RCP 4.5 (left column) and RCP 8.5 (right column) emission scenarios, for the location 45.5°N, 12.5° E (Venice). 59
5.11 Bubble map of the 50-year effective return level (ERL_{50}) trend under the RCP 4.5 scenario. 61
5.12 Bubble map of the 50-year effective return level (ERL_{50}) trend under the RCP 8.5 scenario. 62
5.13 Bubble map of the 50-year effective return level (ERL_{50}) percentage trend under the RCP 4.5 scenario. 63
5.14 Bubble map of the 50-year effective return level (ERL_{50}) percentage trend under the RCP 8.5 scenario. 64

Chapter 1

Introduction

In civil and structural engineering, wind action is one of the most critical environmental forces governing the design of buildings, bridges, and towers. An accurate estimation is fundamental in order to avoid catastrophic structural failures. Underestimating extreme wind speed can lead to severe structural damage or collapse during severe storm events, resulting in unacceptable risks to human life and heavy economic losses. On the other hand, overestimating the wind speed means a dramatic increase in material usage and construction costs. Therefore, defining an accurate "design wind speed" for a specific construction site is a crucial, preliminary step, entrusted usually to the designer through the application of probabilistic return criteria, dictated by national and international building codes.

Traditionally, to evaluate the design wind speed, structural engineering relies on the Extreme Value Theory (EVT). The core idea is to model the tail behaviour of a probability distribution, to predict the frequency and intensity of events that are more extreme than any previously observed. Typically, there are two main approaches: Block Maxima (BM) method, where for each time interval, e.g. every years, the maximum value is extracted, and Peaks-Over-Threshold (POT) method, where all the values above a fixed threshold are considered. The first approach relies on the Generalized Extreme Value (GEV) distribution, while the second one on the Generalized Pareto (GP) distribution. Both approaches are based on a fundamental statistical assumption: stationarity. In this way, the physical process governing extreme winds is constant over the time, and so invariant in the future. Under this paradigm, engineers analyze decades of past anemometric measures to extrapolate future extremes, assuming that the future will be a statistical replication of the past.

The fundamental assumption on stationarity has been severely challenged in recent decades. In particular, environment is rapidly changing due to two main overlapping factors: anthropogenic climate change and rapid urbanization. On a global scale, the increase in greenhouse gas emission has led to a warming

climate, which alters large-scale atmospheric circulation patterns, temperature gradients, frequency and intensity of severe weather system. On a local scale, rapid urbanization alters significantly the surface roughness, since cities are growing vertically and the local aerodynamic exposure changes.

To better understand climate shifts and their impact on local wind hazards, the scientific community relies on hierarchical climate modelling framework. At the global scale, General Circulation Models (GCMs) are employed to simulate the complex, long-term dynamic of the Earth's atmosphere, oceans, and land surface. Their coarse spatial resolution (from 100 to 200 km) is insufficient to capture mesoscale phenomena that dictates extreme wind events, such as storms. Through a dynamical downscaling, Regional Climate Models (RCMs) are obtained and resolve the climate dynamics at a finer spatial grid (from 2 to 12 km, as seen in the Euro-CORDEX framework). This high-resolution downscaling is crucial for structural wind engineering, as it allows for a more realistic representation of local surface roughness and orographic effects. Furthermore, the prediction of future climate has to take into account the uncertainties associated with future human activities and policies. These uncertainties have been formalized through Representative Concentration Pathways (RCPs), such as RCP 4.5 (an intermediate stabilization scenario) and RCP 8.5 (a high-emission, worst case scenario). By forcing the RCMs with these RCPs, researchers can generate rigorous projections of future meteorological variables, providing the necessary physical foundation for non-stationary statistical modelling.

To bridge the gap between a changing climate and structural reliability, classical EVT must be adapted into a Non-Stationary Extreme Value Analysis (NEVA) framework. In a non-stationary context, the parameters governing the extreme value distributions are no longer treated as fixed scalar, but they can vary with the time or with specific physical and meteorological covariates, such as temperature or surface roughness.

The introduction of a NEVA framework introduces some challenges. While stationary models rely on Maximum Likelihood Estimation (MLE) for parameter fitting, this is no longer possible in a non-stationary approach. In this way, the Bayesian inference framework is seen as a possible solution. By employing advanced computational techniques like Markov Chain Monte Carlo (MCMC) algorithms, Bayesian frameworks allow researchers to combine prior knowledge with observational data to generate posterior distributions of the extreme value parameters. In this way, robust estimations and uncertainty bounds can be performed.

Despite the widespread applications of NEVA in environmental fields, such as coastal engineering, hidrology for extreme precipitation or temperature, its application to structural wind engineering is relatively unexplored. In this way,

the primary objective of this thesis is to address this gap by developing , implementing and validating a statistical modelling framework for non-stationary wind hazard assessment. By leveraging advanced Bayesian inference techniques and non-stationary probability distributions, this research aims to provide a robust methodology to evaluate wind hazard for civil structures under changing climatic conditions. To achieve these objectives, the present work is organised as follows:

- Chapter 2 (State of the Art) provides a comprehensive review of the existing literature. It explores the transition from classical stationary extreme value theory to non-stationary applications.
- Chapter 3 (Statistical Model) details the mathematical and theoretical framework developed for this study.
- Chapter 4 (Application to Synthetic Wind Speed Time Series) presents the validation of the proposed statistical model using synthetically generated data. By testing the framework against controlled time series with known underlying non-stationary trends, the accuracy, robustness, and limitations of the Bayesian inference approach are evaluated.
- Chapter 5 (Application to Real World Wind Speed Time Series) applies the validated methodology to historical anemometric measurements and future climate projections across Italian Peninsula, to estimate past and future non-stationary return levels for wind speeds.
- Chapter 6 (Conclusions and Future Perspectives) summarizes the key findings of the research, discusses the engineering implications, and outlines potential avenues for future research in the field of climate-aware structural engineering.

Chapter 2

State of the Art

2.1 Scope and Boundaries of the Review

As outlined in the Chapter 1, evaluating wind hazards under changing climatic conditions is a deeply multidisciplinary challenge. The scientific literature surrounding this topic encompasses different fields, ranging from atmospheric physics and physical climate modelling to computational fluid dynamics (CFD) and structural vulnerability assessment. It is necessary to clearly define the boundaries of the present literature review. The focal point of this review is strictly methodological and statistical. The following sections are dedicated to tracing the evolution of extreme value modelling in engineering practice.

2.2 Stationary Extreme Value Analysis

To predict the occurrence and magnitude of extreme events, traditional statistical methods—which primarily focus on the central tendency of data—are generally inadequate. Instead, a robust theoretical framework capable of modelling the tail behaviour of probability distributions is required. It is within this context that Extreme Value Analysis (EVA) emerged as the standard approach. As comprehensively detailed in the milestone works by Coles [1] and Beirlant et al. [2], classical EVA is deeply rooted in the fundamental assumption of stationarity, where the physical processes generating the extreme events are assumed to remain constant over time. Within this stationary paradigm, the statistical inference of extremes is traditionally carried out through two primary methodological approaches: the Block Maxima (BM) method and the Peaks-Over-Threshold (POT) method.

2.2.1 The Block Maxima Approach and the GEV Distribution

In classical statistics, the asymptotic behaviour of sample means is governed by the Central Limit Theorem (CLT). The CLT states that the distribution of properly normalized sums of independent and identically distributed (i.i.d.) random variables converges to a Normal distribution, regardless of the underlying parent distribution. However, when the focus of the analysis shifts from the average behaviour to the tails of the distribution, a different asymptotic framework is required.

Let $X_1, X_2, \ldots, X_n$ be a sequence of i.i.d. random variables (e.g., daily wind speeds) with a common, yet often unknown, cumulative distribution function (CDF) F. The sample maximum is defined as $M_n = \max(X_1, \ldots, X_n)$, with the exact probability distribution of the maximum given by $P(M_n \leq z) = F^n(z)$.

Since F is typically unknown in engineering applications, this formula is not useful in practice. Moreover, a small error in the estimation of F might lead to significant errors in F^n for large n. Therefore, it is necessary to look at the asymptotic behaviour of F^n as $n \to \infty$. Defining x_F as the upper endpoint of the distribution, i.e., $x_F = \sup\{x : F(x) < 1\}$, it follows that

$$\lim_{n\to\infty} F^n(z) = 0 \quad \forall z < x_F,$$

so the distribution of M_n degenerates to a point mass on x_F.

Similar to the logic of the CLT, to solve this problem, the sequence M_n must be linearly normalized. The goal is to find sequences of normalizing constants $\{a_n > 0\}$ and $\{b_n\}$ such that:

$$M_n^* = \frac{M_n - b_n}{a_n}. \tag{2.1}$$

The entire range of possible limit distributions for M_n^* is given by the *Extremal Types Theorem* (Fisher and Tippett, 1928), and can be formalized in a unified way by the *Generalized Extreme Value (GEV) Theorem* [1]:

Theorem 2.2.1 *If there exist sequences of constants* $\{a_n > 0\}$ *and* $\{b_n\}$ *such that*

$$P\left(\frac{M_n - b_n}{a_n} \leq z\right) \to G(z) \quad \text{as } n \to \infty, \tag{2.2}$$

where G *is a non-degenerate distribution function, then* G *is a member of the Generalized Extreme Value (GEV) family:*

$$G(z) = \exp\left\{-\left[1 + \xi\left(\frac{z-\mu}{\sigma}\right)\right]^{-1/\xi}\right\}, \tag{2.3}$$

defined on the set $\{z : 1 + \xi(z-\mu)/\sigma > 0\}$.

The GEV distribution is fully specified by three parameters:

- $\mu \in \mathbb{R}$ is the location parameter, representing the center of the extreme value distribution;
- $\sigma > 0$ is the scale parameter, denoting the spread or variability of the extremes;
- $\xi \in \mathbb{R}$ is the shape parameter, also referred to as the Extreme Value Index (EVI).

The tail behaviour is mainly characterised by the shape parameter ξ which determines three particular types of extreme value distributions based on their asymptotic decay rates. For $\xi > 0$ the distribution is in the Fréchet family (heavy-tailed) and has polynomial decay. If $\xi < 0$, it corresponds to the Weibull family (short-tailed), where the upper tail is bounded and the distribution converges to a finite maximum limit. The limit of (2.3) for $\xi \to 0$ provides the Gumbel family (light-tailed) with an exponential tail decay:

$$G(z) = \exp\left\{-\exp\left[-\left(\frac{z-\mu}{\sigma}\right)\right]\right\}, \quad -\infty < z < \infty. \tag{2.4}$$

To translate the probabilistic framework of the GEV distribution into practical engineering parameters, it is necessary to introduce the concepts of the quantile function and the tail quantile function.

For a given cumulative distribution function F, the quantile function $Q(p)$ is defined as the generalized inverse of F:

$$Q(p) = \inf\{x \in \mathbb{R} : F(x) \geq p\},$$

where $p \in (0,1)$ is a specified probability. In extreme value analysis, the focus is inherently on the upper tail of the distribution, where p approaches 1. To mathematically formalize the study of this specific region, as proposed by Beirlant et al. [2], it is convenient to introduce the tail quantile function $U(t)$, defined as:

$$U(t) = Q\left(1 - \frac{1}{t}\right),$$

for $t > 1$.

This mathematical formulation holds profound physical and engineering significance. By evaluating the tail quantile function at $t = T$, where T represents time in years, the function yields the T-year return level, z_T. From a physical perspective, the return period T is defined as the expected average time interval between consecutive occurrences of an event exceeding the magnitude z_T. Conversely, its reciprocal $p = 1/T$ represents the annual probability of exceedance.

For the GEV distribution $G(z)$, the return level z_T associated with a return period T can be explicitly calculated by evaluating its tail quantile function (i.e., inverting the CDF). As detailed by Coles [1], since the shape parameter ξ dictates the tail behaviour, the expression for z_T takes different forms depending on whether $\xi \neq 0$ (Fréchet or Weibull domains) or $\xi = 0$ (Gumbel domain):

$$z_T = \begin{cases} \mu - \frac{\sigma}{\xi}\left[1 - \left\{-\log\left(1 - \frac{1}{T}\right)\right\}^{-\xi}\right], & \text{for } \xi \neq 0 \\ \mu - \sigma \log\left\{-\log\left(1 - \frac{1}{T}\right)\right\}, & \text{for } \xi = 0 \end{cases} \tag{2.5}$$

In structural wind engineering, determining the correct return level is a step of enormous importance. The aerodynamic force exerted by the wind on a structure is proportional to the square of the wind velocity. Therefore, the return level z_T is universally adopted by international building codes as the "design wind speed" (e.g., $T = 50$ years for standard structures, corresponding to an annual exceedance probability of 2%). Any inaccuracy in modelling the tail quantile function results in an underestimation of the design wind loads, compromising the structural safety and increasing the risk of catastrophic failure under extreme storm events.

The Domain of Attraction Problem

When the normalized maxima of a sequence converge to a specific GEV distribution G governed by a shape parameter ξ, the parent distribution F is said to belong to the *maximum domain of attraction* of G, denoted by $F \in \mathcal{D}(G_\xi)$.

The conditions for a distribution to belong to a specific domain of attraction can be analytically verified through the von Mises conditions. Following the formalization by de Haan and Ferreira [3], the sufficient condition is based on the hazard function $r(x) = f(x)/[1 - F(x)]$.

Theorem 2.2.2 (von Mises Condition) *Suppose F has a positive and twice differentiable probability density function f in a left neighbourhood of its upper endpoint $x_F \leq \infty$. If the limit of the derivative of the reciprocal of the hazard function satisfies*

$$\lim_{x \uparrow x_F} \frac{d}{dx}\left[\frac{1 - F(x)}{f(x)}\right] = \lim_{x \uparrow x_F} \frac{d}{dx}\left[\frac{1}{r(x)}\right] = \xi, \tag{2.6}$$

then $F \in \mathcal{D}(G_\xi)$.

If the von Mises condition holds, standard EVA literature ([3]) provides explicit formulas to compute the normalizing constants a_n and b_n. The location parameter b_n is the $(1 - 1/n)$-quantile of F, while the scale parameter a_n is the reciprocal of

the hazard function evaluated at b_n:

$$b_n = F^{-1}\left(1 - \frac{1}{n}\right), \quad a_n = \frac{1}{r(b_n)} = \frac{1 - F(b_n)}{f(b_n)}. \tag{2.7}$$

In wind engineering, the parent distribution of the mean wind speed is traditionally modelled using the Weibull distribution, which is not the same Weibull as the Type III GEV. Let F be a Weibull parent distribution with scale parameter $A > 0$ and shape parameter $B > 0$:

$$F(x) = 1 - \exp\left[-\left(\frac{x}{A}\right)^B\right], \quad f(x) = \frac{B}{A}\left(\frac{x}{A}\right)^{B-1} \exp\left[-\left(\frac{x}{A}\right)^B\right], \tag{2.8}$$

with $x_F = \infty$. Evaluating the von Mises condition yields $\frac{d}{dx}[1/r(x)] = \frac{A^B}{B}(1-B)x^{-B}$. Since $B > 0$, as $x \to \infty$ the limit is exactly 0, firmly placing the Weibull distribution in the Gumbel domain ($\xi = 0$). Applying eq. (2.7), the theoretical parameters for the limiting Gumbel distribution of the block maxima are:

$$b_n = A(\ln n)^{1/B}, \quad a_n = \frac{A}{B}(\ln n)^{\frac{1-B}{B}}. \tag{2.9}$$

2.2.2 The Peaks-Over-Threshold Approach and the GP Distribution

While the Block Maxima approach is easily implementable, it suffers from practical limitations: by extracting a single maximum per block, it discards a vast amount of potentially critical information (e.g., the second highest wind speed in a stormy year). To overcome these inefficiencies, the Peaks-Over-Threshold (POT) approach extracts all the independent observations that exceed a sufficiently high threshold u.

The BM and POT approaches are not separate statistical theories, but rather two different characterizations of the same extreme tail behaviour. If block maxima follow a GEV distribution, then the distribution of threshold excesses is correspondingly dictated by the Generalized Pareto (GP) distribution.

Theorem 2.2.3 (Pickands–Balkema–de Haan) *Let $X_1, X_2, \ldots, X_n$ be a sequence of i.i.d. random variables with distribution F. If there exist constants $a_n > 0$ and b_n such that $P(\frac{\max(X_1,\ldots,X_n)-b_n}{a_n} \leq z) \to G(z)$, where G is a GEV distribution with parameters (μ, σ, ξ), then, for large enough u, the distribution function of $(X - u)$ conditional on $X > u$ converges to the Generalized Pareto distribution:*

$$H(y) = 1 - \left(1 + \frac{\xi y}{\tilde{\sigma}}\right)_+^{-1/\xi}, \quad y > 0, \tag{2.10}$$

where the subscript + denotes the positive part, and the scale parameter $\tilde{\sigma}$ is functionally related to the GEV parameters by:

$$\tilde{\sigma} = \sigma + \xi(u - \mu). \tag{2.11}$$

A key implication of theorem 2.2.3 is that the shape parameter ξ of the GP distribution is identical to that of the corresponding GEV distribution. The value of this shape parameter strictly dictates the tail behaviour of the threshold excesses, reducing the generalized formulation to several well-known classical families:

- Exponential ($\xi \to 0$): eq. (2.10) reduces to the Exponential distribution, $H(y) = 1 - \exp(-y/\tilde{\sigma})$;
- Pareto ($\xi > 0$): the GP distribution is directly equivalent to the classic heavy-tailed Pareto distribution;
- $\xi < 0$: the distribution exhibits a short tail with an absolute upper bound at $y = -\tilde{\sigma}/\xi$.

While the POT method provides a more efficient use of data, its practical implementation introduces a significant analytical challenge: the choice of the threshold u. This selection entails a delicate bias-variance trade-off. Choosing a threshold that is too low violates the asymptotic basis of the extreme value model (the limit theorem no longer holds), introducing bias in the estimates. Conversely, setting a threshold that is too high retains too few exceedances, leading to high variance and uncertainty in the parameter estimation ([1]).

In practical engineering terms, maintaining consistency between the models allows for a robust formulation of the T-year return level z_T. In the POT framework, accounting for the unconditional probability of exceedance $\zeta_u = P(X > u)$ and the average number of observations per year n_y, the return level is derived by inverting the GP distribution:

$$z_T = \begin{cases} u + \frac{\tilde{\sigma}}{\xi}\left[(T n_y \zeta_u)^{\xi} - 1\right], & \text{for } \xi \neq 0 \\ u + \tilde{\sigma}\ln(T n_y \zeta_u), & \text{for } \xi = 0 \end{cases} \tag{2.12}$$

2.2.3 Rates of Convergence: Ultimate and Penultimate Approximations

The theoretical foundations of both the Block Maxima (BM) and Peaks-Over-Threshold (POT) approaches are intrinsically asymptotic. The GEV distribution arises as the limit law of suitably normalized maxima as $n \to \infty$, while the GP distribution emerges as the limiting distribution of threshold exceedances as the threshold u approaches the upper endpoint x_F. In practical engineering applications,

however, neither condition is ever fully attained: datasets are finite and thresholds cannot be chosen arbitrarily high without drastically reducing the available sample size.

Consequently, the finite-sample behaviour of extremes may differ substantially from the asymptotic model. The magnitude of this discrepancy depends on the rate at which the parent distribution converges to its limiting extreme-value law. Whereas the Central Limit Theorem is accompanied by Berry–Esseen bounds establishing a convergence rate of order $O(n^{-1/2})$ under suitable conditions, no universal convergence rate exists in Extreme Value Analysis. Instead, the rate of convergence depends strongly on the tail properties of the parent distribution and may be extremely slow.

When the limiting GEV or GP model is adopted directly, the resulting approximation is referred to as the *ultimate approximation*. To improve finite-sample performance, Fisher and Tippett (1928) introduced the concept of a *penultimate approximation*, in which the asymptotic shape parameter is replaced by a sample-size- or threshold-dependent quantity that captures the local curvature of the tail.

Penultimate Approximation for the GEV Distribution

Consider a parent distribution F belonging to the Gumbel domain of attraction. Define

$$u(x) = -\log[-\log F(x)],$$

and let

$$v(t) = u^{\leftarrow}(t)$$

denote its generalized inverse. The local deviation from exact Gumbel behaviour is quantified by the curvature function

$$\eta(t) = \frac{v''(t)}{v'(t)}.$$

Under suitable von Mises regularity conditions, the convergence of normalized maxima can be improved by replacing the asymptotic Gumbel limit G_0 with a GEV distribution whose shape parameter varies with sample size [4].

Theorem 2.2.4 (Penultimate Approximation for the Gumbel Domain) *Suppose that the parent distribution satisfies the following asymptotic limits:*

1. $\lim_{x\to\infty} \eta(x) = 0$ *(von Mises' first-order condition),*
2. $\lim_{x\to\infty} \frac{v'''(x)}{v''(x)} = 0$ *(von Mises' second-order condition),*
3. $\lim_{x\to\infty} \frac{\eta''(x)}{\eta'(x)} = 0$ *(von Mises type penultimate condition).*

Then, there exists a sequence $\{\gamma_n\}_{n\geq 1}$ *such that*

$$F^n(a_n x + b_n) - G_{\gamma_n}(x) = o(F^n(a_n x + b_n) - G_0(x)), \tag{2.13}$$

uniformly over $x \in \mathbb{R}$*. More specifically, with the exact location and scale norming constants given by* $b_n := v(\ln n)$ *and* $a_n := v'(\ln n)$*, and defining the penultimate shape parameter as* $\gamma_n := \eta(\ln n)$*, it follows that:*

$$\lim_{n\to\infty} \frac{F^n(a_n x + b_n) - G_{\gamma_n}(x)}{\eta'(\ln n)} = \frac{x^3 G_0'(x)}{6}, \tag{2.14}$$

uniformly for all $x \in \mathbb{R}$*.*

The theorem shows that the penultimate GEV distribution provides a more accurate approximation than the ultimate Gumbel model.

For a Weibull parent distribution

$$F(x) = 1 - \exp\left[-\left(\frac{x}{A}\right)^B\right],$$

the transformed quantile function is

$$v(t) = A\, t^{1/B}.$$

Differentiating gives

$$v'(t) = \frac{A}{B} t^{1/B-1},$$

and

$$v''(t) = \frac{A}{B}\left(\frac{1}{B} - 1\right) t^{1/B-2}.$$

Therefore,

$$\eta(t) = \frac{v''(t)}{v'(t)} = \frac{1-B}{B}\frac{1}{t}.$$

Evaluating at $t = \ln n$ yields

$$\gamma_n = \frac{1-B}{B \ln n}. \tag{2.15}$$

Penultimate Approximation for the GP Distribution

An analogous framework exists for the Peaks-Over-Threshold approach. For distributions in the Gumbel domain, the ultimate GP approximation corresponds to an exponential tail ($\xi = 0$). However, as shown by Worms [5], the convergence rate can be improved by introducing a threshold-dependent penultimate shape parameter $\Lambda(u)$.

The local scale parameter is naturally defined as

$$\sigma(u) = \frac{1 - F(u)}{f(u)} = \frac{1}{r(u)}. \tag{2.16}$$

For a Weibull parent distribution,

$$\sigma(u) = \frac{A^B}{B} u^{1-B}. \tag{2.17}$$

Worms showed that the penultimate GP shape parameter satisfies

$$\Lambda(u) = \sigma'(u). \tag{2.18}$$

Differentiating the local scale function gives

$$\Lambda(u) = \frac{1 - B}{B} \left(\frac{u}{A}\right)^{-B}. \tag{2.19}$$

If $B > 1$, as expected in wind applications, this result explains why analyses of finite meteorological records frequently produce negative estimates of the GEV or GP shape parameter even though the underlying Weibull parent distribution belongs to the Gumbel domain of attraction. Such negative estimates should not necessarily be interpreted as evidence of a bounded tail; rather, they are often manifestations of the slow pre-asymptotic convergence inherent to Weibull-type wind climates.

2.3 Non-Stationary Extreme Value Analysis

As highlighted in the previous sections, classical EVA relies on the stationarity assumption. However, under a changing climate and evolving aerodynamic roughness due to urbanization, this hypothesis is no longer valid. To address this, the Extreme Value Analysis must be extended into a Non-Stationary Extreme Value Analysis (NEVA) framework.

In a non-stationary context, the parameters governing the extreme value distributions—namely the location, scale, and shape parameters vector $\boldsymbol{\theta} = (\mu, \sigma, \xi)$ for the GEV, or $\boldsymbol{\theta} = (\tilde{\sigma}, \xi)$ for the GP distribution—are no longer treated as strictly constant scalars. Instead, they are modelled as functions of time or other relevant physical covariates. For instance, the location parameter can be expressed as a time-dependent linear function $\mu(t) = \mu_0 + \mu_1 t$, or linked to a large-scale climatic index. Generally, in non-stationary modelling, the shape parameter ξ is kept constant to ensure numerical stability and avoid unrealistic tail behaviour, while variations are absorbed by the location and scale parameters.

2.3.1 Maximum Likelihood Estimation (MLE)

The traditional approach to fit both stationary and non-stationary EVA models is the Maximum Likelihood Estimation (MLE) method. In a general statistical framework, let $\mathbf{y} = (y_1, y_2, \ldots, y_n)$ be a vector of independent extreme observations drawn from a probability density function $f(y|\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is the parameter vector to be estimated. The likelihood function $L(\boldsymbol{\theta}|\mathbf{y})$ represents the joint probability of observing the given sample:

$$L(\boldsymbol{\theta}|\mathbf{y}) = \prod_{t=1}^{n} f(y_t|\boldsymbol{\theta}).$$

In practice, to simplify calculations and prevent numerical underflow, it is customary to maximize the log-likelihood function $l(\boldsymbol{\theta}|\mathbf{y})$:

$$l(\boldsymbol{\theta}|\mathbf{y}) = \sum_{t=1}^{n} \log f(y_t|\boldsymbol{\theta}).$$

In a stationary model, $\boldsymbol{\theta}$ is constant for all observations. However, in a non-stationary NEVA framework, the parameters are modelled as functions of time or other physical covariates $\mathbf{x}_t$. Consequently, the log-likelihood function is modified such that the parameter vector becomes time-varying, $\boldsymbol{\theta}_t = \boldsymbol{\theta}(\mathbf{x}_t)$:

$$l(\boldsymbol{\theta}|\mathbf{y}) = \sum_{t=1}^{n} \log f(y_t|\boldsymbol{\theta}(\mathbf{x}_t)).$$

The MLE method seeks to find the parameter estimates $\hat{\boldsymbol{\theta}}$ that maximize this non-stationary log-likelihood.

While Maximum Likelihood Estimation (MLE) is computationally efficient and has been the most widely adopted procedure in early non-stationary studies [6], it presents significant mathematical and practical limitations. These shortcomings have been extensively detailed in a recent comprehensive systematic review by Radfar et al. [7], highlighting several critical problematic aspects of the MLE approach. Specifically, the review points out that MLE does not always provide a unique mathematical solution and frequently fails or yields highly unreliable estimates when applied to small sample sizes (such as short historical meteorological records).

2.3.2 Bayesian Inference

To overcome the limitations of the MLE approach, a shift towards the Bayesian inference framework is required. Unlike frequentist methods that provide point estimates of the parameters, the Bayesian approach treats the parameters $\boldsymbol{\theta}$ as random variables. According to Bayes' theorem, the posterior probability distribution

of the parameters $p(\boldsymbol{\theta}|\mathbf{y})$, given the observed extreme data $\mathbf{y}$, is proportional to the product of the likelihood function $L(\mathbf{y}|\boldsymbol{\theta})$ and a prior distribution $\pi(\boldsymbol{\theta})$:

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto L(\mathbf{y}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}). \tag{2.20}$$

The prior distribution allows researchers to incorporate physical constraints or expert knowledge before analysing the data.

The closed form of the posterior distribution is known only for a few distributions, and, in general, it requires the resolution of complex, multidimensional integrals. For this reason, Bayesian inference in EVA relies on advanced computational sampling techniques, primarily Markov Chain Monte Carlo (MCMC) algorithms.

The fundamental idea behind MCMC methods is to numerically study the characteristics of a complex probability distribution by drawing samples from it. Specifically, the goal is to construct a Markov chain whose limit distribution exactly matches the target posterior distribution of interest. Starting from an arbitrary initialization of the parameters, the algorithm sequentially generates new samples through a transition probability. Regardless of the starting point, the Markov chain is mathematically designed to eventually converge to the desired posterior distribution.

In practical implementations, the raw output of an MCMC algorithm requires specific diagnostics to ensure validity. Because the initial samples are heavily influenced by the arbitrary starting values, it is standard practice to discard a certain number of early iterations—a phase known as "burn-in"—to ensure that the retained samples strictly belong to the converged distribution. Furthermore, since consecutive samples in a Markov chain are inherently autocorrelated, a "thinning" procedure is frequently applied. By retaining only one sample every k iterations, the serial dependence is reduced, yielding a set of independent and identically distributed samples.

Depending on the specific mathematical structure of the problem, different MCMC algorithms can be employed. When the full conditional distribution of a parameter (i.e., the distribution of one variable given all the others) is known in a closed form, the Gibbs Sampler is typically used to sample one variable at a time. Conversely, in highly complex models where the full conditional is not analytically tractable, the Metropolis-Hastings algorithm is applied. This approach utilizes a proposal distribution to generate a new state, which is then evaluated through an accept-reject mechanism.

The application of MCMC algorithms has led to the development of software tools for non-stationary hazard assessment. Cheng et al. [8] pioneered one of these foundational frameworks by developing the NEVA package. To estimate the non-stationary GEV parameters, NEVA utilizes a Differential Evolution Markov Chain (DE-MC) approach. However, a major practical limitation of this software is its restricted functional flexibility: under non-stationary conditions, the framework

only allows modelling the parameters as linear functions of time, and no other covariates are included in the analysis.

Recognizing that time alone is often insufficient for complex climatic shifts, and that simple linear trends may oversimplify natural variability, Ragno et al. [9] expanded this concept by introducing ProNEVA (Process-informed Nonstationary Extreme Value Analysis). This generalized framework overcomes the limitations of its predecessor in two significant ways. First, it integrates physical drivers—such as CO_2 concentrations, or urbanization indices—directly as covariates into the extreme value distributions. Second, it allows the distribution parameters to be modelled not only as linear functions, but also as exponential or quadratic functions of the user-defined covariates. To estimate the parameters and quantify the associated uncertainty, ProNEVA is powered by a highly efficient hybrid-evolution MCMC algorithm. More mathematical details on this algorithm will be extensively discussed in Chapter 3.

Finally, it is important to note that while MCMC algorithms are exceptionally powerful for temporal non-stationarity, they are not the only available tools for Bayesian inference. In fact, MCMC approaches can become computationally prohibitive when high-dimensional spatial dependencies are introduced. Alternative Bayesian methodologies have been proposed. For example Opitz et al. [10] proposed a flexible Bayesian generalized additive modelling framework using Integrated Nested Laplace Approximations (INLA). By decomposing the estimation of marginal extremes into distinct stages and assuming that non-stationary patterns in the bulk of the distribution are similar to those in the tail, they successfully estimated high spatio-temporal quantiles for precipitation across unobserved locations, demonstrating the vast potential of advanced Bayesian techniques.

2.4 NEVA Application to Climate Variables

Although the systematic application of Non-Stationary Extreme Value Analysis (NEVA) to structural wind engineering is still at its infancy, the methodology has reached a higher level of maturity in other environmental and engineering disciplines. Some application fields are reviewed in the following.

2.4.1 Non stationary Wind Hazard

Turning to the wind literature, much of the current research is driven by the renewable energy sector, which inherently focuses on mean wind speeds rather than extreme, hazardous gusts. For example, Bonanno et al. [11] utilized an ensemble of bias-corrected Euro-CORDEX regional climate models to assess how climate change will affect wind power generation across the Italian peninsula under RCP 4.5 and 8.5 scenarios. While their study provides crucial insights into future wind

producibility, it exclusively analyses monthly and annual mean wind speeds. The authors explicitly acknowledge that potential climate-related changes in intra-daily variability and extreme wind events require further, dedicated investigation.

To accurately model wind fields over complex topography, researchers have increasingly relied on mesoscale reanalysis data and high-resolution projections. Raffa et al. [12] recently developed VHR-PRO_IT, an hourly climate projection dataset over Italy at a remarkably detailed spatial resolution of 2.2 km. Similarly, Raffaele et al. [13] leveraged high-resolution reanalysis models to propose the first meso-γ scale Italian wind map of extreme wind speeds (the 50-year return level, V_{50}). Furthermore, Mo et al. [14] utilized surface observations and reanalysis data to adjust extreme wind speeds in China for time-varying exposure categories due to urbanization.

However, a major limitation persists in these structural hazard assessments: they generally continue to rely on the stationarity assumption. For instance, the estimation of V_{50} in [13] relies on historical reanalysis without projecting future non-stationary climate shifts, potentially underestimating the future aerodynamic loads on structures designed today.

Recognizing the impact of environmental changes on local wind fields, a few recent studies have begun to apply non-stationary EVA to wind records. A notable example is the work by Yang et al. [15], who investigated the impact of rapid urbanization on near-surface wind speeds using 30 years of data from the Beijing meteorological tower. They successfully fitted a non-stationary GEV model where the location and scale parameters were treated as time-varying functions to capture temporal and seasonal variations.

However, studies like this often rely on simpler parameter estimation methods (e.g., standard MLE or least squares), which struggle to robustly quantify the deep epistemic uncertainties of a changing climate. While advanced high-resolution climate projections (RCMs) are now available, and robust Bayesian non-stationary statistical frameworks have been validated in other fields, the structural engineering community still lacks a comprehensive procedure to bridge these domains for wind hazard assessment.

2.4.2 Other non stationary climate hazards

The *coastal environment* subjected to the *rise of water level and wave height* is arguably the sector where non-stationary modelling has seen the most widespread application. Due to global warming, sea-level rise, and anthropogenic interventions, marine signals exhibit strong non-stationarity, directly threatening coastal communities and infrastructures. In their systematic literature review, Radfar et al. [7] analysed decades of NEVA applications in coastal studies. They found that the majority of current research focuses on the non-stationary modelling of extreme

water levels and wave heights. To capture the evolving nature of these hazards, researchers have moved beyond simple linear time trends. Instead, the parameters of the extreme value distributions (predominantly the GEV) are frequently modelled using large-scale macro-climatic indices—such as the North Atlantic Oscillation (NAO) or the Southern Oscillation Index (SOI)—as physical covariates.

Similarly, the fields of *hydrology* have heavily adopted the NEVA framework to manage the increasing frequency and intensity of *severe rainfall events*. As previously discussed, early studies such as Beguería et al. [6] successfully applied non-stationary POT approaches to daily precipitation series, identifying significant temporal trends in the magnitude and intensity of rainfall events. More recently, the focus has shifted towards highly complex spatio-temporal predictions. The work of Opitz et al. [10] on Dutch precipitation extremes exemplifies this trend, utilizing advanced Bayesian models to map high-quantile rainfall across wide geographical areas.

Furthermore, *temperature extremes*, particularly heatwaves, have been extensively analysed under non-stationary assumptions. The original NEVA framework developed by Cheng et al. [8] was validated specifically on annual temperature maxima, proving that stationary return periods heavily underestimate the risk of extreme heat in future scenarios. The generalized ProNEVA framework [9] demonstrated the profound versatility of Bayesian non-stationary modelling across a diverse spectrum of environmental hazards. Specifically, the authors successfully validated their methodology by integrating tailored physical drivers—such as urbanization indices for river discharge, atmospheric CO_2 concentrations for temperature extremes, and temporal covariates for sea-level rise—proving that customized, physics-informed covariates consistently outperform strictly stationary models.

In the light of the above, the primary gap this MSc Thesis seeks to fill is the application of a fully Bayesian, process-informed Non-Stationary Extreme Value Analysis to highly localized wind speed data. By doing so, this work represents a step forward in developing a novel methodology for evaluating wind hazard, explicitly accounting for the deep uncertainties driven by climate change.

Chapter 3

Statistical model

As highlighted in the literature review, assessing structural wind hazards under a changing climate requires abandoning the traditional stationarity assumption. This chapter formalizes a fully Bayesian, non-stationary extreme value framework.

To systematically validate the computational algorithms (which will be applied to synthetic data in chapter 4) and subsequently analyse real-world time series (chapter 5), the statistical model is developed progressively. The formulation starts from a stationary baseline model and advances to a non-stationary generalization where time (t) is adopted as the sole physical covariate to capture the evolving climatic trends. Finally, the chapter details the Hybrid Markov Chain Monte Carlo (MCMC) algorithm utilized to solve the Bayesian inference problem.

While the mathematical formulations and the MCMC sampling described herein have been computationally executed utilizing the ProNEVA Matlab package [9] (fig. 3.1 reports the architecture), the subsequent diagnostic evaluations, although taking inspiration from the same framework, were independently implemented by the author. This external implementation was specifically designed to extract and exploit a more comprehensive set of analytical information, ensuring a rigorous validation of the adopted models.

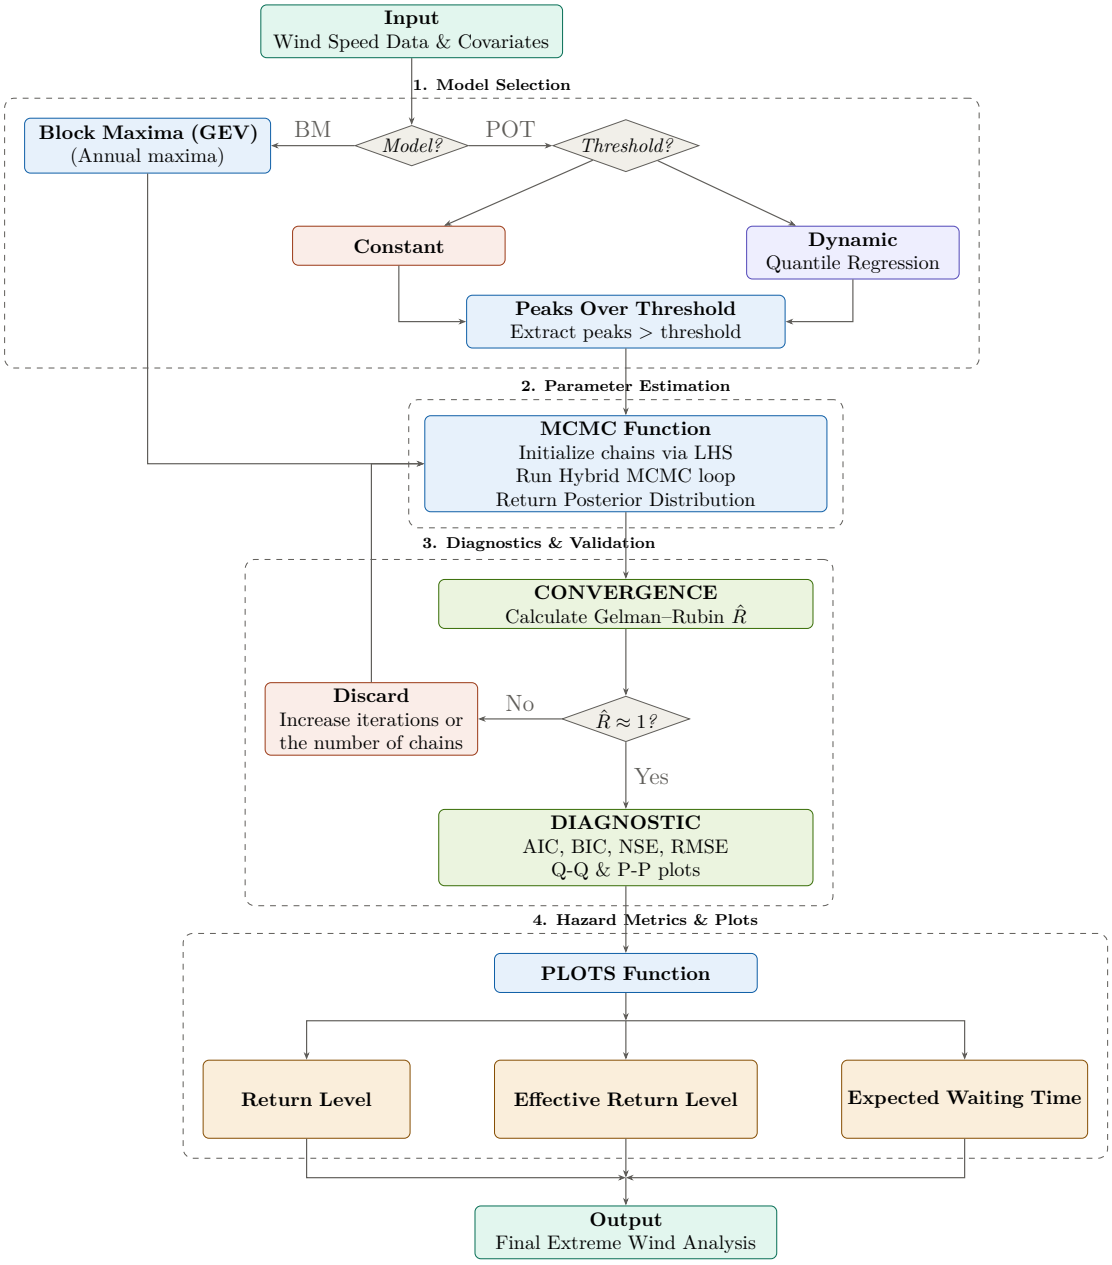

Figure 3.1: Overall architecture and workflow of the ProNEVA software package.

3.1 Formulation of the Extreme Value Models

The framework integrates the two major paradigms of extreme value modelling: the Block Maxima (BM) approach with Generalized Extreme Value (GEV) fitting, and the Peaks-Over-Threshold (POT) approach with Generalized Pareto (GP) fitting.

3.1.1 The Stationary Baseline

In the stationary baseline model, the underlying physical processes generating the wind extremes are assumed to be constant over time. Consequently, the parameter vectors governing the extreme value distributions are treated as time-invariant scalars.

For the BM approach, the distribution of the annual maxima is modelled via the GEV distribution, governed by the parameter vector $\boldsymbol{\theta} = (\mu, \sigma, \xi)$, representing the constant location, scale, and shape parameters, respectively. Similarly, for the POT approach, the threshold exceedances are modelled via the GP distribution governed by $\boldsymbol{\theta} = (\tilde{\sigma}, \xi)$, representing the constant scale and shape parameters. In this way, the number of parameters that needs to be estimated is respectively three and two, for the GEV and GP distributions.

To rigorously validate the computational framework, the parameters of the stationary models are estimated utilizing both the traditional Maximum Likelihood Estimation (MLE) and Bayesian inference via Markov Chain Monte Carlo (MCMC) sampling. This dual approach allows for a comprehensive methodological comparison, exploiting the robust uncertainty quantification provided by the Bayesian posterior distributions against the standard frequentist point estimates.

3.1.2 The Non-Stationary Extension

The necessity of a non-stationary framework is preliminarily assessed through formal statistical hypothesis tests on the underlying wind time series. The following description follows the methodology detailed in the supplementary material of the ProNEVA framework [9].

First, to identify monotonic temporal changes the Mann-Kendall (MK) trend test is typically applied. The MK test evaluates the monotonic trend of a time series $X = \{x_1, \ldots, x_n\}$ based on Kendall's S-statistic, defined as the difference between the numbers of concordant and discordant pairs:

$$S = \sum_{i<j} \operatorname{sign}(x_i - x_j). \tag{3.1}$$

When the variable X is independently and randomly distributed, the expected value is $S = 0$. For large samples, S tends to normality, allowing the null hypothesis of

no monotonic trend to be tested. However, within the specific adopted framework, when assessing the dependence on a physical covariate X_c (such as time t), the traditional MK test is replaced with a zero-slope test for a linear regression model between the observed data $\tilde{Y}$ and the covariate. Given the true slope β_1 and its estimate $\hat{\beta}_1$, the test statistic is formulated as:

$$t_{n-2} = \frac{\hat{\beta}_1 - \beta_1}{s_{\beta_1}} = \frac{\hat{\beta}_1}{s_{\beta_1}}, \tag{3.2}$$

where s_{β_1} is the standard error of the estimate, and n is the number of observations. The null hypothesis $H_0 : \beta_1 = 0$ is evaluated against a Student's t-distribution with $n - 2$ degrees of freedom at a 5% level of significance.

Second, while detecting a trend in the mean is crucial, it is equally important to investigate shifts in the variance. To test the homoscedasticity assumption, the White test is utilized. Assuming a preliminary linear regression model $\tilde{Y} = \beta_0 + \beta_1 X_c + \epsilon$, with fitted values $\hat{Y} = \beta_0 + \beta_1 X_c$, the test focuses on the squared residuals ϵ^2. The heteroscedasticity is evaluated by estimating an auxiliary regression:

$$\epsilon^2 = \delta_0 + \delta_1 \hat{Y} + \delta_2 \hat{Y}^2 + \text{error}. \tag{3.3}$$

The null hypothesis of homoscedasticity posits that the variance does not depend on the covariate, implying $\delta_1 = 0$ and $\delta_2 = 0$. Using the $R^2_{\epsilon^2}$ from this auxiliary regression, the Lagrange Multiplier (LM) statistic is calculated as:

$$LM = n \cdot R^2_{\epsilon^2}, \tag{3.4}$$

which follows a χ^2 distribution with 2 degrees of freedom (χ^2_2). The White test is performed at a 5% significance level.

Rejecting the null hypotheses of either the zero-slope test or the White test provides a robust, data-driven foundation for adopting a non-stationary approach, where both the location and scale parameters are allowed to evolve.

To capture the effect of climate change, the framework is extended into a non-stationary domain. In this setting, the parameters of the extreme value distributions are no longer constant scalars but are modelled as functions of a physical covariate. In the context of this thesis, the covariate adopted is time (t).

To ensure numerical stability and avoid unrealistic tail behaviour—especially when dealing with relatively short meteorological records—the shape parameter ξ is kept constant. The time-dependency is formulated through linear trends:

$$\mu(t) = b_\mu + a_\mu t, \tag{3.5}$$

$$\sigma(t) = \exp(b_\sigma + a_\sigma t). \tag{3.6}$$

The exponential function applied to the scale parameter ensures that $\sigma(t)$ remains strictly positive across the entire temporal domain, satisfying the fundamental mathematical constraints of the GEV and GP distributions. Consequently, for a non-stationary BM analysis, the vector of unknown parameters to be estimated expands from three scalars to five: $\boldsymbol{\theta}_t = (a_\mu, b_\mu, a_\sigma, b_\sigma, \xi)$, while for the POT analysis from two to three.

In the non-stationary POT framework, the threshold selection requires special attention. A fixed threshold might lead to the over- or under-selection of extreme events if the underlying statistical properties of the wind series shift over time. To address this, the model implements a dynamic threshold $u(t)$ estimated via Quantile Regression (QR), ensuring that a consistent proportion of the tail is sampled throughout the historical record.

Unlike standard Ordinary Least Squares (OLS) regression, which estimates the conditional mean of a response variable, Quantile Regression evaluates the conditional quantiles. Let y_i be the wind speed observations recorded at time t_i. The objective of QR is to estimate the τ-th conditional quantile of the wind speed, denoted as $Q_y(\tau|t)$, where $\tau \in (0, 1)$ represents a sufficiently high non-exceedance probability (e.g., $\tau = 0.95$ for the 95th percentile threshold).

Assuming a linear temporal dependence for the dynamic threshold, the function can be expressed as:

$$u(t) = Q_y(\tau|t) = b_u + a_u t. \tag{3.7}$$

The regression coefficients b_u and a_u are estimated not by minimizing the sum of squared residuals, but rather by minimizing an asymmetric absolute loss function (often referred to as the *pinball loss*):

$$\min_{b_u, a_u} \sum_{i=1}^{n} \rho_\tau\Big(y_i - (b_u + a_u t_i)\Big), \tag{3.8}$$

where the loss function $\rho_\tau(e)$ for a residual e is defined as:

$$\rho_\tau(e) = \begin{cases} \tau e & \text{if } e \geq 0 \\ (\tau - 1)e & \text{if } e < 0 \end{cases} \tag{3.9}$$

From a computational perspective, solving this minimization problem involves iterative optimization algorithms. To ensure rapid convergence and numerical stability, the solver requires an appropriate initial guess for the parameters. The numerical search is initialized using the parameter estimates derived from the standard OLS regression, facilitating the solver's efficiency in navigating towards the optimal quantile coefficients.

By minimizing this asymmetric function, the algorithm effectively fits a line $u(t)$ that keeps exactly $(1 - \tau)$ of the observations above it, regardless of the

underlying data trend. Dynamically adjusting the threshold along the time series guarantees that the exceedance rate remains approximately homogeneous over the entire observational period.

3.2 Bayesian Inference and Hybrid MCMC Algorithm

Once the extreme value models are defined, the next critical step is estimating the parameter vector $\boldsymbol{\theta}_t$. This estimation is performed using a fully Bayesian framework.

According to Bayes' theorem, the posterior probability distribution $p(\boldsymbol{\theta}|\mathbf{y})$, given the observed wind extremes $\mathbf{y} = (y_1, \ldots, y_n)$, is proportional to the product of the likelihood function and the prior distribution:

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto L(\mathbf{y}|\boldsymbol{\theta})p(\boldsymbol{\theta}) \tag{3.10}$$

To compute this posterior, it is first necessary to define the prior distributions $p(\boldsymbol{\theta})$, which encode the initial assumptions about the parameters before observing the data. To prevent subjective bias and allow the historical wind records to strictly drive the estimation, "uninformative" or weakly informative priors are deliberately assigned. Specifically, the location parameter is assigned a Normal distribution with a mean and standard deviation matching those of the observed historical series. Conversely, the scale and shape parameters are assigned zero-mean Normal distributions with exceptionally high variances. This configuration, adopted for both the GEV and GP distributions, ensures that the priors do not artificially constrain the parameter space.

Because the analytical integration of the posterior distribution is mathematically intractable for complex non-stationary models, numerical sampling becomes essential. To achieve this efficiently, the adopted framework utilizes a Hybrid Markov Chain Monte Carlo (MCMC) algorithm (as depicted in fig. 3.2), drawing inspiration from the approach proposed by Sadegh et al. [16].

To guarantee a uniform exploration of the parameter space and avoid poor initializations, the algorithm draws an initial set of samples using Latin Hypercube Sampling (LHS). These samples are assigned to multiple parallel chains, selecting the points that maximize the log-likelihood as the starting states.

Once the initialization phase is completed, the algorithm enters the main MCMC loop. The core innovation of this hybrid approach lies in its dynamic proposal generation step, which balances local refinement and global exploration by probabilistically alternating between distinct updating strategies. At each iteration, a random uniform variable $u \sim \mathcal{U}(0,1)$ is drawn to determine the specific strategy.

With a 10% probability ($u < 0.1$), the algorithm performs a Snooker update. This strategy is specifically designed to prevent the chains from getting trapped in local optima, promoting a robust exploration of the entire parameter space. It generates a proposal candidate by projecting two randomly selected chains onto a direction vector defined by the difference between the current chain and a third randomly selected chain. By leveraging the spatial distribution of multiple parallel chains, the Snooker update allows the algorithm to make significant, geometrically guided jumps across the multidimensional posterior landscape.

For the remaining 90% of the iterations ($u \geq 0.1$), the algorithm employs either an Adaptive Metropolis (AM) or a Differential Evolution (DE) update, typically operating within a block-updating (Gibbs) scheme on a subset of random dimension of the parameter dimensions. The choice between the AM and DE depends on the specific chain index and the stage of the simulation:

- The AM update utilizes the historical covariance matrix of the chains ($\hat{\Sigma}$) to adaptively tune a multivariate normal proposal distribution. This ensures a highly efficient exploration of the local posterior space by learning the scale and orientation of the target distribution as the simulation progresses.

- The DE update proposes a new candidate state by adding a scaled difference between two randomly chosen parallel chains to the current state. In this way, the chains can learn from one another, leading to significantly faster convergence.

Finally, once a candidate state $\mathbf{CH}^*$ is proposed by any of these mechanisms, it is subjected to the standard Metropolis-Hastings acceptance criterion. The Metropolis Ratio (MR) is computed, comparing the likelihood and prior probabilities of the candidate state against the current state (with specific mathematical adjustments for the Snooker projection). The candidate is accepted with probability $\max(MR, 1)$; otherwise, the chain is rejected and remains at its current state for that iteration. This robust hybrid mechanism ensures a faster, better-mixed, and more reliable convergence.

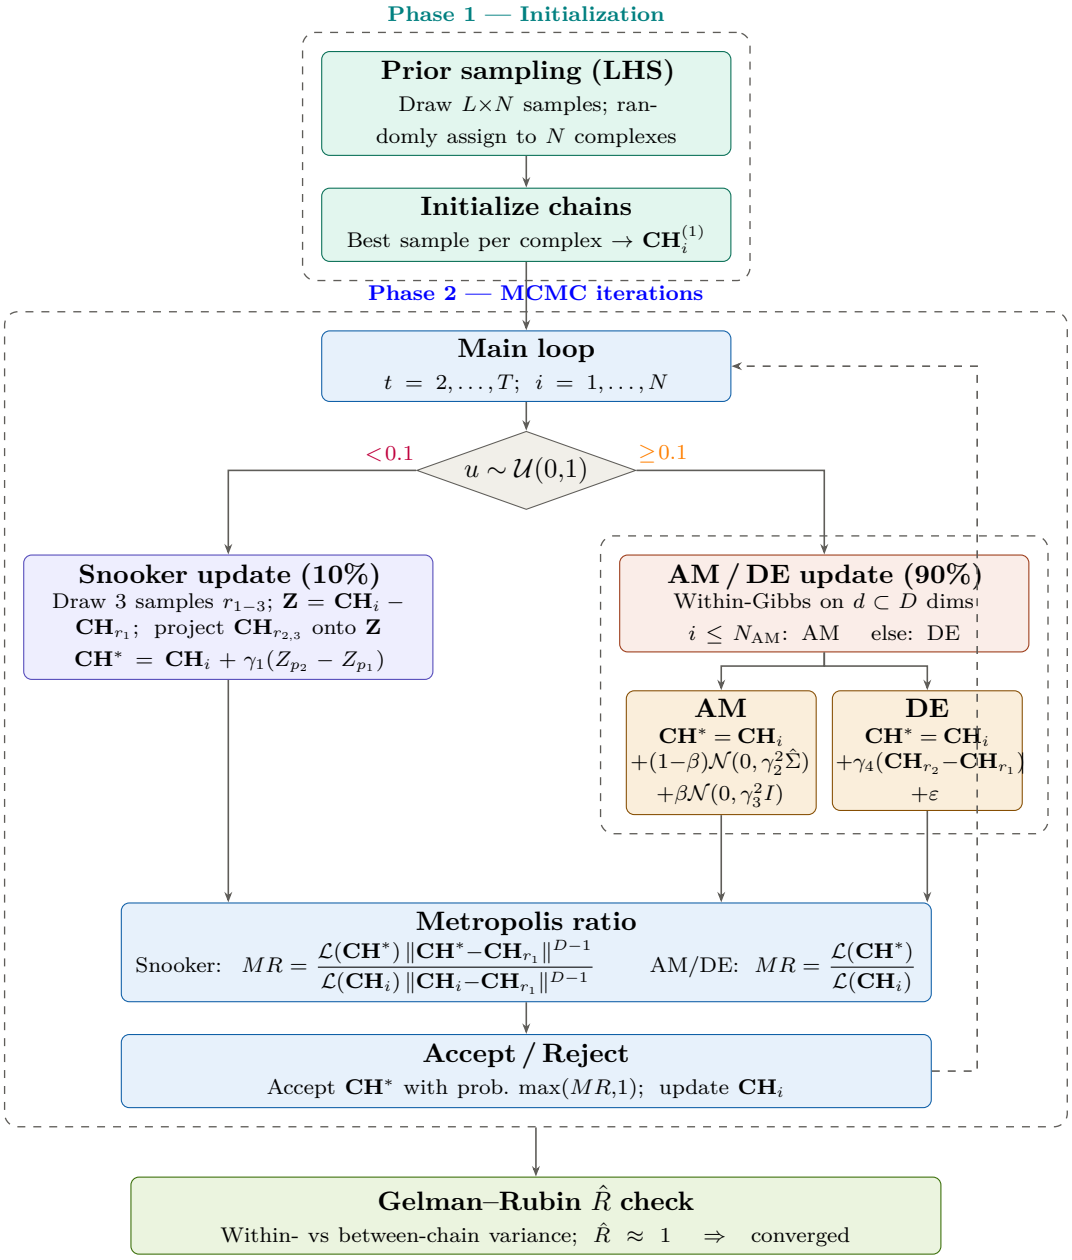

Figure 3.2: Flowchart of the Hybrid MCMC algorithm

3.2.1 Model Diagnostic and Convergence

Sampling from a posterior distribution requires rigorous checks to ensure the results are statistically valid. Initially, the algorithm discards the first portion of the generated samples, known as the *burn-in* period, during which the chains are still moving towards the high-probability regions of the posterior distribution.

To formally verify that the remaining samples have converged to the true stationary posterior, the Gelman-Rubin diagnostic ($\hat{R}$) is computed. This metric evaluates convergence by comparing the variance between multiple parallel chains with the variance within each individual chain.

Let n be the number of iterations per chain and m be the number of parallel chains. The within-chain variance W is defined as the mean of the variances of the individual chains, while the between-chain variance B is calculated from the variance of the chain means. From these components, the pooled posterior variance estimate $\hat{V}$ is computed as:

$$\hat{V} = \frac{n-1}{n}W + \frac{B}{n} \tag{3.11}$$

Finally, the factor $\hat{R}$ is calculated incorporating a degrees-of-freedom correction to account for the finite number of parallel chains:

$$\hat{R} = \sqrt{\frac{m+1}{m}\frac{\hat{V}}{W} - \frac{n-1}{m \cdot n}} \tag{3.12}$$

A value of $\hat{R}$ strictly below 1.2 indicates that the parallel chains are indistinguishable, implying that the algorithm has successfully converged and explored the target distribution thoroughly.

Once convergence is verified, the Goodness-of-Fit is assessed both visually and statistically. The Maximum A Posteriori (MAP) estimate is extracted to represent the best model parameters. Visual diagnostics include the Quantile Plot (Q-Q plot), which compares empirical quantiles against theoretical ones. In non-stationary scenarios, the time-varying observations are standardized to allow plotting on a single stationary reference scale [1].

For the Generalized Extreme Value (GEV) distribution, the non-stationary observations $X(t)$ (characterized by time-varying parameters $\mu(t)$, $\sigma(t)$, and fixed ξ) are transformed into a standard Gumbel distribution. This standardization is achieved through the following transformation:

$$Z(t) = \frac{1}{\xi}\ln\left(1 + \xi\frac{X(t) - \mu(t)}{\sigma(t)}\right) \tag{3.13}$$

The transformed data $Z(t)$ are then sorted in ascending order. To construct the Q-Q plot, these standardized empirical values are plotted against the theoretical

quantiles of the standard Gumbel distribution. These theoretical quantiles are calculated using the empirical non-exceedance probabilities $p_i = \frac{i}{N+1}$:

$$Z_{th} = -\ln(-\ln(p_i)) \tag{3.14}$$

For the Generalized Pareto Distribution (GPD), which is applied to excesses $Y(t)$ over a threshold, the non-stationary observations (characterized by time-varying scale $\sigma(t)$ and fixed shape ξ) are transformed into a standard Exponential distribution with a rate parameter equal to 1. The appropriate transformation is:

$$Z(t) = \frac{1}{\xi} \ln\left(1 + \xi \frac{Y(t)}{\sigma(t)}\right) \tag{3.15}$$

Once transformed and sorted, these values are compared against the theoretical quantiles of the standard Exponential distribution, derived as:

$$Z_{th} = -\ln(1 - p_i) \tag{3.16}$$

By applying these transformations, the originally non-stationary data are converted into independent and identically distributed (i.i.d.) samples. If the time-varying model accurately captures the underlying physical trends, the transformed points will align tightly along the 1:1 diagonal line on the respective Q-Q plots, validating the Goodness-of-Fit.

Additionally, the overall discrepancy between the empirical observations and the theoretical quantiles is quantified using the Root Mean Square Error (RMSE):

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}, \tag{3.17}$$

where y_i are the observed values and $\hat{y}_i$ are the model predictions.

To complement the absolute error metric provided by the RMSE, the relative predictive skill of the model is evaluated using the Nash-Sutcliffe Efficiency (NSE) coefficient:

$$NSE = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}, \tag{3.18}$$

where $\bar{y}$ represents the mean of the observed values. The NSE is a dimensionless metric ranging from $-\infty$ to 1. An NSE value of 1 corresponds to a perfect match between the model and the empirical data. A value of $NSE = 0$ implies that the model's predictions are only as accurate as the simple mean of the observed data, while values below zero indicate that the observed mean is a better predictor than the model itself. Consequently, values closer to 1 denote superior model reliability.

Finally, to address the bias-variance trade-off and perform rigorous model selection—particularly to justify whether the added complexity of a non-stationary

model is justified over the stationary baseline—the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) are evaluated:

$$AIC = 2k - 2\ln(\hat{L}), \tag{3.19}$$

$$BIC = k\ln(N) - 2\ln(\hat{L}), \tag{3.20}$$

where k represents the number of estimated parameters, N is the sample size, and $\hat{L}$ is the maximized value of the likelihood function. Models yielding the lowest AIC and BIC values are preferred, as they offer the optimal balance between goodness-of-fit and model parsimony, severely penalizing unnecessary parameters that could lead to overfitting.

3.2.2 Hazard prediction

The ultimate objective of this statistical framework is to translate complex probability distributions into tangible hazard metrics for structural design.

In a traditional stationary context, the wind hazard is communicated through the Return Level z_T, which corresponds to the wind speed quantile associated with a specific Return Period T (e.g., $T = 50$ years). However, under the non-stationary assumption, evaluating a single, constant return level is mathematically inconsistent, as the probability of exceeding any given threshold evolves continuously with time.

To visualize the continuous trajectory of the climate trend, the *Effective Return Level* is evaluated. This metric dynamically tracks the T-year quantile as a continuous function of time, producing a hazard curve $z_T(t)$ that evolves chronologically alongside the historical record and can be projected into future scenarios.

Considering eq. (2.5) and eq. (2.12), the Effective Return Level for a return period T is calculated by evaluating the tail quantile function conditioned on time, meaning for the GEV

$$z_T(t) = \begin{cases} \mu(t) - \frac{\sigma(t)}{\xi}\left[1 - \left\{-\log\left(1 - \frac{1}{T}\right)\right\}^{-\xi}\right], & \text{for } \xi \neq 0 \\ \mu(t) - \sigma(t)\log\left\{-\log\left(1 - \frac{1}{T}\right)\right\}, & \text{for } \xi = 0 \end{cases} \tag{3.21}$$

and for the GP with a dynamically evolving threshold $u(t)$

$$z_T(t) = \begin{cases} u(t) + \frac{\tilde{\sigma}(t)}{\xi}\left[(Tn_y\zeta_u)^{\xi} - 1\right], & \text{for } \xi \neq 0 \\ u(t) + \tilde{\sigma}(t)\log(Tn_y\zeta_u), & \text{for } \xi = 0 \end{cases} \tag{3.22}$$

Computing these Effective Return Levels across the temporal domain utilizing the MAP estimates yields a continuously updated wind speed profile, effectively converting complex non-stationary statistics into concrete criteria for structural evaluation.

Chapter 4

Application to synthetic wind speed time series

4.1 Generation of Synthetic Time Series

The primary objective of this chapter is to systematically evaluate the performance of the adopted Bayesian framework for non-stationary Extreme Value Analysis (NEVA). To achieve this, the model is tested against synthetic time series with fully controlled and known behaviours, as outlined in table 4.1. This section details the generation process of these synthetic baseline scenarios.

In wind engineering, the parent distribution of the mean wind speed is traditionally modelled using a Weibull distribution. Let $U_{10} \in \mathbb{R}^+$ denote the random variable representing the 10-minute mean wind speed measured at the standard reference height of 10 meters; this variable served as the basis for generating the synthetic time series. To ensure full reproducibility of the computational experiments, the generation process was initialized with a fixed random seed (42).

The synthetic parent distributions were simulated at an hourly resolution (providing 8760 discrete observations per year), with the total temporal horizon varying according to the specific case studies detailed in the subsequent subsections. For the stationary baseline, the time series were generated adopting a Weibull scale parameter $A = 2.38333$ m/s and a shape parameter $B = 1.60047$. Given these parameters, the expected value (mean) and the variance of the hourly wind speed are analytically defined as:

$$\mathbb{E}[U_{10}] = A\,\Gamma\left(1 + \frac{1}{B}\right), \tag{4.1}$$

$$\mathrm{var}(U_{10}) = A^2\left[\Gamma\left(1 + \frac{2}{B}\right) - \Gamma^2\left(1 + \frac{1}{B}\right)\right], \tag{4.2}$$

where $\Gamma(\cdot)$ denotes the Gamma function.

To simulate a non-stationary climate scenario characterized by a progressive intensification of the wind hazard, a linear temporal trend was superimposed on the expected value of the population. Specifically, the slope of the mean wind speed was fixed at $a_{\mathbb{E}} = 0.005$ m/s/yr. By assuming that the atmospheric regime dictates a constant shape parameter B, introducing a linear drift in $\mathbb{E}[U_{10}]$ inherently requires a continuous linear variation of the scale parameter $A(t)$. Observing eq. (4.2), since the variance scales quadratically with A, the time-dependency of the mean automatically induces a quadratic amplification of the variance over time. This mathematically constrained procedure completely defines the non-stationary parent distributions from which the synthetic extremes are subsequently extracted.

Finally, to ground the misspecification test in a physically plausible climate change scenario, a synthetic parent distribution characterized by a constant mean and a time-varying quadratic variance was generated. In the context of climate change, this configuration simulates a widely conjectured environmental condition: an atmospheric regime where the mean wind speed remains stationary, but there are more frequent and severe extreme events. Mathematically, this was modelled by defining the variance as $\mathrm{var}(t) = 5 \cdot 10^{-5} \cdot t^2 + 0.005 \cdot t + \mathrm{var}_0$, where t represents the year. The objective of this test was to assess the robustness of the Bayesian framework when operating in such a misspecified environment. As inferred from the preceding analytical formulas, enforcing a constant mean alongside a quadratic variance forces the parameters of the parent distribution—and, consequently, those of the resulting extreme value distribution—to exhibit strongly non-linear temporal behaviours. Evaluating the framework under these non-linear conditions allows for a rigorous assessment of its practical engineering performance, particularly in the estimation of effective return levels.

Table 4.1: Summary of the temporal evolution combinations for the mean and variance of the *parent* distribution. The cases analysed in this chapter are marked with an **X**

		Variance Evolution	
		Constant	**Quadratic varying**
Mean Evolution	**Constant**	**X**	**X**
	Linearly increasing	**-**	**X**

4.2 Demonstrative Application of EVA to a Stationary Random Variable

To evaluate the performances of the selected framework in a stationary scenario, an Extreme Value Analysis (EVA) was performed on the synthetic stationary time series of the random variable U_{10}. The sequence was generated over a simulated temporal projection of 2000 years, utilizing the Weibull parent distribution parameters defined previously. The statistical evaluation was conducted using both the Block Maxima (BM) and Peaks-Over-Threshold (POT) approaches. For the POT implementation, the threshold was strategically set at the 99.97th percentile of the parent distribution. Crucially, this specific value corresponds exactly to the minimum among all the annual maxima extracted for the BM approach, yielding a POT sample size approximately two to three times larger than that of the BM method.

In fig. 4.1 the generated synthetic time series is presented alongside the corresponding probability distributions. The extreme value distributions depicted in the figure are plotted utilizing the Maximum A Posteriori (MAP) parameter estimates derived from the MCMC algorithm.

A comprehensive comparison of the parameter estimates and Goodness-of-Fit (GoF) metrics is provided in table 4.2. Given the extensive length of the simulated time series (2000 years), the sample size of the extremes is exceptionally large. Consequently, as expected from asymptotic statistical theory, the frequentist Maximum Likelihood Estimation (MLE) and the Bayesian MCMC approaches yield remarkably similar parameter estimates. Furthermore, both numerical estimation methods align almost perfectly with the analytical parameters analytically derived from the Weibull parent distribution ((2.15), (2.19)).

The GoF metrics confirm the robustness of the statistical fits: the Nash-Sutcliffe Efficiency (NSE) exceeds 0.99 in all scenarios, indicating near-perfect predictive skill. Although both methods perform exceptionally well, the MCMC approach yields slightly lower AIC and BIC values, suggesting a marginally superior balance between model fit and parsimony.

Finally, fig. 4.2 displays the Quantile-Quantile (Q-Q) plots comparing the empirical data against the fitted models for both the GEV and GP distributions. The tight alignment of the markers along the reference diagonal and the almost perfect overlap of the data points confirm that both the MLE and MCMC frameworks successfully capture the tail behaviour of the underlying stationary process.

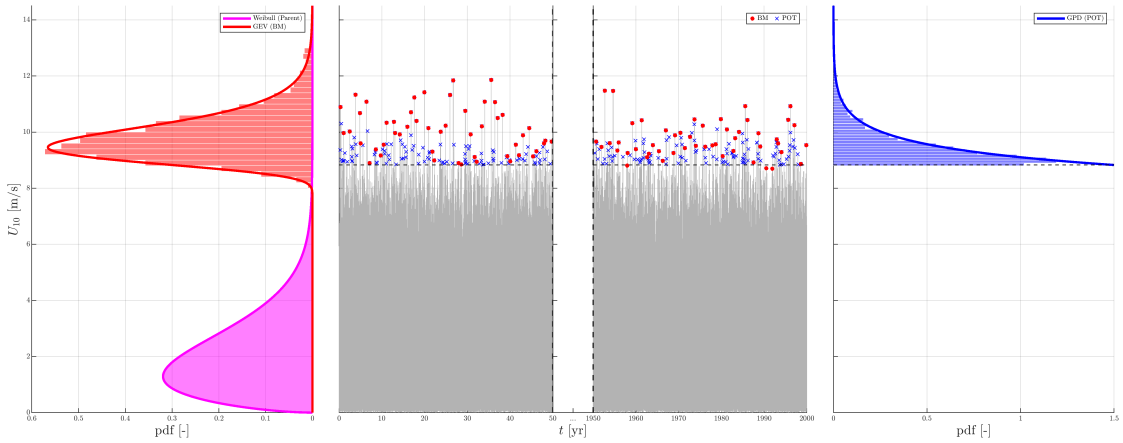

Figure 4.1: Synthetic time series generated from a stationary random variable on a 2000-year projection. Left: parent (Weibull) and GEV distributions. Middle: time series with annual maxima in red and peaks over threshold in blue. Right: GP distribution.

Table 4.2: Comparison of parameter estimates and Goodness-of-Fit (GoF) metrics for the 2000-year stationary synthetic time series. The parameters are estimated via MLE and MCMC, alongside the analytical formulas. For the GP distribution, the location parameter refers to the fixed threshold u.

	Block Maxima (GEV)			**Peaks Over Threshold (GP)**		
	MLE	**MCMC**	**Analytical**	**MLE**	**MCMC**	**Analytical**
Parameters						
Location (μ)	9.4470	9.4424	9.4569	8.8279	8.8279	8.8279
CI width	*0.0647*	*0.0668*	–	–	–	–
Scale (σ)	0.6593	0.6505	0.6509	0.6657	0.6670	0.6785
CI width	*0.0465*	*0.0448*	–	*0.0508*	*0.0525*	–
Shape (ξ)	−0.0263	−0.0232	−0.0413	−0.0233	−0.0283	−0.0461
CI width	*0.0617*	*0.0634*	–	*0.0538*	*0.0553*	–
GoF						
AIC	4592.15	4544.04	–	5994.68	5959.28	–
BIC	4608.95	4560.84	–	6007.81	5972.41	–
RMSE	2.8437	2.9294	–	3.0048	3.1159	–
NSE	0.9975	0.9974	–	0.9983	0.9982	–

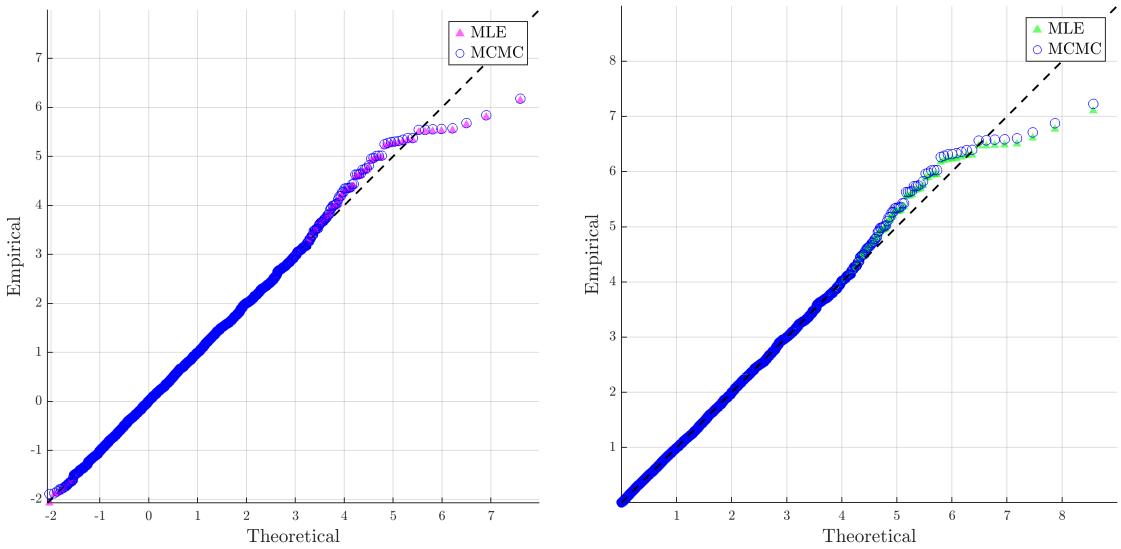

Figure 4.2: Quantile-Quantile (Q-Q) plots comparing the empirical quantiles against the theoretical estimates (MLE and MCMC). Left: GEV distribution. Right: GP distribution.

4.3 Application of NEVA to stationary and non-stationary Random Variables

4.3.1 Stationary time series

The objective of this section is to evaluate the robustness of the non-stationary MCMC framework against overfitting. Specifically, the goal is to quantify the error committed by the algorithm when a stationary random variable is provided, yet the Bayesian analysis is deliberately configured to search for a linear trend in the distribution parameters (location and scale for BM; scale for POT). A perfectly robust model should return slope parameters (a_μ, a_σ) equal to zero. The same stationary Weibull parent distribution defined in section 4.2 is adopted for this study.

For a 2000-year synthetic time series, both the BM and POT frameworks—the latter utilizing a selected threshold at the 99.97th percentile—perform exceptionally well. The results, presented in fig. 4.3 and fig. 4.4, are highly satisfactory, as the estimated slopes of the time-varying parameters are infinitesimally small (on the order of 10^{-5}). Consequently, the predictive probability density functions evaluated at $t = 0$ and $t = 2000$ years closely overlap, with only minor discrepancies observed in the GEV distribution due to the location parameter. The Effective Return Levels (ERL) present a marginally decreasing slope, which becomes more pronounced at higher return periods.

However, environmental time series available in practical engineering applications are significantly shorter, typically spanning around 60 years for simulated data and only 40 years for anemometric measurements. To systematically investigate the model's behaviour as a function of the available temporal horizon and threshold selection, a comprehensive convergence analysis is presented in fig. 4.5.

To generate this figure, the analysis was performed over varying periods—specifically 20, 40, 60, 80, 100, and 200 years. To filter out statistical noise and smooth the resulting curves, 50 random permutations were applied to the time series for each temporal period, before performing the Bayesian framework.

The top panel of fig. 4.5 demonstrates a clear log-log convergence: as the number of sampled years increases, the magnitude of the spurious slopes (for a_μ, a_σ, and the ERL_{50}) drops monotonically. Specifically, the typical time span of the available data is highlighted by a shaded gray rectangle. Within this window, it is evident that the optimal approach is to select a threshold between the 95th and 99th percentiles, as the error in the ERL_{50} slope drops below 3×10^{-3}. This error is almost an order of magnitude lower than the corresponding slopes generated by the GEV model or by the threshold corresponding to the minimum among the annual maxima. Moreover, the convergence rate of the ERL_{50} slope can be evaluated, resulting in a value of approximately -1.5.

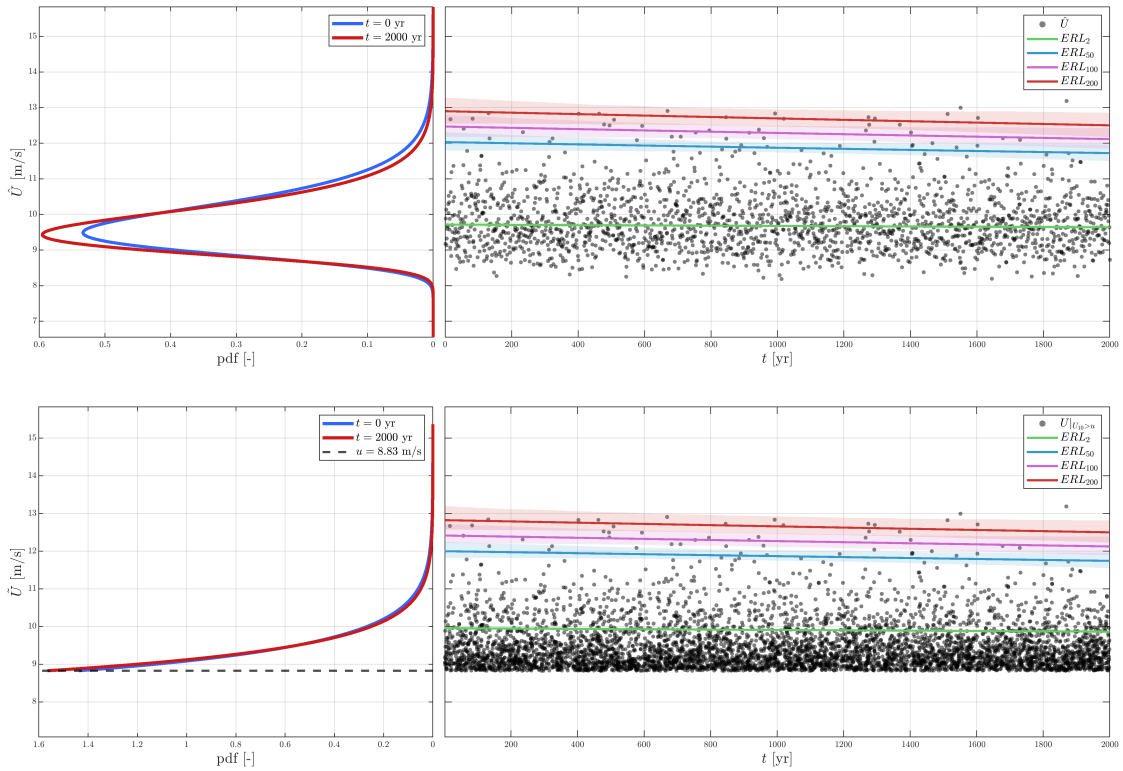

Figure 4.3: Non-stationary MCMC framework results for the GEV model (top row) and the GP model (bottom row). From left to right, the panels illustrate: the predictive probability density functions (pdf), the effective return levels (ERL_T) for $T = 2$, 50, 100, and 200 years.

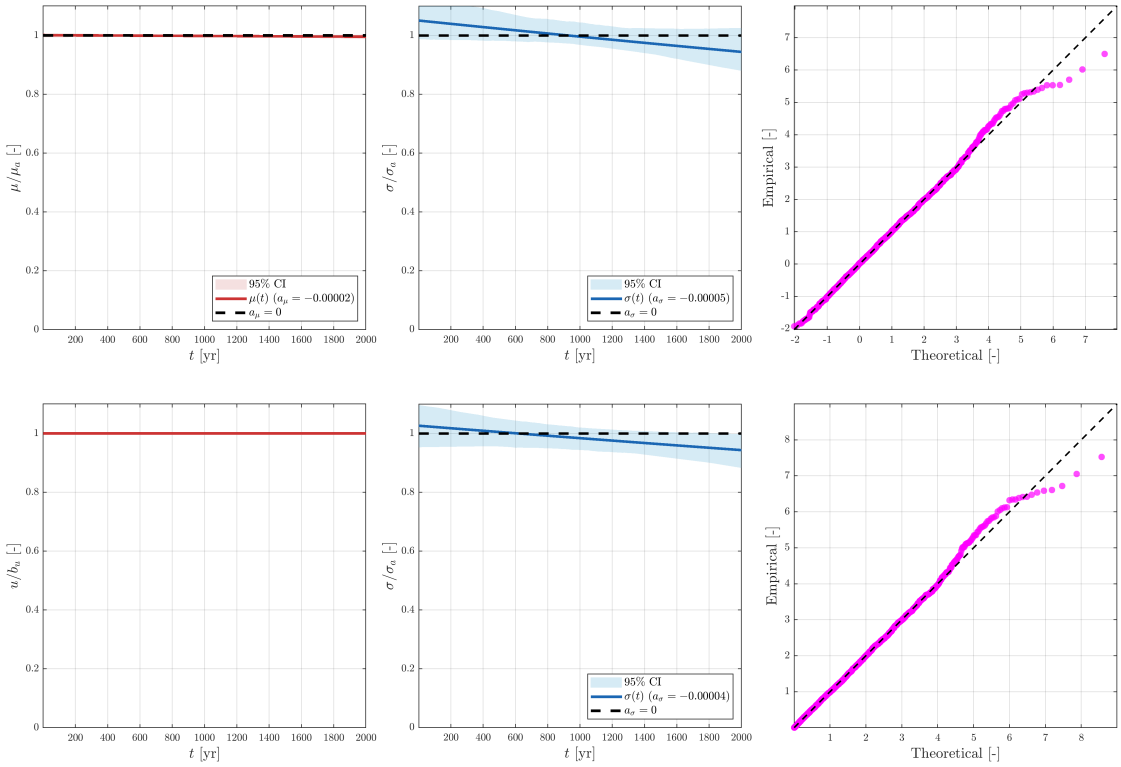

Figure 4.4: Non-stationary MCMC framework results for the GEV model (top row) and the GP model (bottom row). From left to right, the panels illustrate: the temporal evolution of the normalised distribution parameters (location and scale for the GEV; constant threshold and scale for the GPD), and the diagnostic Q-Q plots.

The bottom panel displays the absolute slope errors against the theoretical sample size (i.e., the expected number of extreme events fed to the algorithm, unifying both annual maxima and threshold exceedances). This unified scatter plot reveals a crucial insight: the accuracy of the non-stationary inference is largely independent of the specific extreme value paradigm chosen (BM vs. POT); rather, it is mainly a function of the absolute volume of data. To consistently suppress the spurious-trend error for the ERL_{50} slope, a massive sample size is theoretically required.

Because real-world time series of such magnitude are unavailable, it is necessary to identify an optimal trade-off configuration for a realistic 60-year horizon. The selected configuration, corresponding to the 99th percentile, is identified by a red circle and will be strictly adopted in the subsequent analyses.

4.3.2 Non-Stationary Time Series

Non-Stationary Time Series with linear mean and quadratic variance

To test the performance of the non-stationary POT framework, a synthetic time series representing 100 years of hourly data was generated using a Weibull distribution with the initial baseline parameters. Non-stationarity was introduced into the scale parameter $A(t)$ by imposing a deterministic mean wind speed increase of 0.005 m/s per year.

The extraction of the extreme sample was performed using a time-varying threshold $u(t)$. This threshold was derived via quantile regression at the 99th percentile, yielding a linear increase described by an intercept $b_u = 6.1805$ m/s and a slope $a_u = 0.0147$ m/s/yr. The Bayesian MCMC framework was subsequently employed to estimate the parameters of a non-stationary GPD, assuming a constant shape parameter ξ and a linearly varying scale parameter $\sigma(t)$. As detailed in table 4.3 and visually confirmed by the diagnostic Q-Q plot in fig. 4.6, the numerical fit appears highly robust, displaying excellent performance metrics (e.g., $NSE \approx 0.999$).

According to the penultimate approximation theory, the true pre-asymptotic scale parameter of the excesses $\sigma(t)$ is non-linear, and the true shape parameter $\Lambda(t)$ is strictly dependent on the instantaneous threshold $u(t)$ rather than being a constant, as defined by eq. (2.17) and eq. (2.19). While the linear assumption for the scale parameter σ manages to approximate the true penultimate scale with a negligible mean relative error, forcing a constant shape parameter induces a severe misspecification bias. The MCMC-estimated constant shape ($\xi = -0.0590$) differs significantly from the dynamically evolving penultimate shape $\Lambda(t)$, resulting in a substantial mean relative error around 30%.

Despite the significant analytical error in the shape parameter, the impact on the Effective Return Levels (ERL) is remarkably contained. As shown in the top

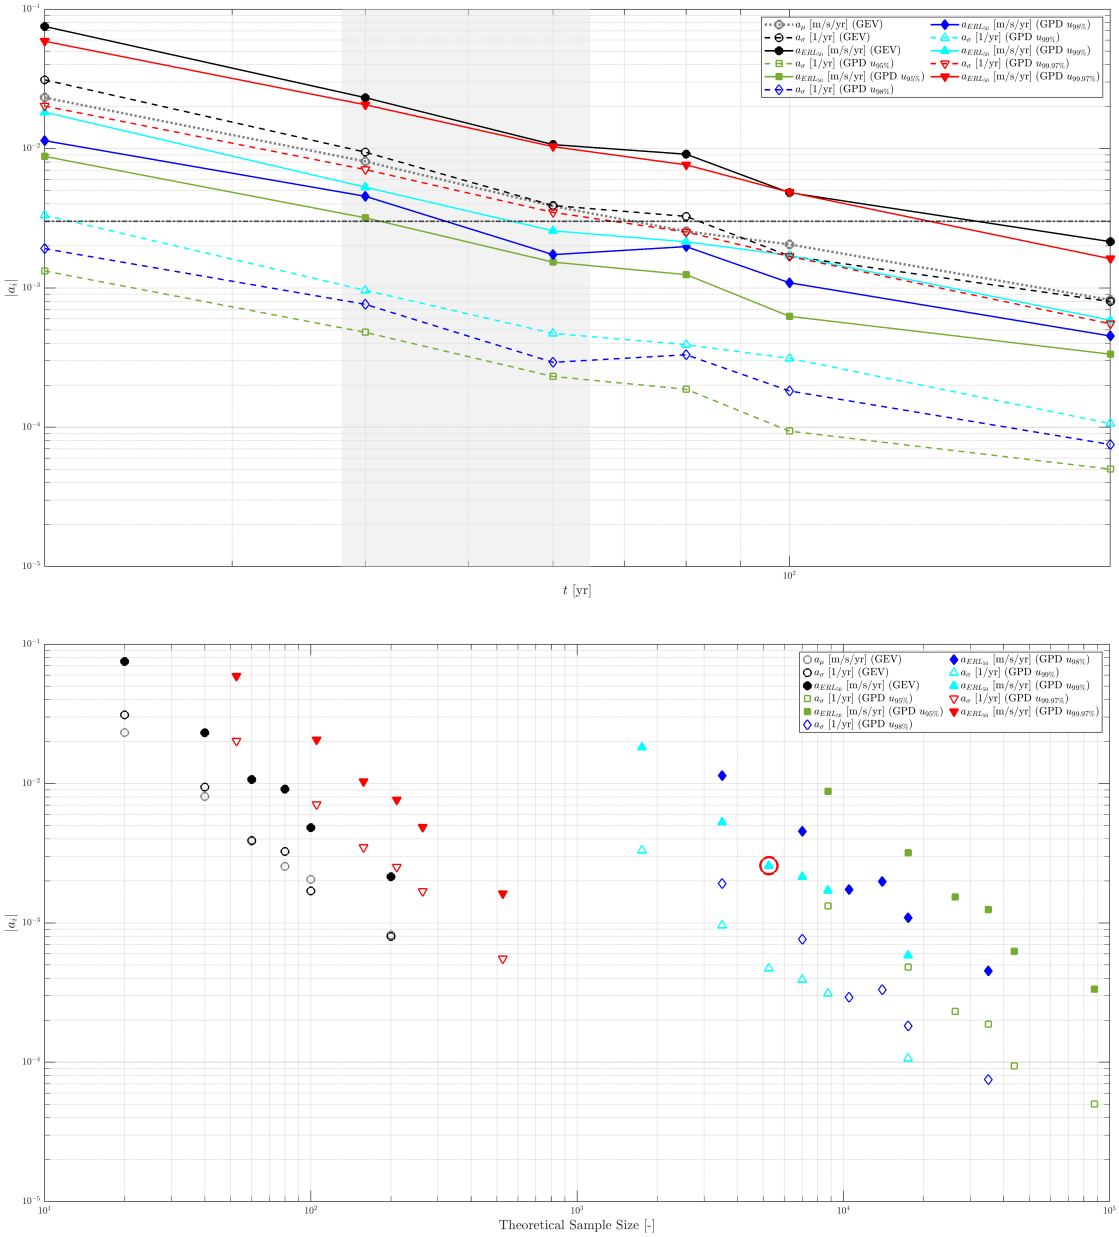

Figure 4.5: Convergence analysis of the parameter and ERL_{50} slopes for the GEV and GP distributions against sampled years (top) and theoretical sample size (bottom).

panel of fig. 4.6, the MCMC-estimated ERLs closely track the exact theoretical quantiles of the non-stationary parent distribution. The mean relative error remains bounded below 1.5%, demonstrating a strong robustness of the model.

Nevertheless, a clear trend emerges: the relative error increases alongside the return period. This behaviour is theoretically expected. Estimating higher return periods requires extrapolating deeper into the extreme tail of the distribution. In this deep-tail region, the true Weibull distribution exhibits a non-linear, pre-asymptotic curvature. Because the standard GPD model forces a rigid, constant shape parameter (ξ) onto the tail, it fails to capture this continuous bending. Consequently, the further the model extrapolates, the wider the gap becomes between the true curving tail and the rigid GPD approximation, progressively increasing the estimation error.

Table 4.3: Summary of Non-Stationary GPD Parameters, and Goodness-of-Fit Metrics, specifically for the test case featuring a linear mean and quadratic variance.

Parameter / Metric	**Estimate**	**95% CI width**
ξ	-0.0590	0.0393
b_σ	0.8461	0.0815
a_σ	0.0020	0.0014
b_u	6.1805	–
a_u	0.0147	–
AIC	15303.18	–
BIC	15324.41	–
RMSE	3.2465	–
NSE	0.9988	–

Non-Stationary Time Series with constant mean and quadratic variance

In conclusion, to further challenge the Bayesian framework and rigorously assess its robustness, the model misspecification case was evaluated. In this scenario, a 100-year synthetic time series was generated from a Weibull parent distribution featuring a constant mean and a time-varying quadratic variance. As in the previous case, the non-stationary POT approach was applied, extracting extremes via a 99th percentile dynamic threshold $u(t)$, which was approximated linearly with an intercept $b_u = 6.0663$ m/s and a slope $a_u = 0.0148$ m/s/yr. The MCMC framework subsequently estimated the parameters of a non-stationary GPD, again enforcing a constant shape parameter ξ and a linearly varying scale parameter $\sigma(t)$, as summarised in table 4.4 and visually depicted in fig. 4.8.

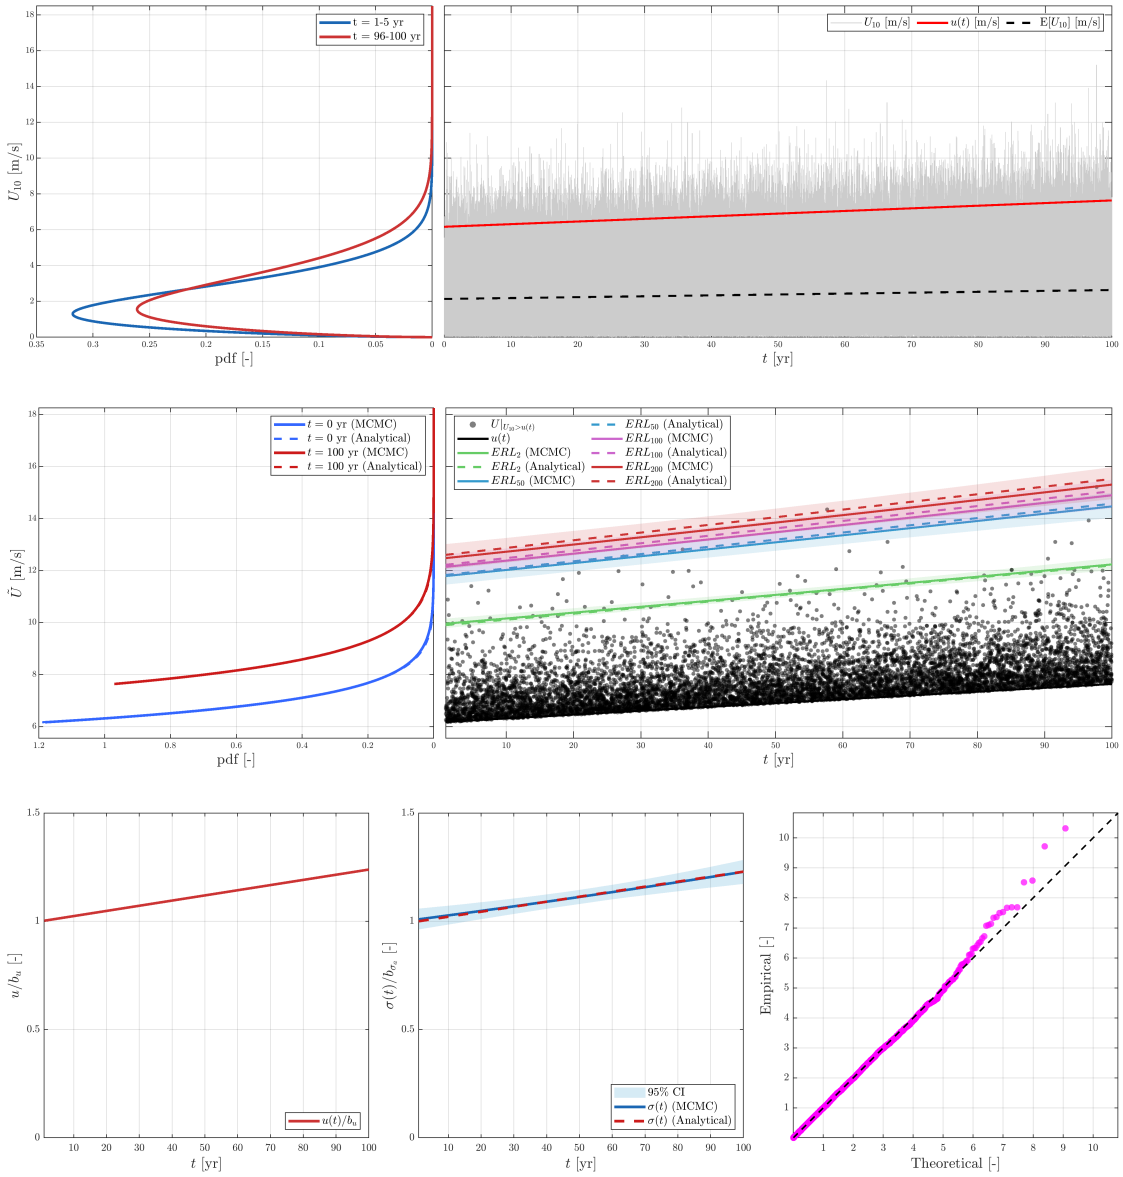

Figure 4.6: Non-stationary MCMC framework results for the GPD model against analytical framework, specifically for the test case featuring a linear mean and quadratic variance. Top row: from left to right the evolution of the GPD from year 0 to year 100 and the effective return level. Bottom row: from left to right, the panels illustrate the temporal evolution of the normalised distribution parameters (dynamic threshold and scale), and the diagnostic Q-Q plot.

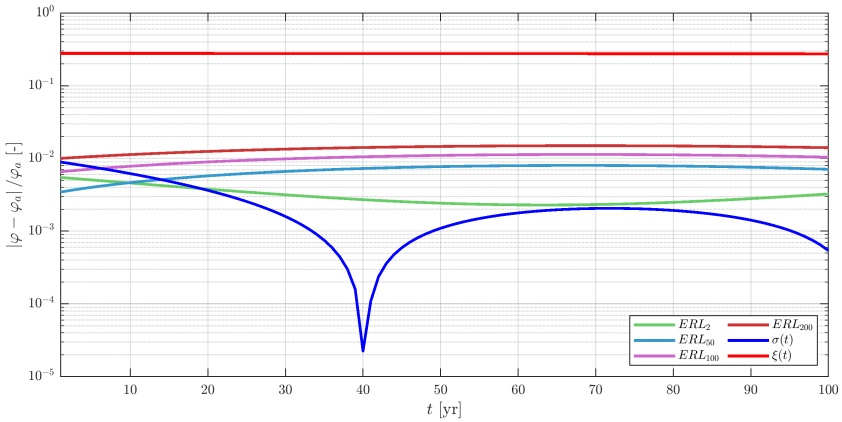

Figure 4.7: Relative errors of the estimated parameters and effective return levels against time, specifically for the test case featuring a linear mean and quadratic variance.

In this specific environment, the structural rigidity of the prescribed model induces a misspecification bias. Because the parent distribution's variance grows quadratically, the true pre-asymptotic scale of the extremes and the resulting analytical Effective Return Levels (ERL) naturally exhibit a non-linear, parabolic trajectory over time. However, the non-stationary framework is intentionally constrained to fit purely linear temporal trends. Forcing a straight line to approximate a parabolic curve inevitably degrades the numerical fit. Consequently, the relative errors plotted in fig. 4.9 are, for some observed years, one order of magnitude higher than those observed in the strictly linear mean scenario.

Despite operating in this severely misspecified environment, the Bayesian framework still captures the general upward trend of the hazard. While the point-wise estimates of the high-return-period ERLs deviate from the true analytical values at the extremes of the time window, the analytical values remain securely confined within the 95% credible intervals computed by the model. By successfully detecting the intensification of the extremes, the framework demonstrates a remarkable degree of practical engineering robustness, even when the underlying physical processes violate the assumed linear models.

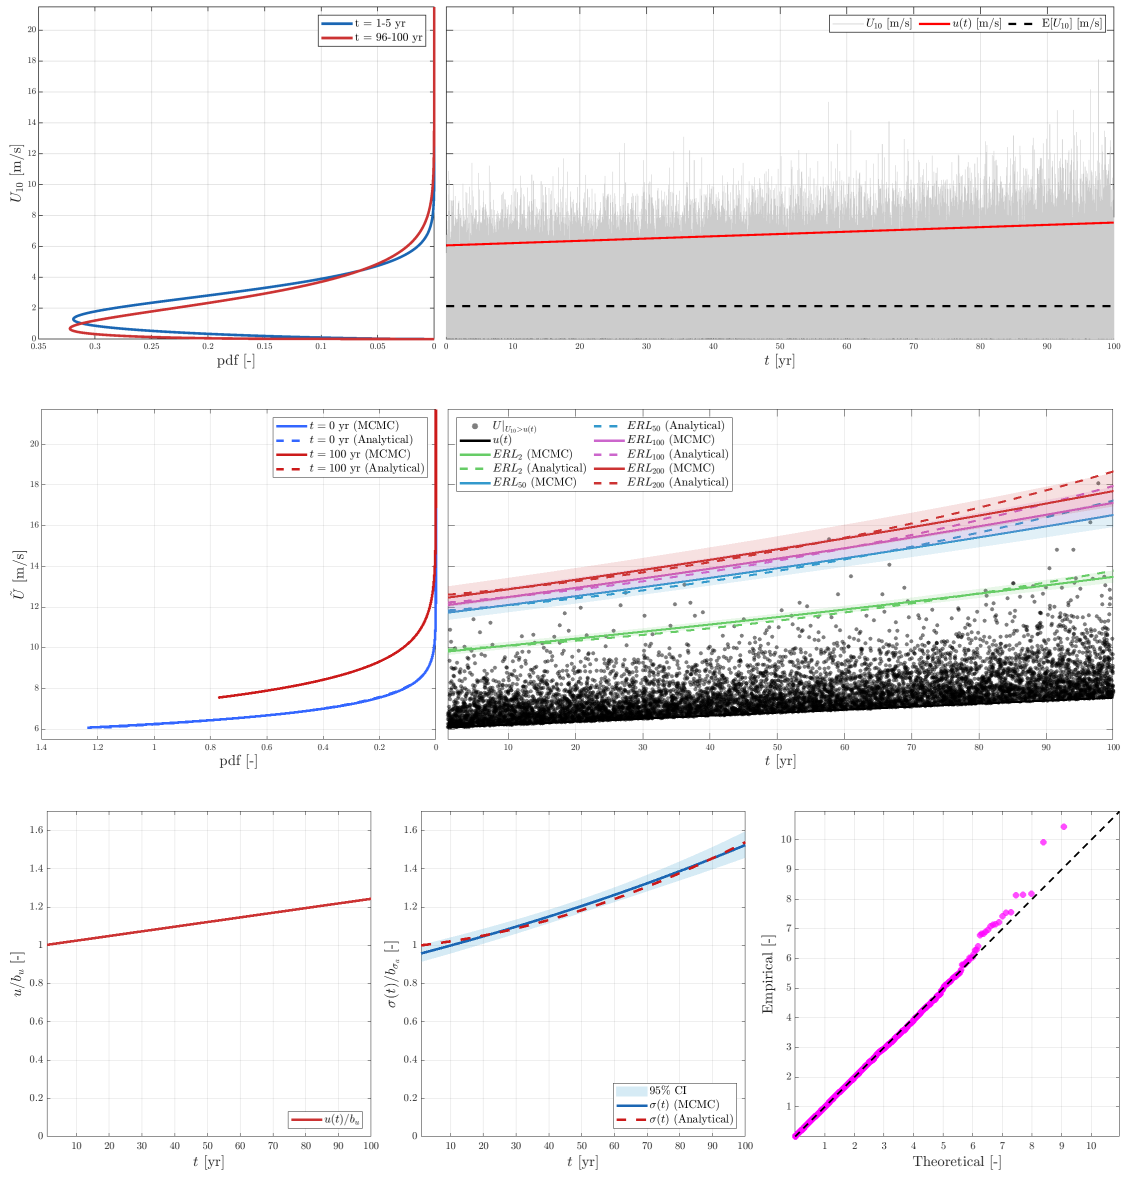

Figure 4.8: Non-stationary MCMC framework results for the GPD model against analytical framework, specifically for the test case featuring a constant mean and quadratic variance. Top row: from left to right the evolution of the GPD from year 0 to year 100 and the effective return level. Bottom row: from left to right, the panels illustrate the temporal evolution of the normalised distribution parameters (dynamic threshold and scale), and the diagnostic Q-Q plot.

Table 4.4: Summary of Non-Stationary GPD Parameters, and Goodness-of-Fit Metrics, specifically for the test case featuring a constant mean and quadratic variance.

Parameter / Metric	**Estimate**	**95% CI width**
ξ	-0.0465	0.0399
b_σ	0.8087	0.0767
a_σ	0.0047	0.0014
b_u	6.0663	–
a_u	0.0148	–
AIC	17102.54	–
BIC	17123.77	–
RMSE	3.1119	–
NSE	0.9989	–

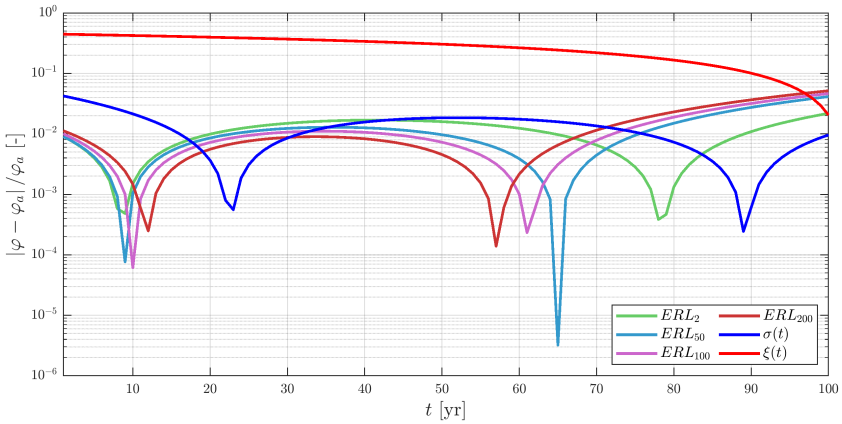

Figure 4.9: Relative errors of the estimated parameters and effective return levels against time, specifically for the test case featuring a constant mean and quadratic variance.

Chapter 5

Application to real world wind speed time series

This concluding chapter bridges the theoretical statistical framework developed in the previous sections with real-world applications, thoroughly investigating extreme wind behaviour under the combined, evolving pressures of anthropic land consumption and climate change. To achieve this, the analysis relies on two different datasets.

First, historical time series were analysed from selected anemometric stations [17], establishing a baseline of extreme wind dynamics under historically documented environmental and land-use changes, based on data provided by the Italian National Institute of Statistics (ISTAT). Second, to model future non-stationarities driven by climate change, a high-resolution meteorological dataset was adopted: the *Climate Projections RCP4.5 and RCP8.5 downscaled at 2.2 km over Italy*, developed within the framework of the Highlander project by the Euro-Mediterranean Center on Climate Change (CMCC) [12], by dynamically downscaling a global climate model. The downscaling process adopted the Representative Concentration Pathway (RCP) 4.5 and 8.5 scenarios defined by the Intergovernmental Panel on Climate Change (IPCC), and was conducted through the COSMO-CLM regional climate model, provided by the CMCC.

5.1 Time series measured at anemometric stations under changing environmental conditions

Analysing historical wind data requires careful consideration of the evolving local environment. Indeed, wind speed is strictly dependent on the surface roughness

of the surrounding region. In particular, within the internal boundary layer, the mean wind speed follows a logarithmic profile:

$$U(z) = \frac{1}{\kappa} u_* \ln\left(\frac{z}{z_0}\right), \tag{5.1}$$

where κ is the Von Karman constant (approximately 0.41), u_* represents the friction velocity (which can be derived from wind speed measurements at multiple heights), and z_0 denotes the surface roughness length, dictated by the local topography. This physical relationship implies that any modification of the surrounding landscape directly alters z_0, thereby modifying the wind speed recorded at a fixed anemometric height. For instance, increasing urbanization in the vicinity of a station systematically increases surface roughness, leading to a localized decrease in the recorded mean wind speed.

Consequently, to investigate the actual impact of such anthropic land consumption on extreme wind speeds, the Mann-Kendall and White tests were performed on the exceedances over the 99th regression quantile on the 10-minute average wind speed historical anemometric measurements, spanning from the beginning of 1981 to the end of 2022, revealing an overall non-stationary trend of the wind speed and of its variance, as shown in table 5.1. To explore this, an overview of two selected anemometric stations is provided. Based on the statistical tests summarized, two representative locations were deliberately chosen to contrast distinct long-term behaviours: Catania Fontanarossa, serving as the almost-stationary case, and Roma Ciampino, representing the non-stationary case.

These specific locations were also investigated by Raffaele et al. [13]. Specifically, Roma Ciampino represents a documented case of progressive roughness increase driven by urbanization, analogous to the rapid development observed in cities like Beijing or Harbin in the PRC (see, e.g., [14, 15]). This urban expansion is visually corroborated by satellite imagery. Figures 5.1 and 5.2 illustrate the land cover at spatial resolutions of 50 km and 2 km, respectively. While the area surrounding Catania shows no significant signs of increasing urbanization, the urban sprawl around Roma Ciampino over the 40-year observation period is clearly visible. This qualitative observation is further quantified by demographic and land-use data provided by the Italian National Institute of Statistics (ISTAT). As shown in fig. 5.3, the percentage of built surface area, denoted by r, remains almost constant for Catania, whereas it exhibits a marked increase during the analysed period for Roma Ciampino.

The proposed non-stationary Bayesian framework was applied to these two specific stations, and the resulting parameter estimations are presented in table 5.2. For the Catania Fontanarossa station, the threshold parameter remains constant over the years, accompanied by a near-zero slope for the scale coefficient. Conversely, Roma Ciampino exhibits distinct non-stationary features in both threshold and

Table 5.1: Mann-Kendall (MK) and White test results (p-value) and percentage variation of the built surface (Δr) for the analysed anemometric stations. The statistical tests refer to the 1981-2022 period, while Δr is calculated for the 1991-2011 period.

Station	**MK**	**White**	Δr [%]
Bari Palese	$1.87 \cdot 10^{-5}$ (H_0 rejected)	$8.63 \cdot 10^{-9}$ (H_0 rejected)	−6.78
Cagliari Elmas	0 (H_0 rejected)	$3.54 \cdot 10^{-11}$ (H_0 rejected)	166.67
Catania Fontanarossa	0.628 (H_0 accepted)	$3.38 \cdot 10^{-4}$ (H_0 rejected)	0.29
Firenze Peretola	$1.46 \cdot 10^{-9}$ (H_0 rejected)	0.019 (H_0 rejected)	−4.27
Genova Sestri	0.100 (H_0 accepted)	0.594 (H_0 accepted)	2.16
Grosseto	$2.05 \cdot 10^{-7}$ (H_0 rejected)	$7.77 \cdot 10^{-16}$ (H_0 rejected)	12.20
Olbia Costa Smeralda	$1.54 \cdot 10^{-7}$ (H_0 rejected)	0.042 (H_0 rejected)	43.10
Pescara	$3.56 \cdot 10^{-5}$ (H_0 rejected)	0.021 (H_0 rejected)	−1.53
Pisa	$2.53 \cdot 10^{-5}$ (H_0 rejected)	$4.15 \cdot 10^{-6}$ (H_0 rejected)	16.42
Reggio Calabria	$1.08 \cdot 10^{-4}$ (H_0 rejected)	0.034 (H_0 rejected)	25.89
Roma Ciampino	$1.77 \cdot 10^{-8}$ (H_0 rejected)	$1.77 \cdot 10^{-7}$ (H_0 rejected)	7.66
Roma Fiumicino	$1.60 \cdot 10^{-6}$ (H_0 rejected)	0.134 (H_0 accepted)	0.00
Torino Caselle	0.860 (H_0 accepted)	0.010 (H_0 rejected)	43.18
Venezia Tessera	0 (H_0 rejected)	$7.80 \cdot 10^{-6}$ (H_0 rejected)	0.00

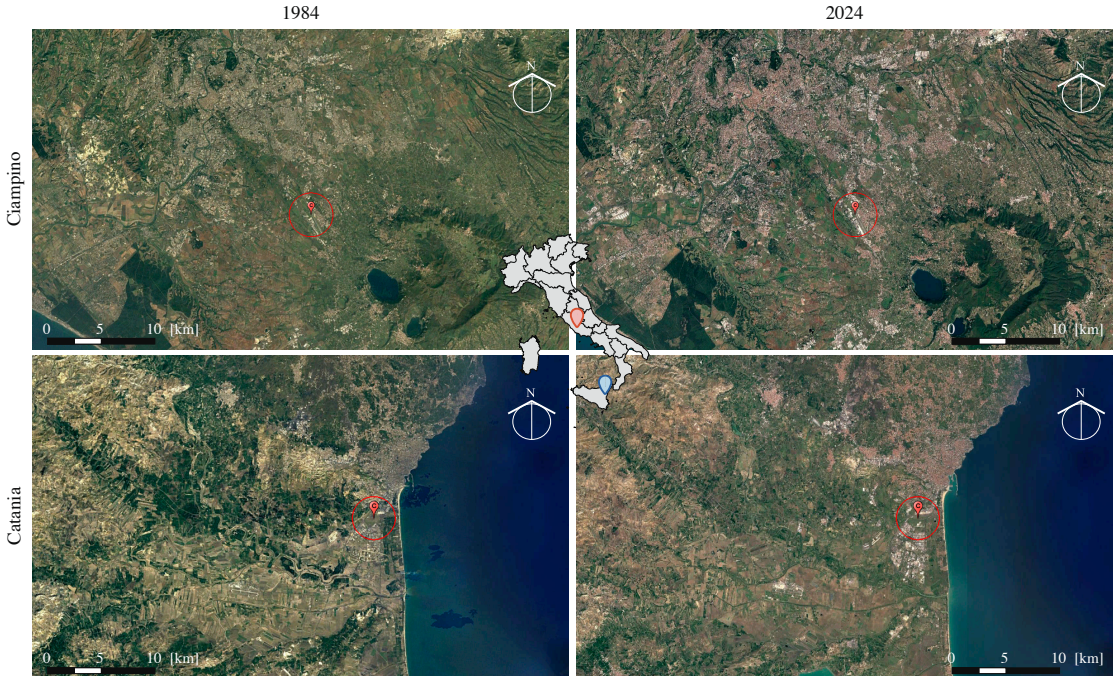

Figure 5.1: View of the land use around Ciampino and Catania in 1984 and 2024 (satellite imagery Landsat / Copernicus).

scale parameters. A comprehensive visual representation of these results—including the parent distribution, the time series, probability density functions (PDF), Effective Return Levels (ERLs), parameter evolution, and Q-Q plots—is provided in figs. 5.4 and 5.5. Consistent with the parameter estimates, the effective return levels for Catania remain relatively stable over time. In contrast, the ERLs for Roma Ciampino decrease progressively, aligning with the increased aerodynamic roughness caused by land consumption.

To qualitatively assess the relationship between urban growth and extreme wind metrics, fig. 5.6 presents the temporal trends of the built surface percentage, r, alongside ERL_{50}, as well as their scatter plot, for the two selected anemometric stations. Since the ISTAT data and the anemometric records cover different time spans, this comparison is restricted to the overlapping period from 1981 to 2011, with the built surface data linearly interpolated to obtain continuous annual values.

The observed inverse relationship between the increasing built surface and the decreasing extreme wind metric for Roma Ciampino is phenomenologically consistent with the expected behaviour of the logarithmic profile introduced in Eq. (5.1). Similarly, the overall constant trend exhibited by Catania Fontanarossa aligns with its negligible urban expansion over the same period.

Table 5.2: Estimated Non-Stationary GPD Parameters (MAP) and Diagnostic Metrics, for the time series measured.

Parameter / Metric	**Catania Fontanarossa**	**Roma Ciampino**
ξ (95% CI)	-0.0391 (0.0544)	0.0034 (0.0585)
a_σ (95% CI)	0.0008 (0.0058)	-0.0123 (0.0069)
b_σ (95% CI)	1.6662 (0.2548)	2.2953 (0.3832)
a_u	0.0000	-0.0762
b_u	9.8000	11.9161
RMSE	8.9545	6.1991
NSE	0.9770	0.9894

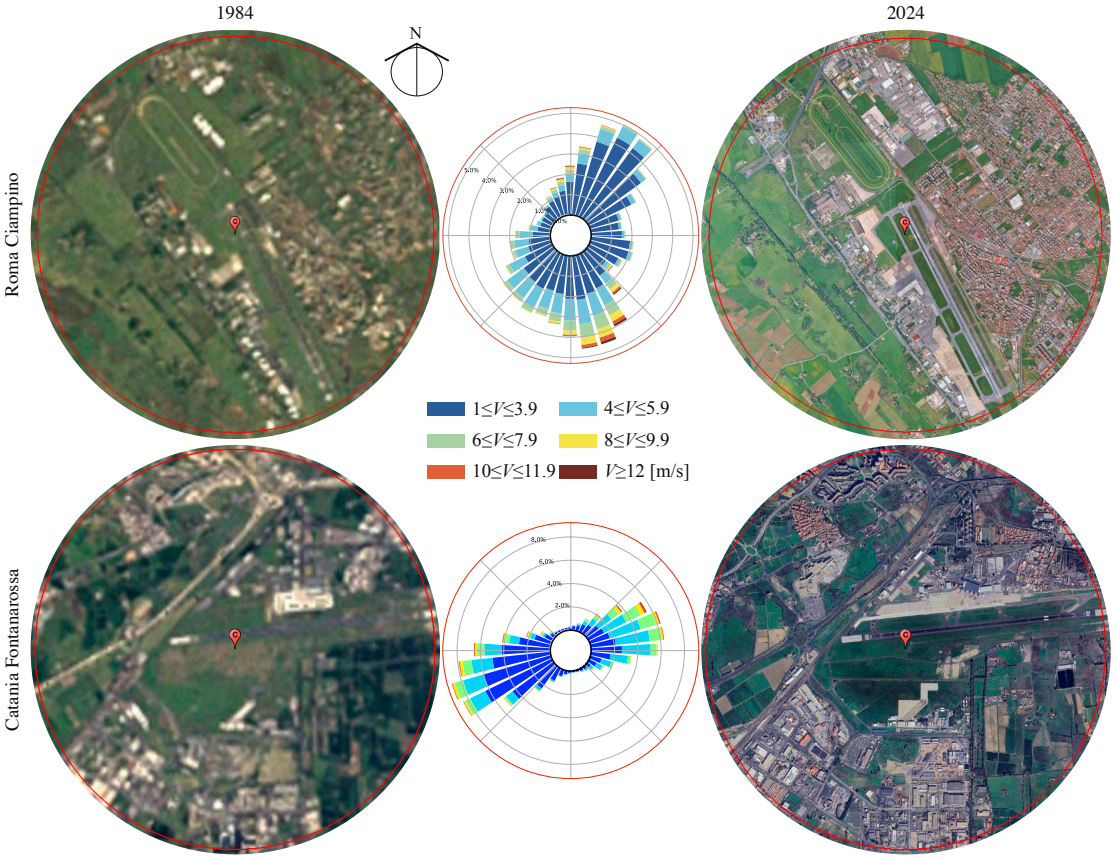

Figure 5.2: Closed view (2 km radius fetch) of the land use around the anemometric stations at Roma Ciampino and Catania Fontanarossa in 1984 and 2024 (satellite imagery Landsat / Copernicus). Wind roses after anemometric records of U_{10} at the stations from Jan 1981 to Dec 2022.

However, these concurrent trends should be interpreted with caution, as a unique dependency of extreme wind speeds on land surface roughness cannot be perfectly isolated. Several intrinsic limitations and confounding factors must be acknowledged. First, the built surface data r represents a macroscopic spatial average over the entire municipality that smooths out the actual inhomogeneous building density. Therefore, it is not a localized measure confined within the strictly relevant aerodynamic fetch (e.g., a 2-km radius), meaning the microscopic aerodynamic effects of buildings immediately surrounding the anemometer are neglected. Second, the current framework does not explicitly separate the effects of land-use modifications from other confounding variables, such as the macro-scale impacts of climate change or inherent instrumental uncertainties (e.g., periodic anemometer

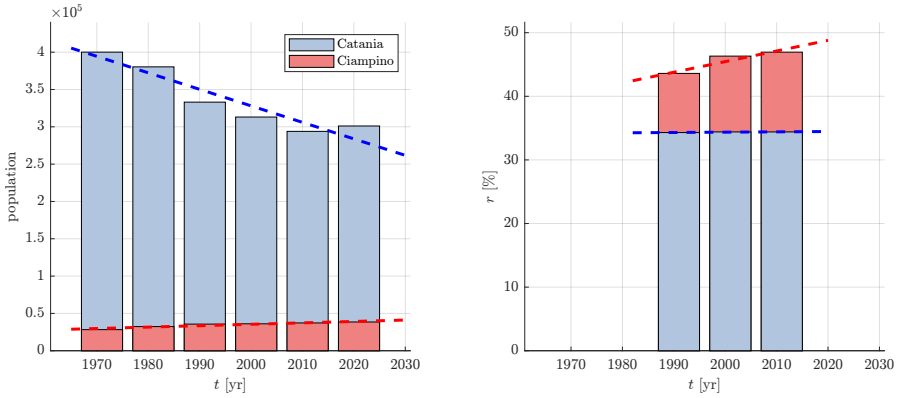

Figure 5.3: Trend of the population and of the percentage of built surface over total one at Catania and Ciampino from 1970 to 2020 (data source: Italian National Institute of Statistics - ISTAT).

replacements, sensor degradation, and temporary malfunctions). Consequently, within this treatment, any variation in the extreme wind metrics actually induced by climatic trends or measurement errors is implicitly absorbed by, and attributed to, the increase of z_0.

In conclusion, to provide a broader geographical perspective on extreme wind behaviour, maps of the Italian Peninsula illustrating the ERL_{50} temporal slope and percentage trend across all 14 analysed stations in table 5.1 are presented in figs. 5.7 and 5.8, respectively. The percentage trend of the ERL_{50} was calculated with the following formula:

$$\Delta ERL_{50} = \frac{ERL_{50}(t_{\text{end}}) - ERL_{50}(t_0)}{ERL_{50}(t_0)} \cdot 100.$$

Overall, the national spatial pattern reveals a prevailing decreasing trend in both slope and percentage across the majority of the territory. However, notable exceptions emerge from this national tendency. Specifically, Torino and Catania exhibit fundamentally constant behaviours, whereas Olbia displays a marked increasing trend of approximately 21%.

Importantly, the inverse relationship between extreme wind metrics and urbanization is not universally consistent across all stations, practically highlighting the methodological complexities discussed in the previous paragraph. For instance Bari exhibits a decrease in both built surface ($\Delta r = -6.78\%$) and ΔERL_{50} (-19%), while Olbia presents a notable increase in extreme winds despite a significant urban expansion ($\Delta r = 43.10\%$). These discrepancies strongly suggest that, in such locations, the aerodynamic effects of local land-use are overshadowed by other

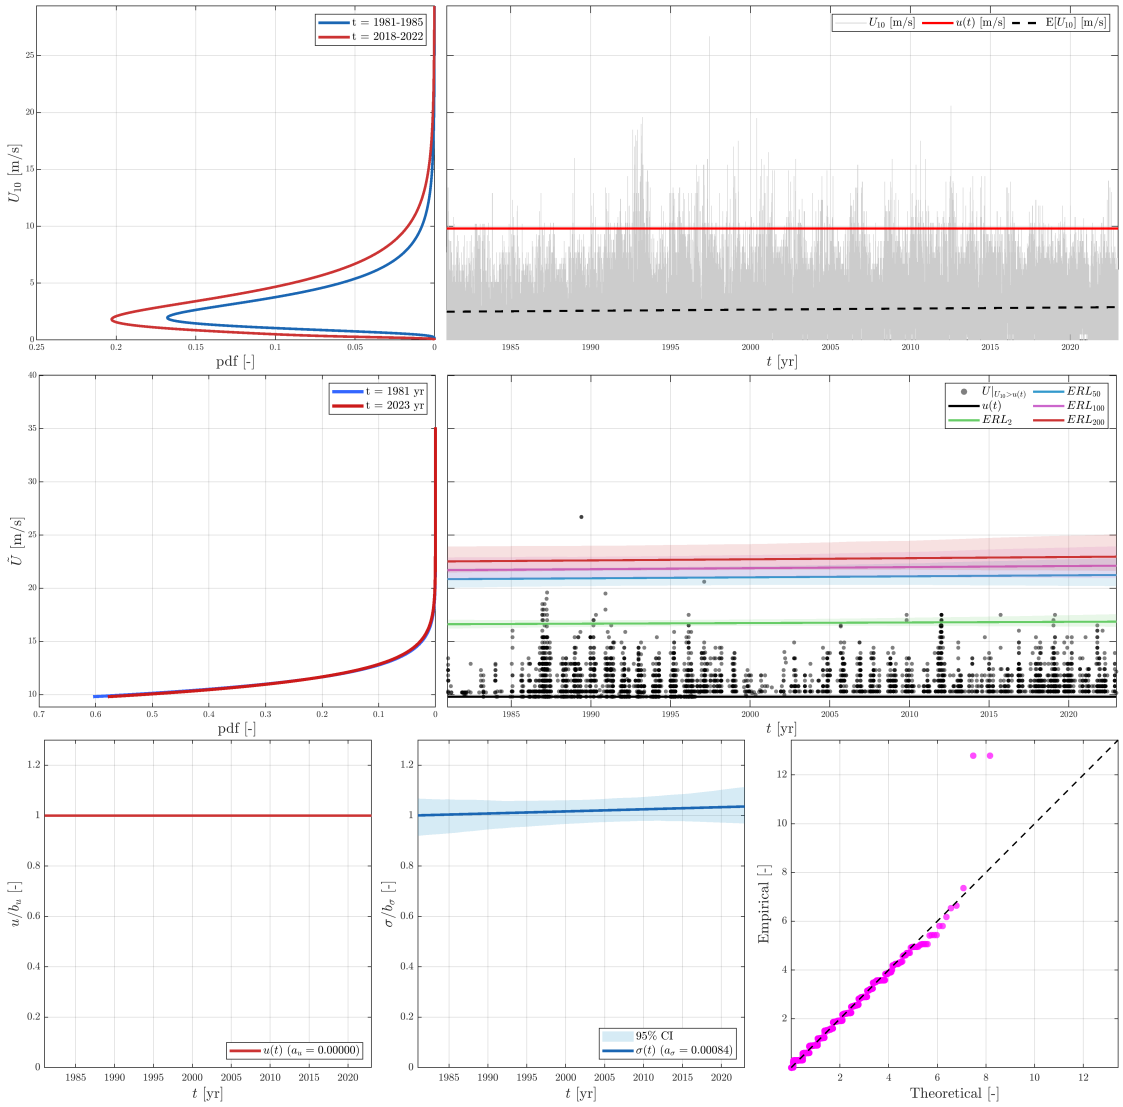

Figure 5.4: Results of the non-stationary Generalized Pareto Distribution (GPD) analysis for Catania Fontanarossa. The top panel displays the temporal evolution of the Probability Density Function (PDF) and Effective Return Level (ERL), while the bottom panel illustrates the parameter evolution and the corresponding Q-Q plot.

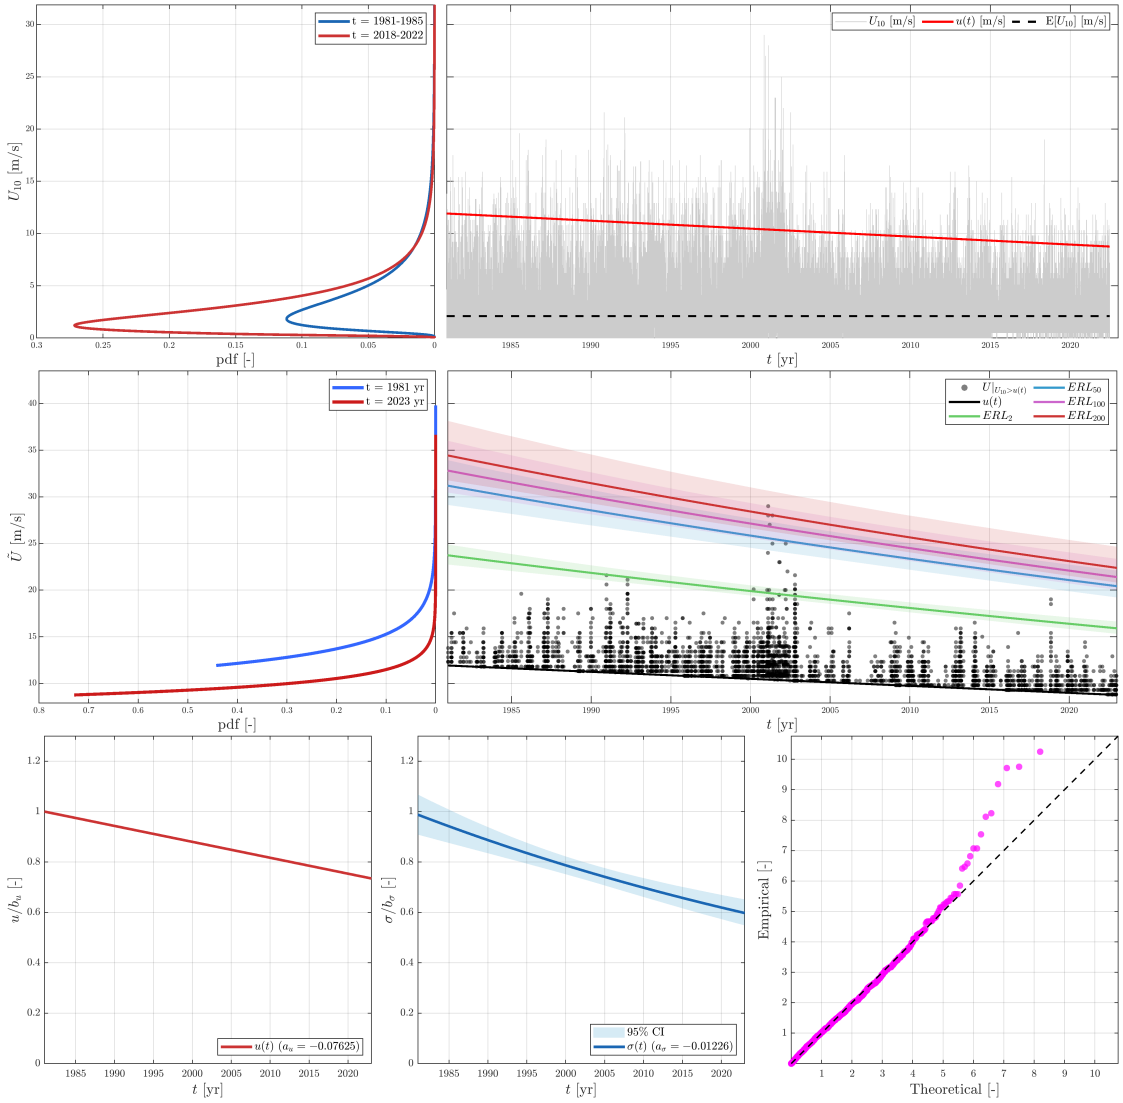

Figure 5.5: Results of the non-stationary Generalized Pareto Distribution (GPD) analysis for Rome Ciampino. The top panel displays the temporal evolution of the Probability Density Function (PDF) and Effective Return Level (ERL), while the bottom panel illustrates the parameter evolution and the corresponding Q-Q plot.

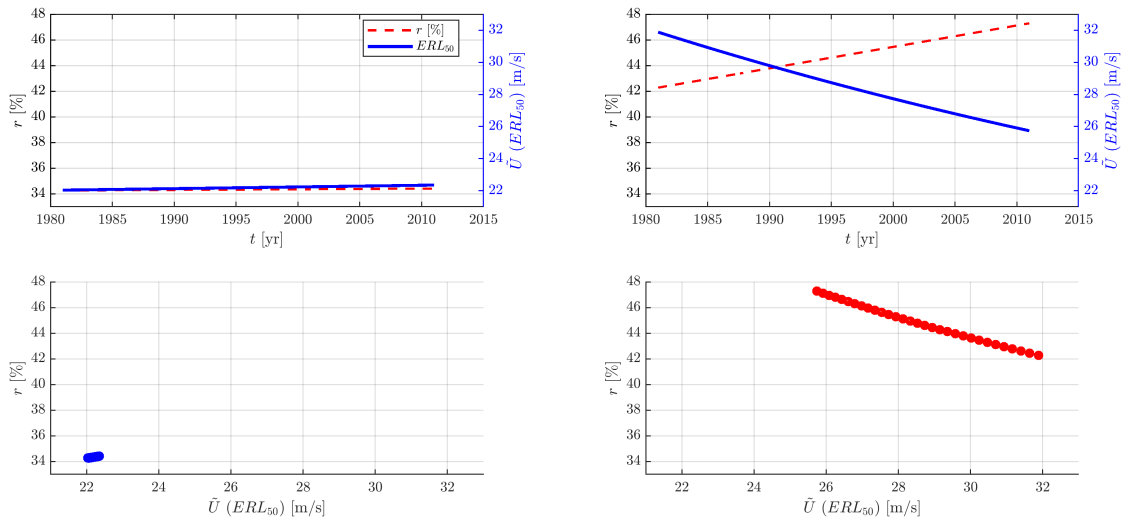

Figure 5.6: Trend of the ERL_{50} and built surface percentage (r) in the top panels (Catania Fontanarossa on the left, Roma Ciampino on the right). Scatter plot of ERL_{50} versus built surface percentage in the bottom panels.

confounding factors, such as meso-scale climatic shifts or unrecorded instrumental errors.

Furthermore, the case of Torino Caselle perfectly illustrates the spatial limitations of using municipality-wide data. Although Torino experienced substantial urban growth ($\Delta r = 43.18\%$), its extreme wind metric remained essentially constant. This apparent paradox is explained by the spatial distribution of the development: major urban expansions occurred kilometres away from the Caselle anemometer. Consequently, the strictly localized aerodynamic fetch remained largely unaltered, completely shielding the anemometric measurements from the macroscopic land-use modifications of the broader municipality.

As these examples demonstrate, it proves inherently challenging to definitively isolate the exact driving causes of the observed wind speed attenuation and correlate them strictly with land-use modifications. Furthermore, a final temporal limitation must be explicitly acknowledged: whereas the extreme wind trends and the corresponding geographical maps encompass the entire 1981–2022 period, the available built surface data used to calculate Δr is restricted to the 1991–2011 time frame. Consequently, the potential aerodynamic impact of any further urban expansion that occurred during the last decade of anemometric measurements remains unaccounted for in this direct comparison.

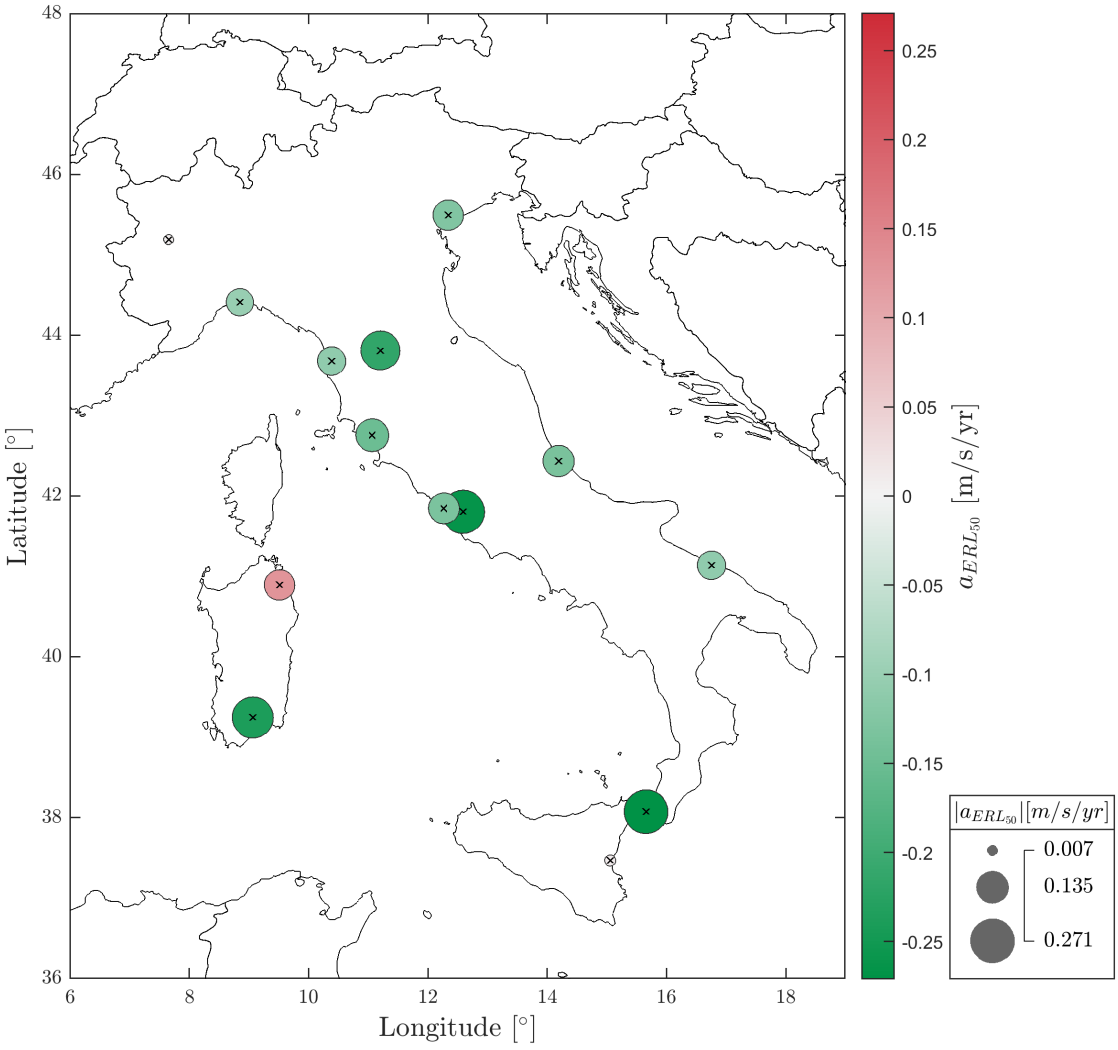

Figure 5.7: Bubble map of the 50-year effective return level (ERL_{50}) trend for the selected anemometric stations.

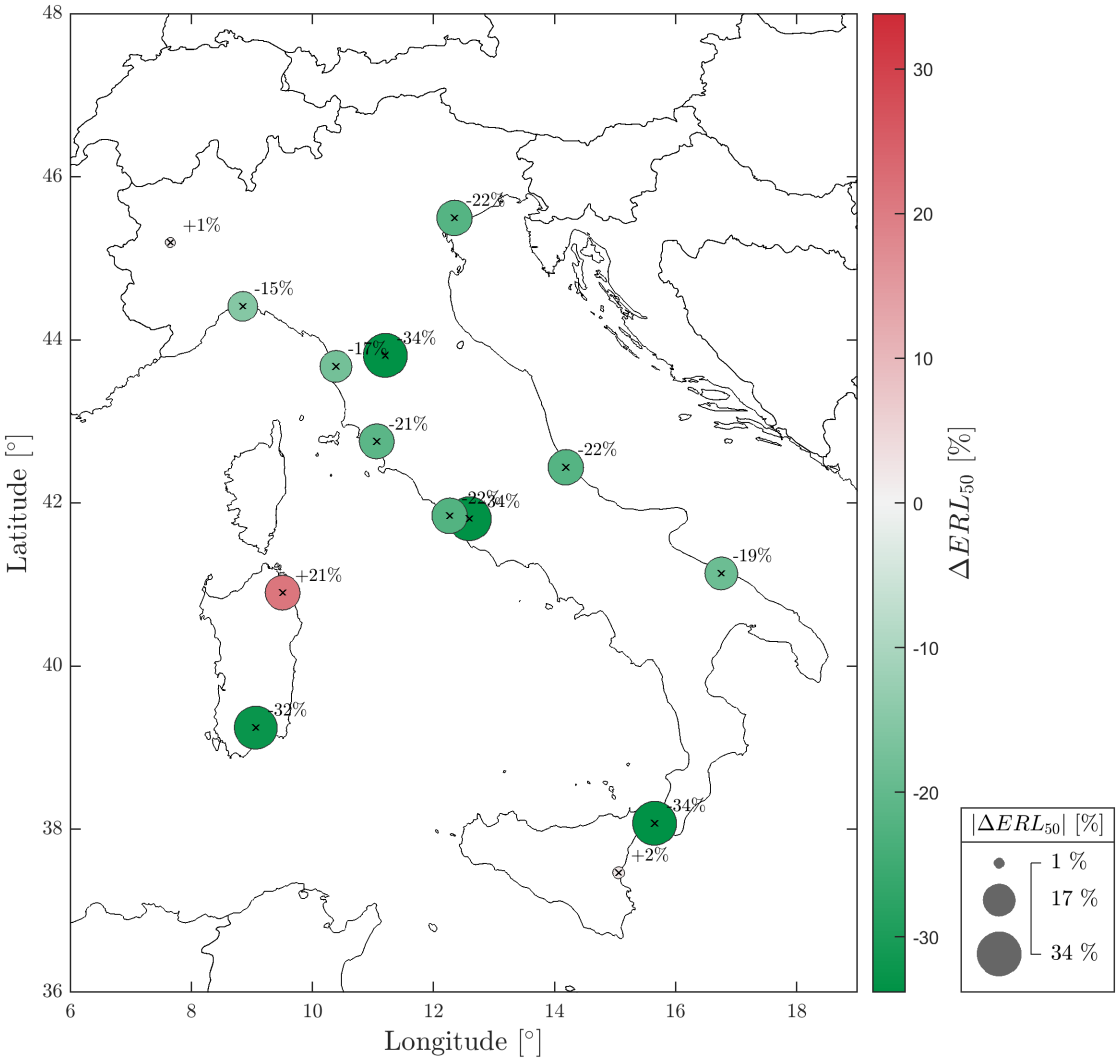

Figure 5.8: Bubble map of the 50-year effective return level (ERL_{50}) percentage for the selected anemometric stations.

5.2 Time series simulated over Italy under climate change

To investigate the future evolution of extreme wind speeds under the influence of climate change, the proposed statistical framework was applied to high-resolution climate projections. The selected dataset comprises projected hourly 10-min average wind speed time series spanning from 2005 to 2070. The analysis initially focused on two representative locations, deliberately chosen to illustrate contrasting behaviours under different emission scenarios: one situated in southern Sardinia and the other in the vicinity of Venice.

Based on the projection dataset, the Mann-Kendall (MK) and White tests were performed on the exceedances over the 99th regression quantile, with the results summarized in table 5.3. For the location in southern Sardinia, the tests indicate a fully non-stationary behaviour under the RCP 4.5 scenario (both H_0 rejected), whereas a stationary behaviour is expected under the RCP 8.5 scenario. Conversely, the Venice location exhibits a decoupled non-stationarity: under RCP 4.5, the extremes show a stationary monotonic trend but non-stationary variance, while under RCP 8.5, they exhibit a non-stationary monotonic trend with a homoscedastic variance.

Subsequently, the non-stationary Bayesian analysis was conducted, and the resulting parameter estimates are reported in table 5.4. A comprehensive visual representation of this analysis—including the time series, Effective Return Levels (ERLs), parent and estimated probability density functions, parameter evolutions, and Q-Q plots—is provided in figs. 5.9 and 5.10, comparing the RCP 4.5 (left column) and RCP 8.5 (right column) scenarios.

The estimated GPD parameters strongly align with the preliminary statistical tests. For southern Sardinia, the results indicate an overall increasing trend in the ERLs under the RCP 4.5 scenario, driven by a positive slope in the scale parameter. Conversely, the projected trend under RCP 8.5 remains almost stationary, with near-zero slopes for both threshold and scale. For the Venice location, an opposite behaviour emerges: an almost stationary trend is predicted under RCP 4.5 (characterized by a constant threshold), while a marked decreasing trend in the estimated metrics emerges under RCP 8.5, primarily driven by notably negative slopes in the threshold and scale parameters.

To gain broader spatial insight into the future evolution of extreme winds across the Italian Peninsula, a novel national-scale mapping approach was developed. Maps for both the RCP 4.5 and RCP 8.5 scenarios were generated using the selected projection data over a high-resolution spatial grid comprising 156 points across Italy. The non-stationary Bayesian inference model was systematically applied to these projections. The Markov Chain Monte Carlo (MCMC) algorithm

Table 5.3: Results of the Mann-Kendall (monotonic trend) and White (heteroscedasticity) non-stationarity tests for the RCP 4.5 and 8.5 projections. For each test, the p-value and test decision regarding the null hypothesis H_0 are reported at a significance level of $\alpha = 0.05$.

Tests	**RCP 4.5**	**RCP 8.5**
Location: 39.5°N, 9.5°E		
MK	0.0005 (H_0 rejected)	0.9044 (H_0 accepted)
White	0.0110 (H_0 rejected)	0.0849 (H_0 accepted)
Location: 45.5°N, 12.5°E		
MK	0.2938 (H_0 accepted)	0.0363 (H_0 rejected)
White	0.0080 (H_0 rejected)	0.3492 (H_0 accepted)

Table 5.4: Estimated GPD parameters (MAP) and performance metrics for each scenario.

Parameter / Metric	**RCP 4.5**	**RCP 8.5**
Location: 39.5°N, 9.5°E		
ξ (95% CI)	-0.1212 (0.0464)	-0.0523 (0.0499)
b_σ (95% CI)	1.4915 (0.1650)	1.4611 (0.1775)
a_σ (95% CI)	0.0022 (0.0026)	-0.0001 (0.0029)
b_u	12.5224	12.5654
a_u	-0.0014	-0.0024
RMSE	3.4439	4.1073
NSE	0.9979	0.9970
Location: 45.5°N, 12.5°E		
ξ (95% CI)	-0.1144 (0.0451)	-0.1310 (0.0424)
b_σ (95% CI)	1.6377 (0.1816)	1.6630 (0.1742)
a_σ (95% CI)	-0.0009 (0.0026)	-0.0012 (0.0025)
b_u	12.3789	12.4431
a_u	0.0000	-0.0095
RMSE	4.1182	4.2953
NSE	0.9970	0.9967

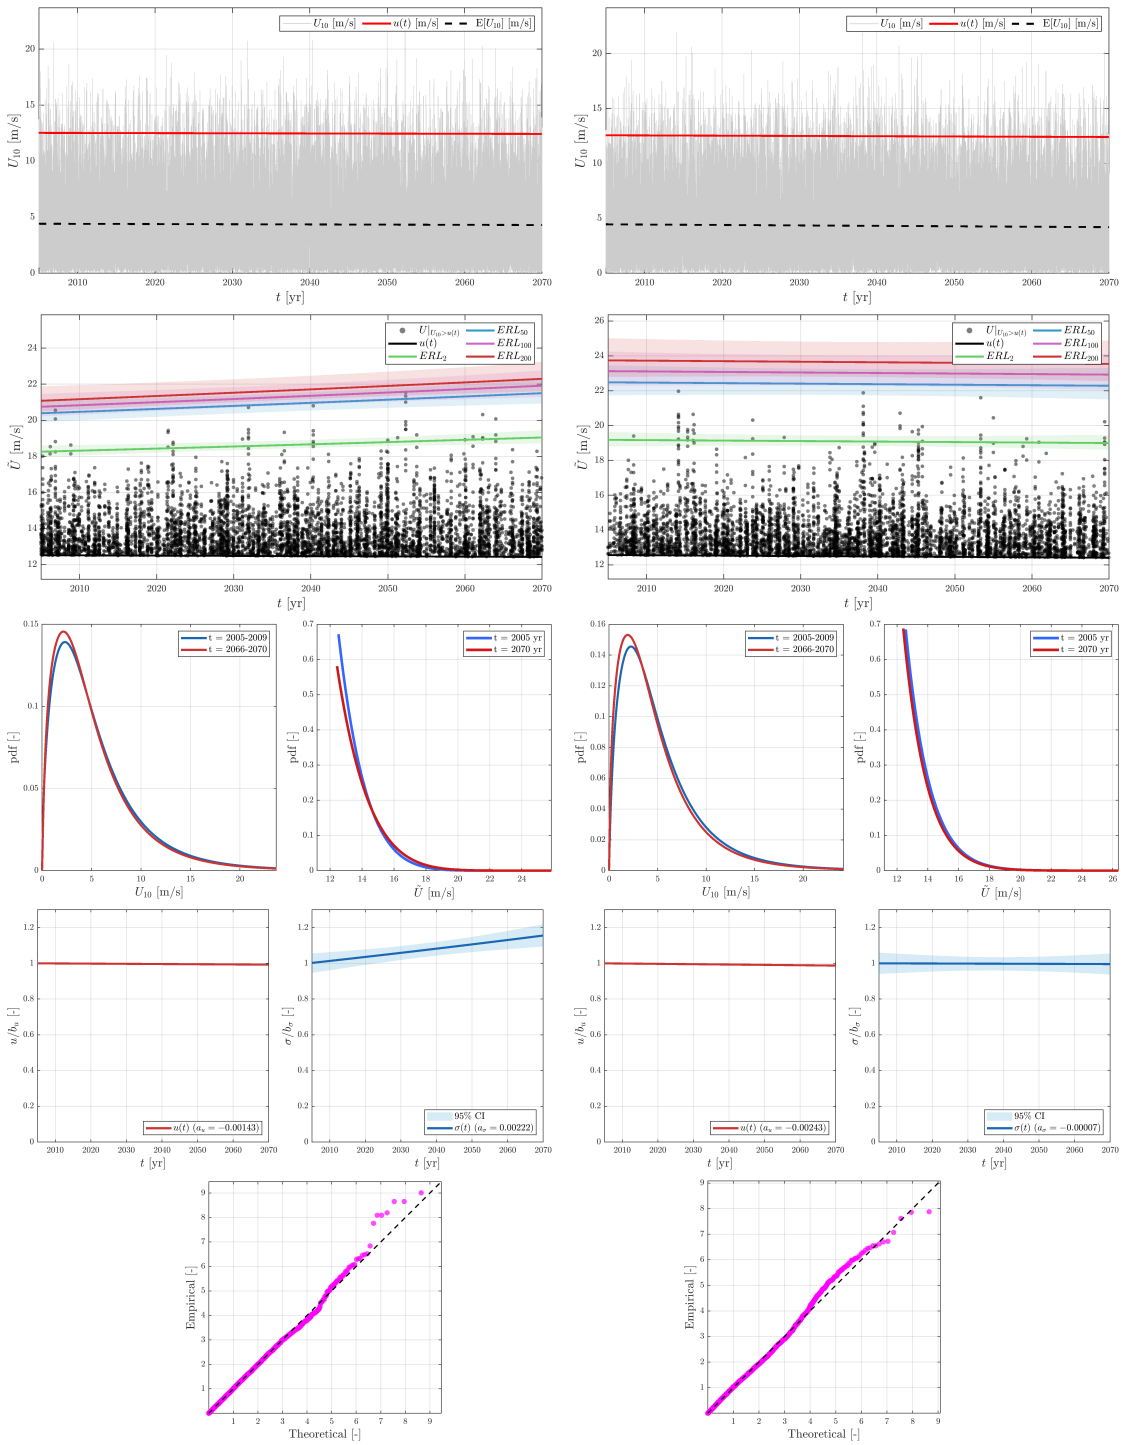

Figure 5.9: Comprehensive non-stationary extreme value analysis comparing RCP 4.5 (left column) and RCP 8.5 (right column) emission scenarios, for the location 39.5°N, 9.5°E (South Sardinia).

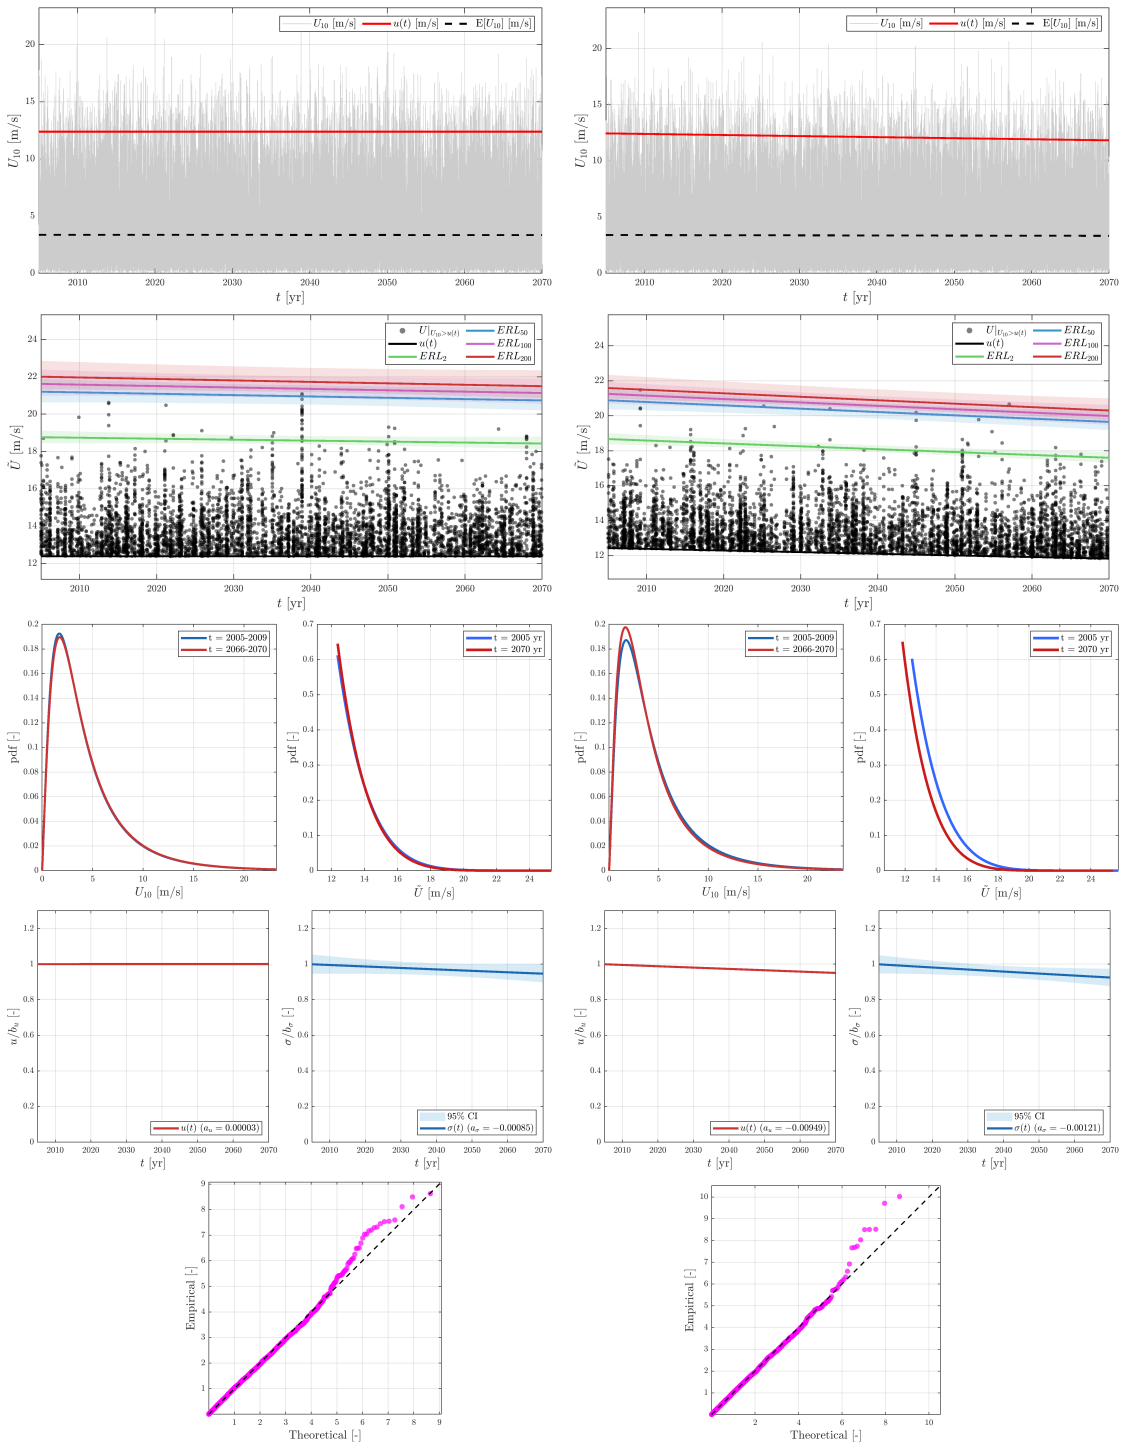

Figure 5.10: Comprehensive non-stationary extreme value analysis comparing RCP 4.5 (left column) and RCP 8.5 (right column) emission scenarios, for the location 45.5°N, 12.5° E (Venice).

was configured with 5 chains, running a maximum of 5000 iterations and an initial burn-in period of 2500 iterations. As shown in table 5.5, the required CPU times remain computationally tractable. The outcomes of this spatial analysis are visualized in the maps presented in figs. 5.11 to 5.14. Although the projection time span is limited to 65 years, distinct differences between the two emission scenarios are clearly discernible.

Focusing specifically on the relative percentage changes (figs. 5.13 and 5.14), several regional patterns emerge. In Northern Italy, the overall trend remains largely consistent across most grid points, displaying a general decreasing trend. However, the Ligurian region presents an intensification of ΔERL_{50} from the RCP 4.5 to the RCP 8.5 scenario, while locations in the Ligurian Sea exhibit the opposite behaviour, shifting towards a decreasing trend. Conversely, points in the Tyrrhenian Sea show a decreasing trend under RCP 4.5 but shift to an increasing trend under RCP 8.5. Along the Adriatic coast, the overall behaviour is largely conserved between the two scenarios, although with a stronger decreasing magnitude at several localized points under RCP 8.5. In Central Italy, locations already experiencing a decrease show a more pronounced negative trend under RCP 8.5, though the overall regional pattern remains stable. In Southern Italy, the trends frequently invert between scenarios. For instance, in Sicily, an increasing trend under RCP 4.5 reverses to a decreasing trend under RCP 8.5. Finally, for the 4 grid points in Sardinia, the trend is completely reversed, specifically shifting from a decrease under RCP 4.5 to an increase under RCP 8.5.

Because the computational domain includes neighbouring cross-border regions, certain dramatic shifts become immediately apparent. A striking example is located in Switzerland, where a positive trend of +18% under RCP 4.5 completely reverses to −15% under RCP 8.5. An analogous behavioural reversal is observed at several points along the French border. This dramatic shift can likely be attributed to the complex orography of the Alps, which poses inherent challenges and introduces higher uncertainties in high-resolution regional climate modelling.

In conclusion, the RCP 8.5 scenario seems to produce more intense spatial gradients—reflected visually by deeper map colours—indicating more severe increasing or decreasing trends compared to the RCP 4.5 scenario. Ultimately, this analysis demonstrates that climate change does not induce a uniform, unidirectional shift in extreme wind speeds across the domain. Rather, the future trajectory—whether intensifying or attenuating—is highly dependent on localized geographical and environmental boundary conditions.

Table 5.5: Summary of CPU execution times for the non-stationary extreme value analysis under RCP4.5 and RCP8.5 scenarios across a 156-point grid over Italy.

Metric	**RCP4.5**	**RCP8.5**
Average CPU time (s)	97.01	96.91
Min CPU time (s)	41.30	91.38
Max CPU time (s)	120.17	102.45
Total time (s)	15134.21	15118.32

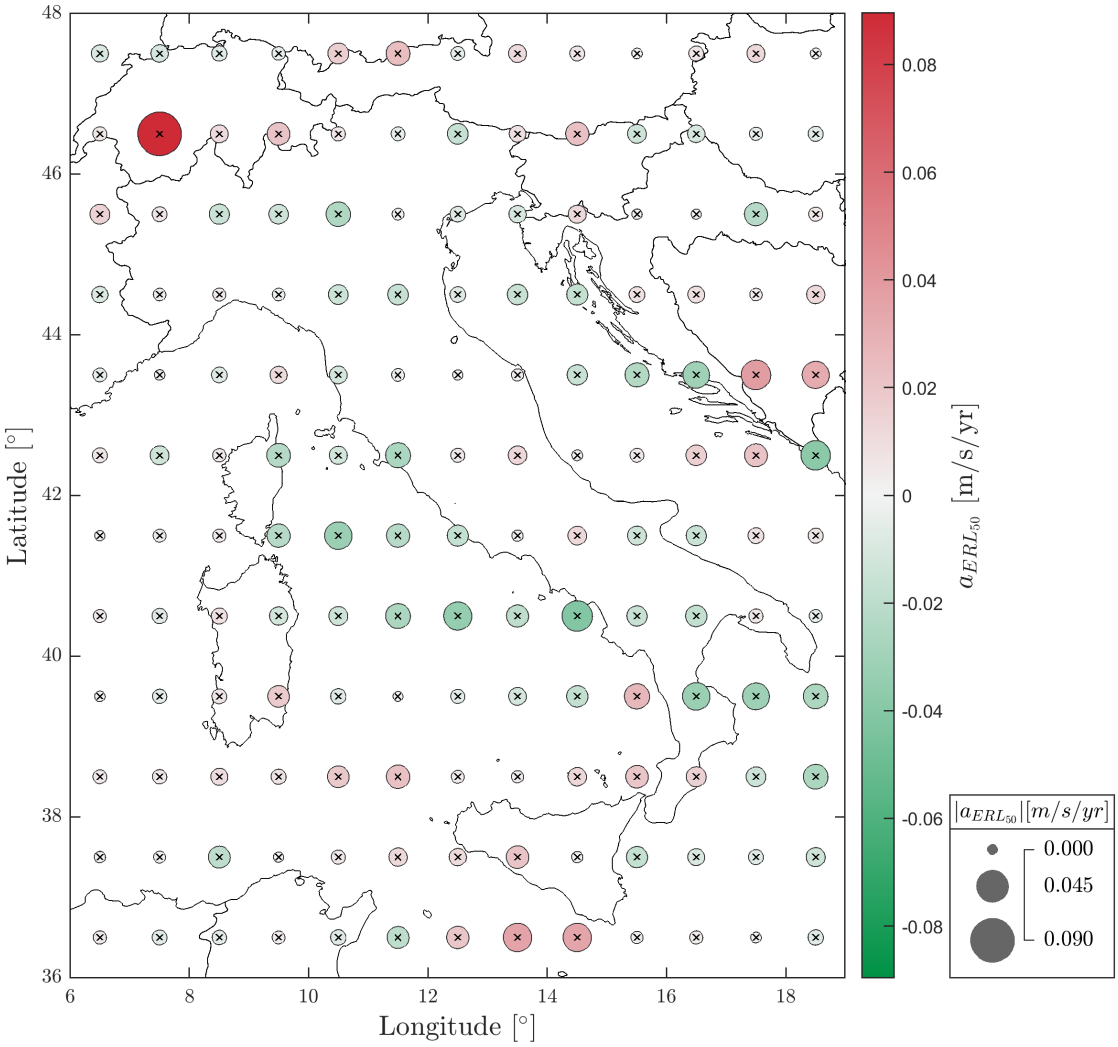

Figure 5.11: Bubble map of the 50-year effective return level (ERL_{50}) trend under the RCP 4.5 scenario.

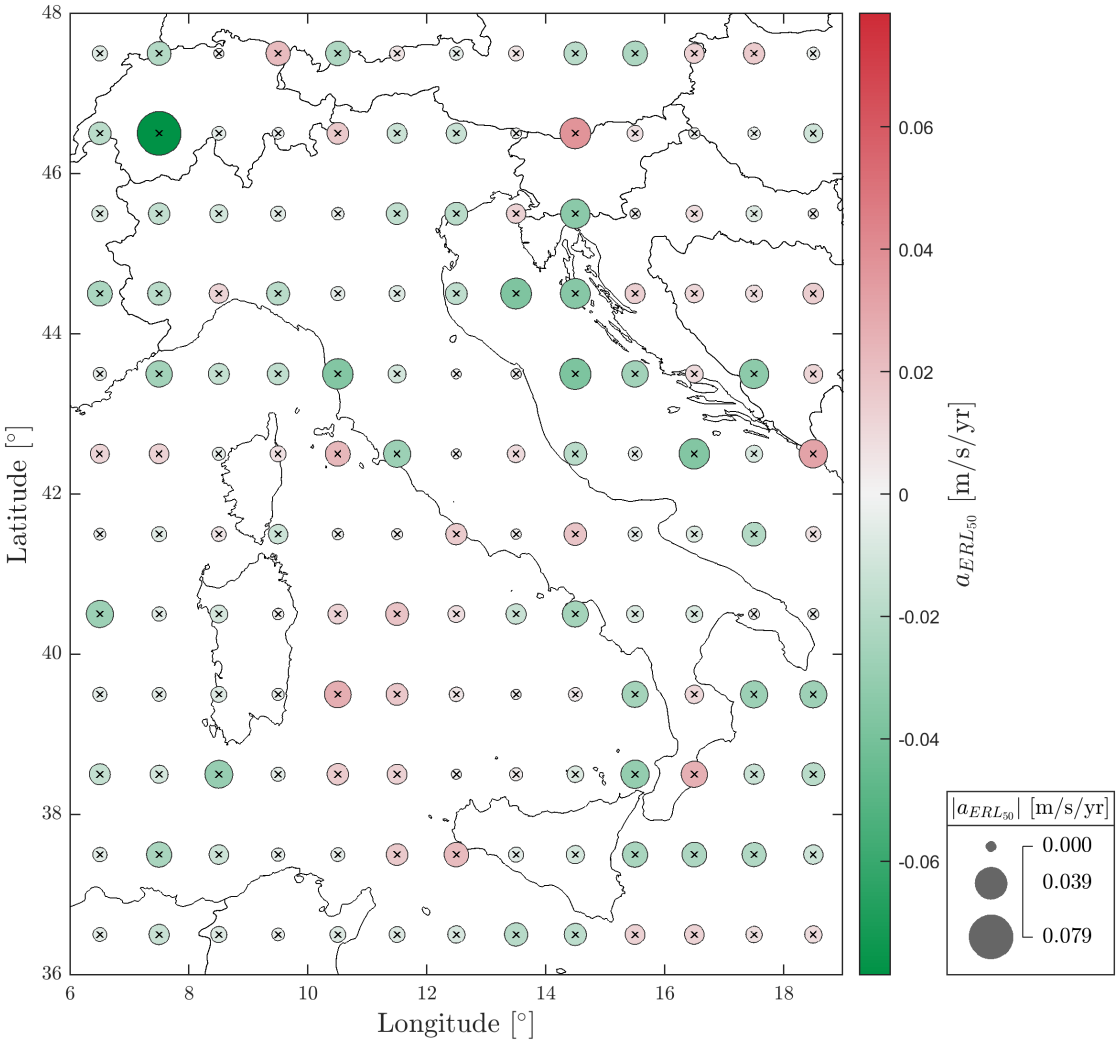

Figure 5.12: Bubble map of the 50-year effective return level (ERL_{50}) trend under the RCP 8.5 scenario.

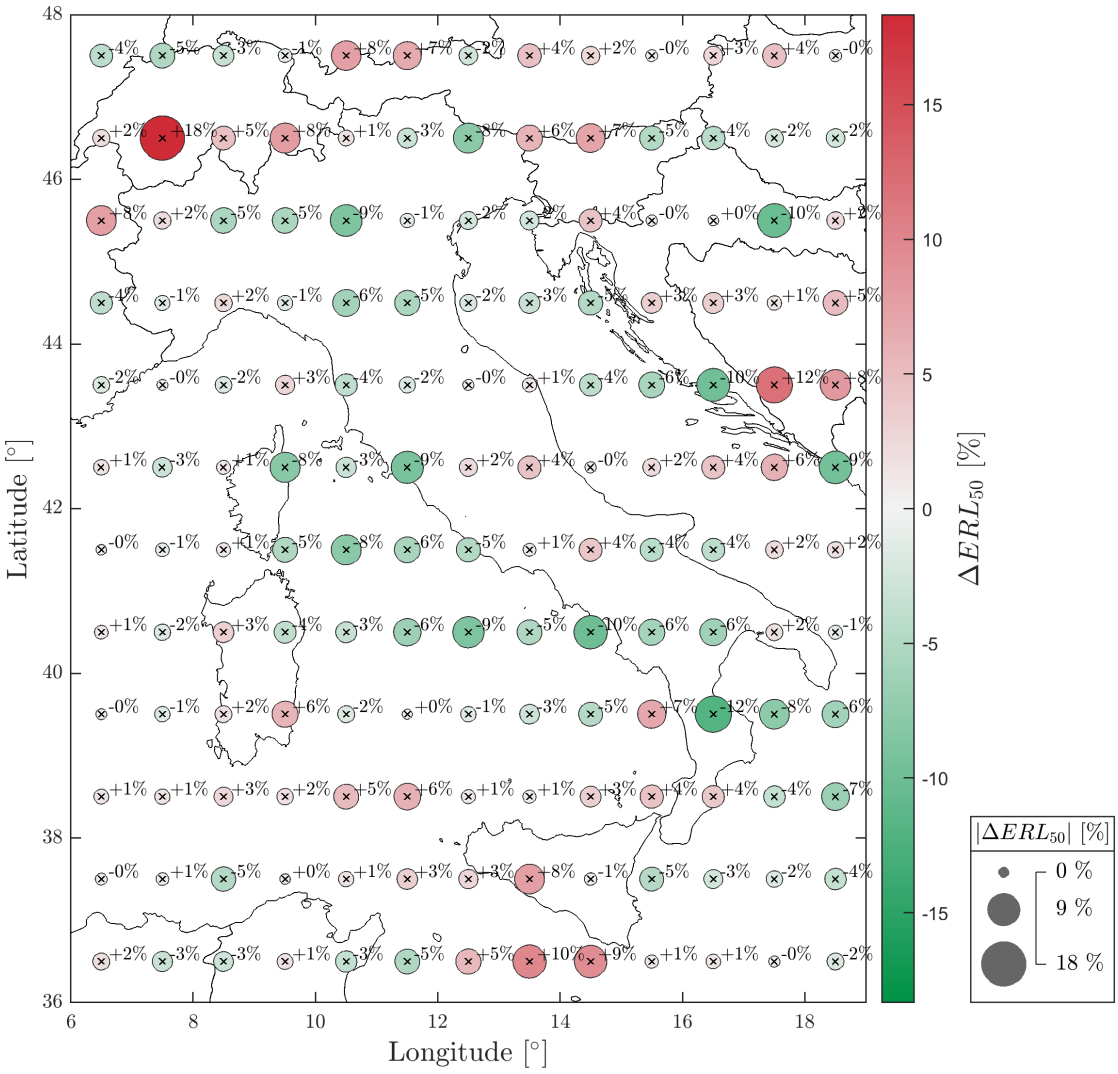

Figure 5.13: Bubble map of the 50-year effective return level (ERL_{50}) percentage trend under the RCP 4.5 scenario.

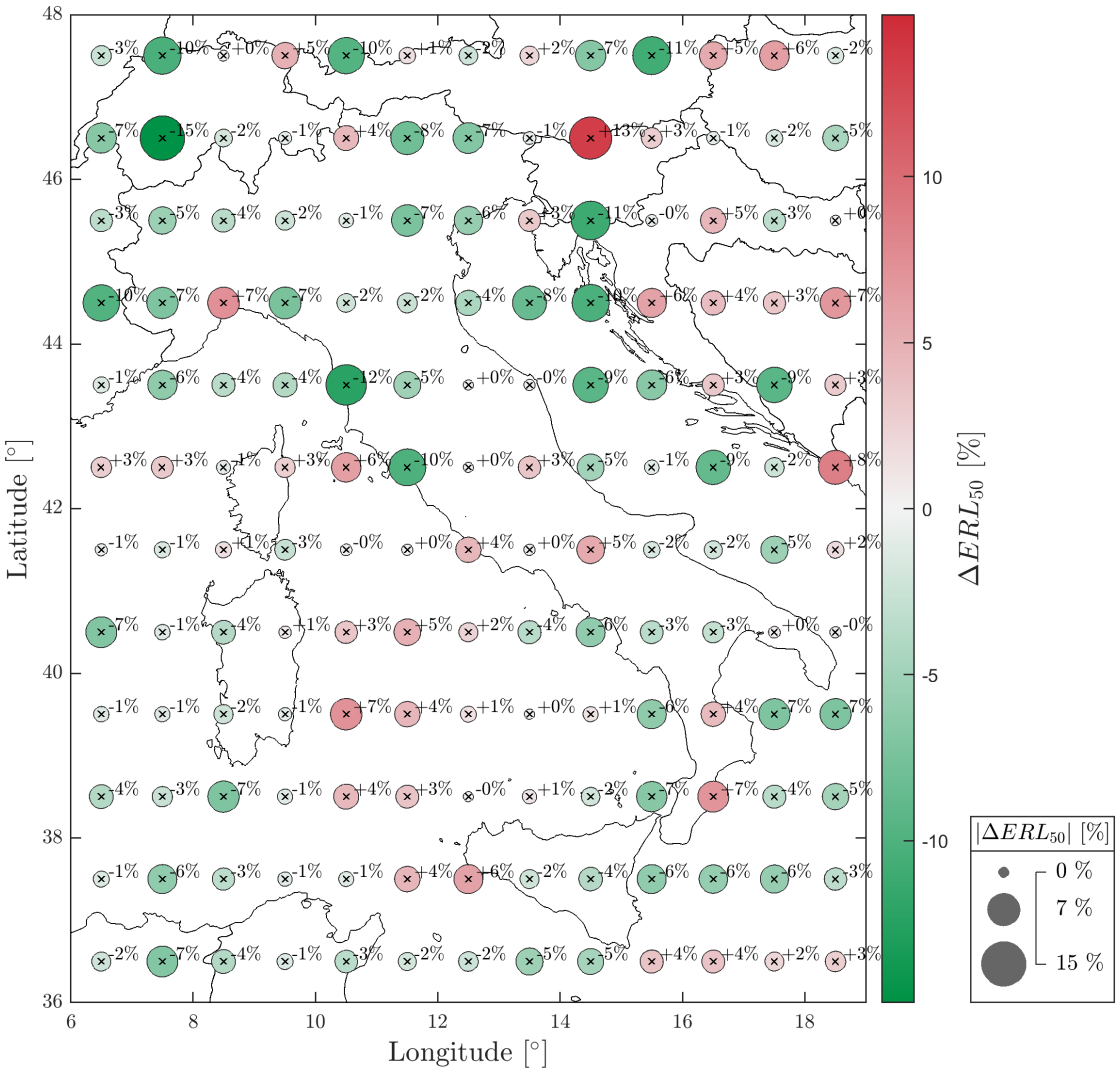

Figure 5.14: Bubble map of the 50-year effective return level (ERL_{50}) percentage trend under the RCP 8.5 scenario.

Chapter 6

Conclusions and Future Perspectives

One of the fundamental assumptions in traditional wind engineering is the stationarity of extreme wind speeds. This hypothesis is increasingly challenged by anthropic environmental changes. The primary goal of this MSc Thesis is to apply a robust statistical framework to investigate and quantify non-stationarities, providing a novel perspective on extreme wind speed metrics.

First, a non-stationary Bayesian inference model is applied to synthetic time series. This preliminary phase confirms the model's ability to capture and predict extreme events under controlled conditions, yielding statistically sound results.

Second, the impact of increasing urbanization and the corresponding change in surface roughness is investigated by applying the verified model to historical anemometric measurements. The spatial analysis of 14 selected stations across the Italian Peninsula reveals a prevailing decreasing trend for the 50-year Effective Return Level (ERL_{50}). This widespread attenuation overall agrees with the general increase in aerodynamic roughness induced by progressive land consumption, although urbanization cannot be identified as the sole driving factor.

Subsequently, the framework is applied to high-resolution future climate projections over the Italian Peninsula, considering both RCP 4.5 and RCP 8.5 emission scenarios. The resulting maps of future ERL_{50} trends highlight the marked spatial non-uniformity of extreme wind dynamics. Crucially, the analysis demonstrates that the most severe emission scenario (RCP 8.5) does not necessarily lead to a generalized increase in extreme wind speeds. Rather, it seems to induce sharper spatial gradients and localized trend reversals, strongly modulated by the region's complex orography.

These findings have significant implications for structural design. The traditional stationary hypothesis risks either underestimating or overestimating design wind

speeds, depending on the local environmental trajectory. Coversely, the wind maps obtained in the present MSc Thesis potentially allow to tune the design wind speed based on the design site location ensuring structural safety and reliability over the structure and infrastructure's design life.

Despite these contributions, this study acknowledges certain methodological limitations. Specifically, non-stationarity was modelled assuming strictly linear temporal trends for the distribution parameters. While effective, this simplification may sometimes fail to capture highly non-linear dynamics in real-world time series. Future research perspectives should investigate whether incorporating more complex trend models can improve predictive accuracy without falling into the trap of over-parametrization and overfitting. Another limitation lies in the reliance on a single high-resolution climate projection model. Future research would greatly benefit from an ensemble modelling approach to systematically reduce intrinsic climate uncertainties. Furthermore, the available projections extend only up to 2070; while representing the current state-of-the-art for this spatial resolution, this time horizon poses a constraint for asymptotic analyses of very long return periods. Finally, projection datasets that dynamically couple global emission scenarios with evolving models of expected urban surface roughness are not yet available. Investigating the physical interaction between the effects of climate change on wind speed and the progressive land urbanization, and its statistical outcome, remains a critical open challenge for the future of wind engineering.

Bibliography

[1] Stuart Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. Springer London, 2001. DOI: 10.1007/978-1-4471-3675-0 (cit. on pp. 4, 5, 7, 9, 26).

[2] Jan Beirlant, Yuri Goegebeur, Johan Segers, and Jozef Teugels. *Statistics of Extremes: Theory and Applications*. Chichester, UK: John Wiley & Sons, 2004. ISBN: 978-0-471-97647-9 (cit. on pp. 4, 6).

[3] Laurens de Haan and Ana Ferreira. *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering. Springer, 2006 (cit. on p. 7).

[4] Paula Reis, Luísa Canto e Castro, Sandra Dias, and M. Ivette Gomes. «Penultimate approximations in statistics of extremes and reliability of large coherent systems». In: (2013) (cit. on p. 10).

[5] Rym Worms. «Penultimate approximation for the distribution of the excesses». In: *ESAIM: Probability and Statistics* 6 (2002), pp. 21–31 (cit. on p. 11).

[6] Santiago Beguería, Marta Angulo-Martínez, Sergio M. Vicente-Serrano, J. Ignacio López-Moreno, and Ahmed El-Kenawy. «Assessing trends in extreme precipitation events intensity and magnitude using non-stationary peaks-over-threshold analysis: a case study in northeast Spain from 1930 to 2006». In: *International Journal of Climatology* 31 (2011), pp. 2102–2114. DOI: 10.1002/joc.2218 (cit. on pp. 13, 17).

[7] Soheil Radfar, Panagiota Galiatsatou, and Thomas Wahl. «Application of nonstationary extreme value analysis in the coastal environment - A systematic literature review». In: *Weather and Climate Extremes* 41 (2023), p. 100575. DOI: 10.1016/j.wace.2023.100575 (cit. on pp. 13, 16).

[8] Linyin Cheng, Amir AghaKouchak, Eric Gilleland, and Richard W Katz. «Non-stationary extreme value analysis in a changing climate». In: *Climatic change* 127.2 (2014), pp. 353–369 (cit. on pp. 14, 17).

[9] Elisa Ragno, Amir AghaKouchak, Linyin Cheng, and Mojtaba Sadegh. «A generalized framework for process-informed nonstationary extreme value analysis». In: *Advances in Water Resources* 130 (2019), pp. 270–282. DOI: 10.1016/j.advwatres.2019.06.007 (cit. on pp. 15, 17, 18, 20).

[10] Thomas Opitz, Raphaël Huser, Haakon Bakka, and Håvard Rue. «INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles». In: *Extremes* 21 (2018), pp. 441–462. DOI: 10.1007/s10687-018-0324-x (cit. on pp. 15, 17).

[11] Riccardo Bonanno, Francesca Viterbo, and Riva Giuseppe Maurizio. «Climate change impacts on wind power generation for the Italian peninsula». In: *Regional Environmental Change* 23.15 (2023), pp. 1–11. DOI: 10.1007/s10113-022-02007-w (cit. on p. 15).

[12] Mario Raffa, Marianna Adinolfi, Alfredo Reder, Gian Franco Marras, Marco Mancini, Gabriella Scipione, Monia Santini, and Paola Mercogliano. «Very High Resolution Projections over Italy under different CMIP5 IPCC scenarios». In: *Scientific Data* (2023). DOI: 10.1038/s41597-023-02144-9 (cit. on pp. 16, 44).

[13] Lorenzo Raffaele, Luca Bruno, and Elisabetta Colucci. «Reanalysis-based mesoscale wind maps for the design of structures and infrastructures with an application to Italy». In: *Journal of Wind Engineering & Industrial Aerodynamics* 253 (2024). DOI: 10.1016/j.jweia.2024.105844 (cit. on pp. 16, 45).

[14] H.M. Mo, H.P. Hong, and F. Fan. «Estimating the extreme wind speed for regions in China using surface wind observations and reanalysis data». In: *Journal of Wind Engineering and Industrial Aerodynamics* 143 (2015), pp. 19–33. DOI: 10.1016/j.jweia.2015.04.005 (cit. on pp. 16, 45).

[15] Shihan Yang, Xiaodong Ji, Yifan Tao, Jingjing Jia, Zeyu Zhang, and Li Jiang. «Study on Non-Stationary Wind Speed Models and Wind Load Design Parameters Based on Data from the Beijing 325 m Meteorological Tower, 1991-2020». In: *Applied Sciences* 14.10 (2024). DOI: 10.3390/app14103979 (cit. on pp. 16, 45).

[16] Mojtaba Sadegh, Elisa Ragno, and Amir AghaKouchak. «Multivariate Copula Analysis Toolbox (MvCAT): Describing dependence and underlying uncertainty using a Bayesian framework». In: *Water Resources Research* 53.6 (2017), pp. 5166–5183. DOI: 10.1002/2016WR020242 (cit. on p. 23).

[17] National Centers for Environmental Informations, Integrated Surface Dataset (Global), last accessed: October 2023. URL: https://www.ncei.noaa.gov/access/search/data-search/global-hourly (cit. on p. 44).