Propagation of wind direction uncertainty into wake modelling

Li-Hsuan Chang

Delft University of Technology

TU Delft

Propagation of wind direction uncertainty into wake modelling

by

Li-Hsuan Chang

Supervisor(s):
Simon Watson — TU Delft
Mark C. Kelly — DTU Wind & Energy System
Gibson Kersting — RWE Clean Energy
Alessandro Sebastiani — RWE Renewables, Denmark

Faculty: Faculty of Aerospace Engineering
Education: Master of Science

TU Delft

Abstract

Wind direction uncertainty is often overlooked in wind resource assessment, despite its potential to cause significant deviations in annual energy production (AEP). This study quantifies directional uncertainty arising solely from wind data, decomposed into three components: measurement uncertainty, long-term wind rose difference (LTWRD), and interannual variability (IAV). Two offshore sites with contrasting layouts and inflow regimes were analysed using statistical and resampling-based approaches.

The results show total directional uncertainties of 0.71% and 4.39% of AEP, with LTWRD emerging as the dominant contributor at the more directionally sensitive site. No consistent pattern in the composition of directional uncertainty was observed across sites, highlighting the need for site-specific assessments.

Given that typical total AEP uncertainties for wind projects average 8–10%, the findings indicate that directional uncertainty can represent a substantial fraction of the total, especially at sites with strong wake sensitivity and uni-directional inflow. Therefore, wind direction should be explicitly included in uncertainty analyses, and long-term correction procedures should address wind direction as well as wind speed.

Acknowledgements

I would like to express my deepest gratitude to all my supervisors, Mark, Gibson, Alessandro, and Simon, for their valuable feedback, patience, and support throughout this journey. Each of you contributed in different ways, and I truly appreciate your involvement.

In particular, I would like to extend my heartfelt thanks to Gibson and Mark for their unwavering commitment and additional guidance. Gibson, thank you for dedicating so much time to discussing the challenges I faced and for always doing your best to support me, whether by helping me access computational resources, providing specific data, or offering thoughtful feedback. Your encouragement and willingness to support my ideas, even during your busiest moments, truly meant a lot. Mark, thank you for your timely responses and for consistently offering constructive, insightful feedback. Your expertise has significantly shaped the direction and quality of this thesis.

Alessandro, thank you for being so approachable and responsive, and for being the only supervisor I could regularly meet in person. I greatly appreciated our conversations, not only about my thesis, but also about the industry and ongoing projects. I would also like to thank Simon, who, despite his busy schedule, always made time to attend meetings and responded promptly to my requests, even during his holidays.

This endeavour would not have been as smooth (or even possible) without the support from RWE, which provided the data and resources essential for my research. I am sincerely grateful for RWE's contribution and support.

I am also thankful to my friends and office mates, especially Iris, EJ, Yao, Spyros, and Teng. Iris, thank you for your late-night support and all the good food. EJ, I appreciate your help and guidance since I arrived in Denmark in 2023, it made settling in much easier. Yao, thank you for the emotional support, our weekly deep talks, and of course, the great meals. Spyros, I value your feedback and the hard work we shared at the RWE office. And most importantly, Teng. I am truly grateful to have shared this journey with you at TU Delft. From our application process to now, at the end of this master's program, you've always been by my side. Thank you for always answering my calls, lifting my spirits, and keeping me motivated with your constant support and friendship.

Finally, I would like to thank my family, especially my mother, brother and my dog. Thank you for your unwavering belief in me and for supporting my choices without hesitation. Your love and encouragement gave me the confidence to fully dedicate myself to my studies. I am forever grateful for your unconditional support.

There are many others who helped me throughout this journey, far too many to name here. To all of you: thank you for your encouragement, friendship, and the countless memories. This journey would not have been the same without you.

Contents

1 Introduction

1.1 Background

Wind energy plays a central role in the global transition to renewable energy. Annual wind power production continues to grow, with Europe alone expected to install an additional 187 GW of capacity between 2025 and 2030 to meet the EU's 42.5% renewable energy target. This expansion will bring total installed capacity in the EU to approximately 425 GW by 2030 [1].

Such rapid growth places an increasing emphasis on the accuracy and consistency of annual energy production (AEP) estimates. The reliability of AEP calculations is fundamental to project financing, operational planning, and energy policy development. High AEP uncertainty can negatively impact debt-service coverage ratios, raise financing costs, and in extreme cases threaten project viability [2].

At the same time, modern wind farms are growing in size, and turbines are increasingly exposed to harmful wake interactions. Accurate modelling of inflow conditions, particularly wind direction, is therefore essential, since even small directional shifts can significantly alter wake propagation patterns and total farm output.

Uncertainty in AEP estimation is generally grouped into five categories [3]: (1) Wind data, (2) Wind model, (3) Power conversion, (4) Bias, and (5) Losses. The thesis focuses exclusively on wind data uncertainty, as it forms the foundation for all subsequent modelling steps.

1.2 Wind Data Uncertainty in AEP and the Role of Wind Direction

Most research on wind data uncertainty has focused on wind speed because of its direct link to turbine power production via the power curve. Numerous deterministic and probabilistic methods have been developed to propagate wind speed measurement error into AEP uncertainty [4, 5]. Even with high-quality instrumentation, wind speed measurement uncertainty alone is rarely below 2%, and when propagated through the cubic power curve relationship, this translates to AEP uncertainties of 3–6% [3]. These studies have examined not only the magnitude of wind speed errors but also their interactions with other variables such as turbulence intensity (TI), hub-height selection, and correlations between uncertainty components.

By contrast, wind direction uncertainty remains comparatively underexplored and is often neglected in AEP assessments [6, 7]. While wind speed determines the magnitude of available energy, wind direction strongly governs wake interactions and spatial variability in wind farm performance. Neglecting directional uncertainty can therefore lead to systematic underestimation of wake losses and total AEP uncertainty.

Djokić et al. [8] demonstrated that explicitly incorporating wind direction uncertainty improved predictive performance in models that otherwise assumed perfect directional data. Brandis et al. [9], using mesoscale modelling (NEWA dataset, 1989–2018 [10]) for the German Bight, reported mesoscale wind direction changes exceeding 11° in 50% of cases. These variations, relevant over spatial scales of tens to hundreds of kilometres, can markedly influence wake model performance.

Porté-Agel et al. [11] quantified this sensitivity by simulating the Horns Rev offshore wind farm under multiple inflow angles with a large eddy simulation (LES) model. They found that a 10° shift from the worst-case (full wake) inflow increased total farm power output by 43%. While striking, this result is a deterministic sensitivity analysis and does not constitute a full uncertainty quantification, as it does not incorporate the probability distribution of inflow directions.

Gaumond et al. [12] partly addressed this gap by estimating wind direction uncertainty and its influence on model–data agreement for Horns Rev. However, their work focused on model uncertainty related to direction, rather than wind data uncertainty itself. Other direction-related studies have focused on yaw misalignment or sensor-comparison errors [13, 14, 15], for example, comparing mast-mounted wind vanes with nacelle yaw sensors. Overall, there is no widely adopted framework for incorporating wind direction uncertainty into AEP uncertainty budgets within the wind data category.

1.3 Uncertainty Quantification of Wind Rose Characteristics

Quantifying uncertainty in wind direction (and, more generally, in the wind-rose frequency distribution) is less standardised than for wind speed. Two main methodological families can be identified:

1. Circular statistical metrics
   Because wind direction is a circular variable bounded between 0° and 360°, conventional statistical measures such as linear standard deviation are unsuitable. Several alternative formulations exist [16, 17], with the Yamartino method [18] shown to perform well in many cases [19, 20]. However, these methods typically assume unimodal or symmetric wrapped-normal distributions [21] and may fail to represent multi-modal or highly skewed distributions, which are common in offshore sites.

2. Frequency-distribution comparisons
   This approach quantifies differences in wind direction frequency distributions between two datasets (e.g., short-term mast measurements vs. long-term reanalysis) by comparing their sector-wise frequencies. A common metric is the wind rose deviation (WRD) [22], defined as the sum of absolute differences in sector frequencies, similar to metrics applied in wind rose bias-correction methods [23, 24]. These methods preserve the difference in distributional shape but are sensitive to the choice of sector bin size and to sample size limitations. As noted by Kjeller Vindteknikk [22], the representativeness of a long-term directional variation cannot be fully characterised by a single statistic such as a correlation coefficient or the WRD alone.

Despite these limitations, both approaches are widely used in practice, yet no industry consensus exists on which is most appropriate for integrating wind rose uncertainty into AEP calculations.

1.4 Research Goal and Thesis Outline

The aim to this thesis is to quantify wind direction uncertainty from wind data and assess its contribution to AEP uncertainty. The study addresses two key research questions:

- How does uncertainty in wind direction propagate into AEP through wake modelling, and to what extent does measurement bias contribute to the overall AEP uncertainty?
- What is the most appropriate way to quantify uncertainty in wind direction: using circular statistical metrics or frequency distribution differences?

The analysis is carried out for two offshore sites with contrasting inflow regimes and layouts, enabling assessment of site-specific behaviour. The findings aim to guide wind farm developers in selecting measurement strategies, highlight the importance of long-term correction methods, and inform whether directional uncertainty should be systematically integrated into total AEP uncertainty budgets.

The remainder of this thesis is structured as follows:

- **Chapter 2** – Theoretical background
- **Chapter 3** – Methodologies for quantifying directional uncertainty
- **Chapter 4** – Case study applications and results
- **Chapter 5** – Discussion, additional uncertainty considerations, and recommendations
- **Chapter 6** – Conclusions and future work

2 Basis and Framework

2.1 Site Overview and Datasets

This study investigates wind conditions and uncertainty quantification at two offshore wind project sites. For confidentiality reasons, the sites are anonymised as *Site A* and *Site B*, with sufficient descriptive context retained for the purposes of this thesis. *Site A* is an operational offshore wind farm located in the Baltic Sea between Denmark and Germany. In contrast, *Site B* is a pre-construction offshore site situated off the western coast of the United States and is currently in the planning phase.

2.1.1 Site Characteristics and Measured Dataset

Both sites were selected based on the availability of site-specific measured data and long-term reanalysis datasets. For *Site A*, 35 years of hourly reanalysis data and 5.5 years of 10-minute resolution met mast measurements are available. *Site B* is supported by 20 years of hourly reanalysis data and 1.5 years of 10-minute resolution met mast data. All measured datasets are sourced from the RWE database and are subject to confidentiality agreements. The availability of both long-term reanalysis data and higher-frequency measurements provides a basis for investigating directional uncertainty in wind conditions, as shown in Figure 2.1.

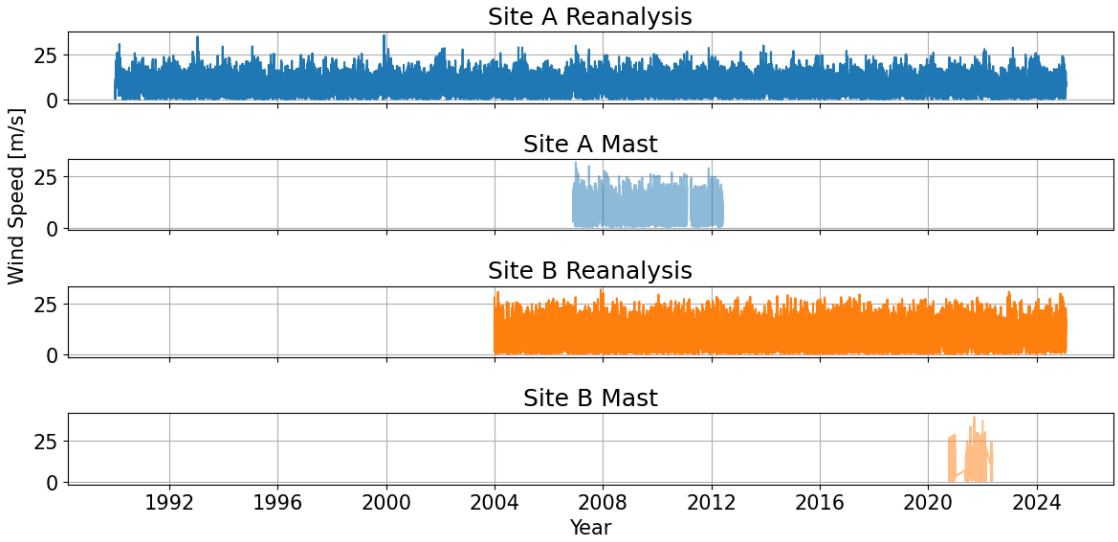

Figure 2.1: The availability of wind data from measurement and reanalysis of *Site A* and *Site B*.

It is worth noting that while the measured dataset at *Site B* nominally spans 1.5 years, a substantial data gap between late December 2020 and the end of May 2021 reduces the effective duration of usable data to approximately 1.1 years. Relying on such a limited measurement period for AEP calculation and uncertainty quantification introduces several challenges for *Site B*. First, the reduced temporal coverage compromises representativeness, making the results more sensitive to seasonal imbalances and abnormal events. Second, the smaller dataset limits the diversity of resampled sequences in bootstrapping, potentially leading to an underestimation of uncertainty. Finally, the shortened concurrent

period weakens the statistical correspondence between measured and reanalysis data, which may impair the accuracy of future long-term corrections.

2.1.2 Reanalysis Dataset

This study uses the ERA5 reanalysis dataset [25], which provides a physically consistent reconstruction of historical weather conditions by assimilating a wide range of observations into a global forecast model. ERA5 delivers hourly atmospheric data on a global grid at $0.28^\circ \times 0.28^\circ$ resolution (approximately 31 km) [26], with wind variables available at multiple vertical levels, including 100 m above ground level, making it suitable for long-term climatological assessments at turbine hub heights.

ERA5 is widely used in wind resource assessment, yet previous studies report substantial underestimation of wind power when using ERA5 wind speeds directly, up to 20% according to Wilczak et al.[27]. Comparisons with other reanalyses (e.g., MERRA-2) indicate that ERA5 provides more accurate large-scale patterns[28], especially offshore, where the absence of orographic effects improves correlation with site measurements [29]. However, like other reanalyses, ERA5 cannot fully resolve local wind speed or wind direction variations. While no specific studies have quantified ERA5's wind direction accuracy at local scale, similar resolution constraints apply.

Moreover, ERA5's hourly resolution does not capture high-frequency fluctuations such as rapid direction shifts or wind veer events. However, since this study focuses on wind direction uncertainty at longer time scales (e.g., interannual variability), such high-frequency limitations are not expected to have a major influence.

For consistency, only measured wind direction datasets with documented quality control and at least one year of continuous coverage were compared against ERA5. In this study, wind direction from ERA5 is used only after spatial downscaling to site level; no bias correction or directional adjustment has been applied. Unless otherwise stated, all references to reanalysis data in the following sections refer to ERA5 after this downscaling step.

2.1.3 Wind Direction Characteristics

A second motivation for selecting these sites is their contrasting wind direction characteristics. *Site A* locates in the Baltic Sea, leading to more uniformly distributed wind directions with seasonal bi-directionality. In contrast, the *Site B* site is dominated by a persistent, strongly uni-directional inflow. These differences in wind rose patterns allow for a comparative assessment of how directional variability in fluence uncertainty quantification. Figure 2.2 illustrates the wind rose distributions derived from met mast measurements at each site, corresponding to turbine hub heights of 100 meters (*Site A*) and 170 meters (*Site B*), respectively.

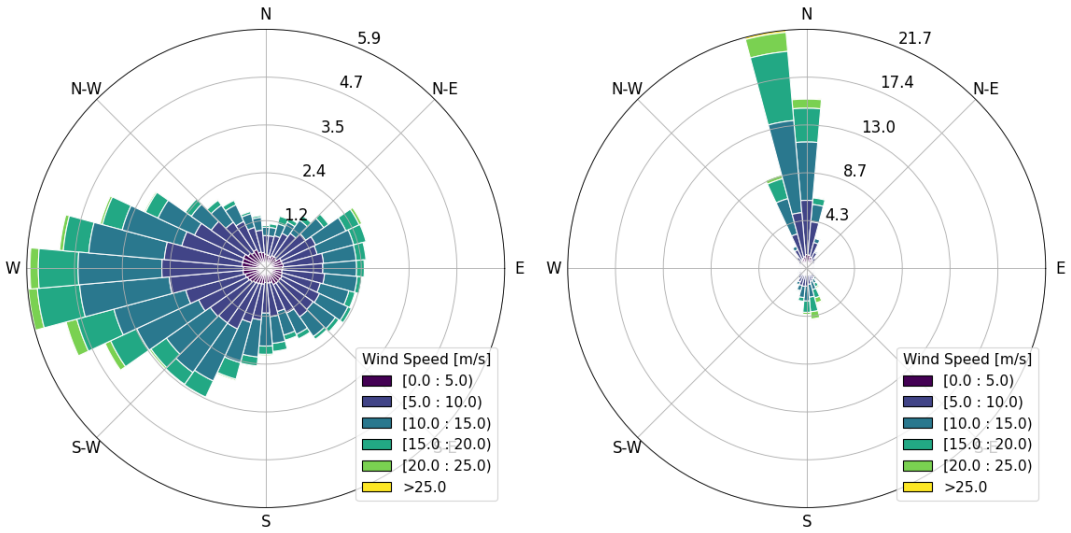

Figure 2.2: Wind roses of 10-minute measured wind speed and direction at hub height for *Site A* (100 m, left) and *Site B* (170 m, right).

2.1.4 Wind Farm Layouts

The wind farm configurations of both sites are shown in Figure 2.3. The turbine locations have been normalised by rotor diameter to preserve confidentiality and show turbine spacing. While terrain information is not considered for offshore wind farms in this study, and due to confidentiality, terrain features are not visualised. This simplification is justified as the layout simulations in Pywake are conducted under the assumption of flat terrain.

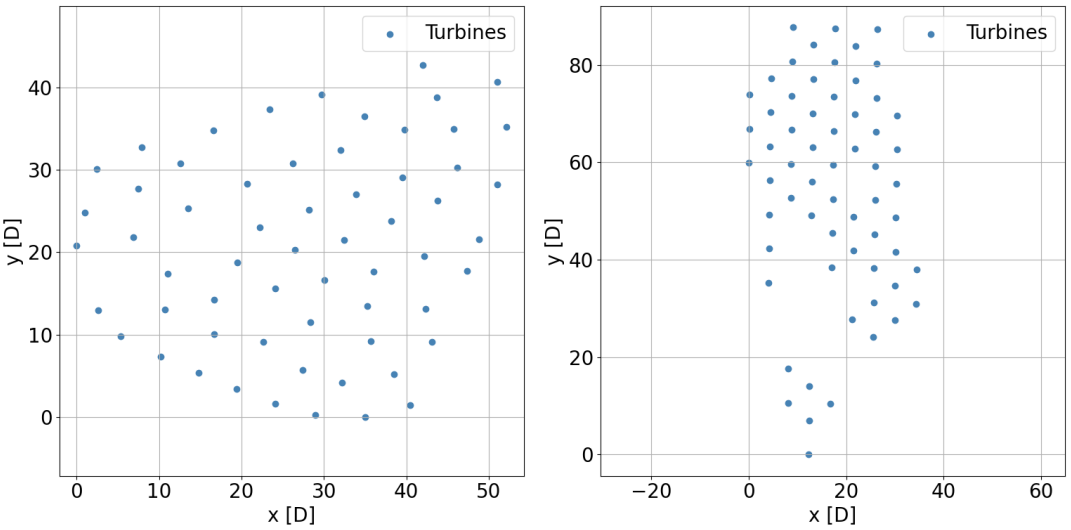

Figure 2.3: Wind farm layout at *Site A* (left) and *Site B* (right), normalised by their rotor diameter D (154 meter and 292 meter, respectively).

As illustrated in Figure 2.3, the two wind farm layouts exhibit markedly different spatial configurations when normalised by rotor diameter. *Site A* displays a relatively symmetric grid-like arrangement, particularly symmetric about the y-axis. This symmetry, when coupled with the broader and more uniformly distributed wind rose observed at *Site A*, suggests that the site is likely to be less sensitive to small directional shifts in inflow. In contrast, *Site B* features a denser and more staggered turbine arrangement with clear asymmetry

and tighter inter-turbine spacing. The dominant inflow direction at *Site B* aligns closely with the long axis of the layout, resulting in a high probability of wake interactions.

Notably, inter-row spacing in the prevailing inflow direction at *Site B* is approximately 5–7 rotor diameters (D), within the lower end of the typical recommended range. For reference, the Horns Rev wind farm off the coast of Denmark averages about $7D$ spacing [30], while downstream spacing is often optimised between $6D$ and $10D$ depending on site-specific constraints and control strategies [31]. Spacing models, developed by Meneveau et al., suggest that larger spacings (around $15D$) can improve cost efficiency [32]. The compact configuration at *Site B* amplifies directional sensitivity, small inflow shifts can substantially alter the number of turbines exposed to upstream wakes, with corresponding impacts on energy yield.

By comparison, *Site A* has wider turbine spacing of roughly 8–$10D$ along its dominant westerly inflow, reducing wake overlap and directional sensitivity.

Together, these two sites provide complementary test cases for evaluating directional uncertainty in offshore wind resource assessment. *Site A* represents a well-documented, operational site with long-term measurements and a less waked layout, serving as a baseline for conventional offshore conditions. *Site B*, by contrast, presents an extreme case of strong uni-directional inflow, compact turbine spacing, and limited measurement duration. This pairing enables assessment of directional uncertainty across both typical offshore conditions and more challenging, directionally sensitive configurations.

2.2 Sources of Wind Direction Uncertainty

Reliable wind resource estimation requires a comprehensive understanding of its associated uncertainties. In the context of wind direction, these uncertainties stem from various sources, including measurement bias, interannual variability (IAV), and long-term wind rose deviation (LTWRD), which is the discrepancy between directional frequency distributions in short-term measurements versus long-term reference datasets.

According to the IEC 61400-1 framework and as summarised by Lee and Fields [6], energy production uncertainty during resource assessment stage can be categorised into five main components: measurement, historical wind resource, spatial extrapolation, project evaluation period variability, and plant performance. This section focuses on the uncertainty component and details the selection and justification of the uncertainty sources incorporated into this study.

2.2.1 Measurement Uncertainty

A key objective of this study is to understand whether wind direction measurement uncertainty inherent to standard met mast setups significantly contributes to the overall directional uncertainty in wind resource assessments. One of the research questions addressed is whether there is a justified need to upgrade instrumentation for higher precision, or if the uncertainty associated with lower-accuracy equipment is already sufficiently small relative to other sources.

While high-accuracy sensors and calibration protocols can be used to improve measurement precision and accuracy, systematic errors may still arise from practical factors such as mast alignment, sensor orientation, and flow distortion around the mast structure. These errors typically do not vary randomly over time but are instead persistent and site-specific, making it important to quantify their impact carefully. Including this source in the total directional uncertainty budget is therefore essential for understanding the extent to which measurement uncertainty could bias energy yield estimates, especially when

combined with or compared to other sources of uncertainty. All detailed calculations, assumptions, and individual contributions to the measurement uncertainty are presented in Section 3.1.

2.2.2 Historical Wind Resource

While long-term reanalysis data are essential for capturing climatological trends in wind resource assessment, their treatment differs depending on the variable of interest. For wind speed, long-term correction techniques such as Virtual Met Mast or Measure-Correlate-Predict are commonly applied to bridge the gap between short-term measurements and long-term reference data [33]. However, for wind direction, such corrections are not typically performed. Instead, wind direction in reanalysis datasets is implicitly refined during the downscaling process. The downscaling models improve spatial representation but do not explicitly calibrate wind direction against local measurements.

Although reanalysis datasets are widely used for directional analysis, they introduce additional uncertainty due to spatial mismatches, model limitations, and lack of site-specific validation [34, 35, 36]. The resulting deviation between the directional distributions (wind roses) of short-term measurements and long-term reanalysis is referred to this study as historical resource assessment uncertainty.

To characterise this uncertainty in wind direction, it is decomposed into two independent components: **Inter-Annual Variation (IAV)** and **Long-Term Wind Rose Difference (LTWRD)** [37]. IAV accounts for natural year-to-year shifts in the wind rose distribution, while LTWRD reflects systematic differences between the short-term measured and long-term reanalysis directional data. To avoid double-counting of annual variation, seasonal-trend decomposition is applied to datasets to isolate and remove long-term patterns prior to LTWRD uncertainty estimation.

Unlike wind speed, wind direction is a circular variable, making conventional metrics like Pearson correlation and linear standard deviation inappropriate. Moreover, wind direction affects AEP indirectly through altering the distribution of wind speeds and their directional frequency. For this reason, IAV is not expressed in degrees, since there is no universally transferable scalar metric for circular, often multimodal, wind direction distributions, as discussed in Section 3.3. In the case of LTWRD, seasonal components are first removed, and Moving Block Bootstrapping (MBB) is then applied to the residuals over the concurrent period to quantify the associated uncertainty, as detailed in Section 3.2.

Due to their distinct sources and statistical properties, IAV and LTWRD are treated as independent contributions to the total historical wind direction uncertainty and combined accordingly.

2.2.3 Spatial Extrapolation Uncertainty

Spatial extrapolation uncertainty pertains to the errors introduced when vertically or horizontally projecting wind measurements from the mast height to the turbine height or across wind farm layout. Vertically, wind veer occurs due to the Coriolis force, frictional forces, and pressure gradients. In offshore environments, the low surface roughness substantially reduces vertical veer. An NREL studied off the coast of Massachusetts reported offshore veer rates ranging from $0.03°m^{-1}$ to a maximum of about $0.1°m^{-1}$ [38]. Consequently, when extrapolating from the anemometer up to hub height, the directional change is typically under one degree, which is negligible compared to other sources of uncertainty such as IAV or instrumental bias. Horizontally, wind direction tends to remain coherent over offshore distances due to the absence of complex terrain effects.

While this component could be incorporated into the total uncertainty budget for complete-

ness in future work, it is excluded from the present study, which focuses solely on inherent uncertainties in the wind data itself.

2.2.4 Project Evaluation Period Variability

Uncertainty in the project evaluation period arises due to the assumption that the modelled operational period aligns with the long-term site average. The uncertainty in the future wind resource can be divided into two components: (1) that due to normal year-to-year variability in the wind climate, which is also called interannual variability, and (2) that due to the risk of long-term climate change. The interannual variability of the reference data period has already been accounted for in the historical wind resource assessment uncertainty. In addition, potential changes in wind direction due to climate change fall outside the scope of this study and are therefore treated as external and independent factors. Thus, period variability uncertainty has been excluded from the current study, but uncertainty from climate change could be incorporated in future research if necessary.

2.2.5 Plant Performance Uncertainty

This study specifically addresses uncertainties related to wind direction and therefore excludes uncertainties associated with plant performance. Such plant performance uncertainties cover factors including, but not limited to, turbine and balance-of-plant availability, electrical efficiency losses, turbine performance, and downtime resulting from extreme environmental conditions. These sources of uncertainty, while relevant to overall energy production assessments, are beyond the scope of this analysis.

In this study, the full directional uncertainty propagated into wake modelling consists of three independent terms: **measurement uncertainty**, **LTWRD uncertainty**, and **IAV uncertainty**. Each component is quantified as the standard deviation of its impact on AEP. Assuming that these sources of uncertainty are statistically independent and have approximately Gaussian distribution, they are combined using the root-sum-square (RSS) method to yield the total standard uncertainty on AEP.

2.3 Numerical Methods

This section describes the data pre-processing process, the wake model used in this research, and the AEP calculation in *PyWake* [39].

2.3.1 Pre-Processing: Site Conditions

To simulate wind farm performance using *PyWake*, site-specific wind conditions must be defined in terms of wind speed, wind direction, turbulence intensity, and the joint probability of these combinations. This setup is handled by *XRSite* class in *PyWake*, which maps ambient conditions to turbine relative coordinates (e.g., crosswind and downwind directions) and propagates wake effects accordingly.

In the selected study sites, strong winds predominantly arriving from specific directions. This anisotropy is evident in the joint probability distribution of wind speed and wind direction. As illustrated in Figure 2.4, wind speed probability density functions are weighted by the directional occurrence frequency, demonstrating how different sectors contribute to the overall wind climate.

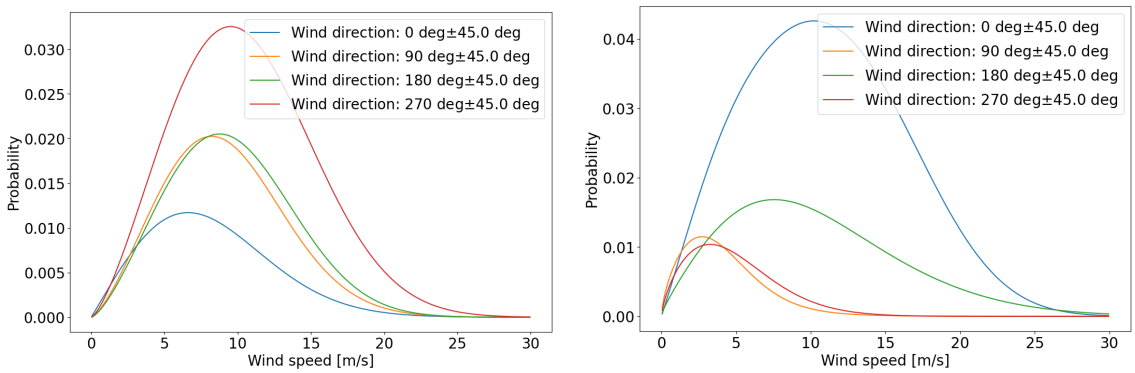

Figure 2.4: Wind speed probability distribution weighted by wind direction frequency for *Site A* (left) and *Site B* (right).

To capture the direction-dependent wind characteristics, the input wind data is binned into 2° directional sectors, resulting in 180 sectors over the full 360° range. Within each sector, a Weibull distribution is fitted to the wind speed data to describe the probability density function of wind speeds. This sector-wise Weibull characterisation allows the simulation to reflect the directional variation of wind energy potential across the site.

The Weibull distribution is defined by two parameters: the scale parameter A, which correlates with the mean wind speed, and the shape parameter k, which controls the spread of the distribution. These parameters are estimated using Maximum Likelihood Estimation (MLE). Although Ordinary Least Squares (OLS) fitting is sometimes preferred in cases of small sample size due to its lower bias, the large number of data points in each sector justifies the use of MLE in this study [40]. To ensure the reliability of the Weibull fits, a minimum sector width of 2° is enforced. This guarantees a sufficient number of observations in each sector, avoiding unreliable parameter estimation due to undersampling.

2.3.2 Wake Model

There is currently no universally accepted model that can fully capture all flow phenomena and precisely predict energy production across diverse wind farm configurations. While computational fluid dynamics approaches offer high-fidelity representations of flow behaviour, they are computationally expensive and often impractical for large-scale AEP assessments. In contrast, engineering wake models provide a more practical balance between computational efficiency and accuracy. However, these models depend on a combination of sub-models, including wake deficit, superposition, and blockage models. The selection of sub-models must be tailored to the specific site conditions and layout. A model configuration that performs well in one context may not be suitable in another.

In this study, the *PyWake* framework is adopted as the simulation platform due to its modular structure and flexibility. The sub-model combinations and parameters setting used in this work follow the recommendations outlined in the PyWake documentation [41]. The following section outlines the specific model setup used in this work and the rationale behind these choices.

2.3.2.1 Wake Deficit Model

The wake model employed in this study is TurbOPark, a turbulence-optimised Gaussian-type engineering wake model developed by Ørsted. TurbOPark is designed to improve the prediction of wake losses by incorporating the influence of both ambient and wake-generated turbulence on the wake expansion process. It dynamically adapts the wake growth rate as a function of the downstream turbulence, resulting in improved directional

sensitivity and wake recovery estimates in offshore wind farms.

The decision to use a Gaussian-type model follows general modelling guidance from DTU PyWake documentation [41]. Compared to top-hat models, Gaussian models avoid discontinuities and better conserve momentum in isolated wakes. Among the Gaussian formulations available in PyWake, TurbOPark is distinguished by its dynamic wake expansion governed by local turbulence intensity (TI), rather than using a fixed or overly simplified wake growth rate. Following the relation from Pedersen et al. [42], TurbOPark models models the wake expansion rate as a function of local turbulence intensity $I(x)$:

$$\frac{dD_w}{dx} = AI(x) \tag{2.1}$$

where D_w is the wake diameter, x is the downstream distance from the turbine, and A is the wake expansion coefficient. The wake expansion rate is proportional to the local turbulence intensity. A value of $A = 0.04$ is adopted, based on Ørsted's validation against 19 offshore wind farms using SCADA data, as discussed in Nygaard et al. [43].

To calculate $I(x)$, TurbOPark uses the Frandsen wake turbulence model [44] to combine the ambient turbulence intensity I_0 with the additional turbulence generated in the wake $I_w(x)$ as:

$$I(x) = \sqrt{I_0^2 + I_w^2(x)} \tag{2.2}$$

The wake-generated turbulence $I_w(x)$ is empirically modelled as[42]:

$$I_w(x) = \frac{1}{c_1 + c_2 \frac{x/D}{C_T(V_{in})}} \tag{2.3}$$

where D is the rotor diameter, C_T is the thrust coefficient of the turbine as a function of the incoming wind speed V_{in}, and $c_1 = 1.5$ and $c_2 = 0.8$ are the two empirical constants given in [43, 42] based on wind tunnel and field data.

TurbOPark's wake shape inherently accounts for long-term averaging, integrating the effects of wake meandering, shear-layer growth, and small-scale mixing into a single, smooth profile. This makes it particularly suitable for estimating the impact of wind direction uncertainty on the same base. Moreover, the model's greater sensitivity to directional changes, as compared to traditional Jensen-type models, improves its capability in capturing the AEP deviation caused by shifts in directional frequency distributions.

A sensitivity analysis of the wake expansion coefficient A and turbulence model parameters c_1 and c_2 is discussed in Section 5.4.2.

2.3.2.2 Superposition Model

To account for the combined influence of multiple upstream wakes on a downstream turbine, a superposition model must be applied. This model determines how individual wake deficits interact and accumulate at the rotor plane of a given turbine. The choice of superposition model is closely tied to the definition of effective wind speed. In this study, the setting *use_effective_ws=True* is enabled, which means that the effective wind speed at each turbine accounts for the local wake-modified flow field rather than assuming a uniform freestream velocity throughout the array. Under this setting, the *LinearSum* model provides a coherent and physically correct approach to combining wake deficits, as it avoids the risk of overestimating cumulative losses or producing negative wind speeds in highly waked regions.

Alternative superposition methods recommended in the documentation, such as *SquaredSum*, may offer improved approximations of wake effects in very deep wind farm arrays.

However, these methods assume a uniform freestream wind speed across all turbines, which is a condition that is not satisfied when using effective wind speed formulations. Additionally, SquaredSum-based methods are known to artificially amplify wake losses by effectively double-counting overlapping wake deficits, which can lead to overestimation of power loss for dense wake interactions.

For these reasons, *LinearSum* was selected as the superposition model in this study. It ensures compatibility with the chosen wake deficit model and the physical interpretation of effective wind speed at each turbine.

2.3.2.3 Blockage Model

In addition to wake effects, blockage refers to the deceleration of the flow upstream due to flow obstruction. This phenomenon leads to a region of reduced wind speed ahead of the turbines, causing the local flow field to differ from free-streamed conditions. If unaccounted for, blockage can result in systematic overestimation of energy production in wind farm simulations.

Studies from DTU estimate that blockage-related overprediction can lead to power output increases of up to 2% in aligned turbine layouts [45]. To mitigate this bias, a blockage correction is incorporated into the modelling framework. In this study, the *SelfSimilarityDeficit* model from the PyWake package is employed to represent blockage effects. The *SelfSimilarityDeficit* model is a simplified analytical induction model that captures the upstream flow deficit using the self-similarity assumption for flow expansion. It provides a physically motivated correction to inflow condition upstream of turbines, thereby improving the accuracy of predicted AEP.

2.3.2.4 Engineering Wind Farm Model

In summary, this study adopts a modular modelling approach using PyWake, combining three sub-models: the **TurbOGaussianDeficit** for wake deficit, **LinearSum** for wake superposition, and **SelfSimilarityDeficit** for blockage effects. This configuration is referred to as *Nygaard_2022_adjusted* in Table 2.1. The *use_effective_ws=True* setting is enabled to ensure that each turbine's performance is evaluated based on the locally modified wind field rather than a uniform freestream value. This ensures internal consistency between the deficit and superposition models, as recommended in the PyWake documentation.

To integrate these components, the **All2AllIterative** engineering wind farm model is employed. Among the three general models available in PyWake:*PropagateDownwind*, *PropagateUpDownIterative*, and *All2AllIterative*, only *All2AllIterative* fully accounts for both wake and blockage effects by iteratively solving the mutual interactions among all turbines. In each iteration, wake deficits from all turbines are superposed at every turbine location, and the resulting effective wind speeds are updated until convergence is reached or the maximum number of iterations is exceeded.

Although PyWake includes predefined composite models such as *Nygaard_2022*, which bundle wake, turbulence, and superposition models according to published literature [42], the default *Nygaard_2022* implementation relies on the *PropagateDownwind* solver. This method estimates the wind speed at each turbine based solely on upstream influences, thereby neglecting upstream-facing blockage effects. For this reason, a modified version of *Nygaard_2022* is constructed using the *All2AllIterative* solver to properly capture the full array of wake and blockage interactions.

While this study does not primarily aim to benchmark wake models, three industry-relevant engineering models (summarised in Table 2.1) are included for comparison. These are: the adjusted *Nygaard_2022* model, which applies the updated TurbOPark framework

Table 2.1: Wake Model Configuration Overview

Name	Wind Farm Model	Wake Deficit	Blockage	Superposition
Nygaard_2022_adjusted	All2AllIterative	TurboGaussianDeficit	SelfSimilarityDeficit	LinearSum
Nygaard_2020_Jensen	All2AllIterative	TurboNOJDeficit	SelfSimilarityDeficit	LinearSum
Ott_Nielsen_2014_Blockage	All2AllIterative	FugaDeficit	FugaDeficit	LinearSum

with Gaussian wake deficit and blockage representation; the original *Nygaard_2020* TurbOPark model based on the Jensen wake formulation [46]; and the *OttNielsen_2014* implementation, which incorporates the Fuga model to account for both wake and blockage effects [47]. The comparative results of these models are discussed in Section 5.4.1.

2.3.3 AEP calculation

AEP is computed using the *PyWake* framework, which models wind farm performance by combining a wind resource characterisation with turbine-specific power curves and wake interactions. AEP is calculated by integrating the expected power output over all wind directions and wind speeds, weighted by the joint probability distribution $f(U|\theta)$, the wind speed distribution conditioned on wind direction.

PyWake handles this using a sector-wise Weibull fit, as introduced in Section 2.3.1. For each direction θ, the turbine power output $P(U)$ is evaluated based on the local inflow conditions and then weighted by $f(U|\theta)$. The total AEP is computed by aggregating across all directional sectors:

$$AEP = T \sum_{\theta=0^\circ}^{360^\circ} F(\theta) \sum_{i=1}^{N_{\text{turb}}} \int_{U_{\text{cut-in}}}^{U_{\text{cut-out}}} P_i(U) f(U|\theta)\, dU \tag{2.4}$$

where T is the number of seconds in a year, $F(\theta)$ is the directional frequency for sector θ, $P_i(U)$ is the power curve applied at turbine i, and $f(U|\theta)$ is the wind speed distribution for that sector.

Unlike uniform inflow models, this study enables effective wind speed in *PyWake*, which accounts for local wind conditions at each turbine based on wake propagation and flow deficits. As a result, each turbine may experience a different inflow speed even within the same site condition, leading to a non-uniform power output across the farm. The total AEP, calculated as the sum of all turbines, thus become more directionally sensitive, since the extent of wake losses and local wind speed variation depends on wind direction.

3 Uncertainty Quantification

As concluded in Section 2.2.1, the uncertainty in wind direction are primarily from three sources: measurement uncertainty, LTWRD uncertainty, and IAV uncertainty. In this chapter, the quantification method of three uncertainty sources will be introduced respectively.

3.1 Measurement Uncertainty

This section quantifies how systematic measurement uncertainty in wind direction propagates into uncertainty in AEP. Directional errors from measurements primarily manifest as a fixed angular offset in the measured wind rose. To determine a reasonable value for this offset and assess its impact on AEP, the analysis proceeds in two steps. First, a conservative estimate of the directional measurement uncertainty is obtained by aggregating individual error sources according to standard uncertainty propagation principles. Next, a sensitivity analysis is conducted to evaluate how a fixed directional offset influences the resulting AEP. All details are presented in the two following subsections.

3.1.1 Estimation of Directional Measurement Uncertainty

Based on the ECN report [48], the systematic uncertainty in wind direction measurement includes the following components: (1) wind vane accuracy, (2) sensor resolution, (3) mast alignment error, (4) flow distortion, and (5) data acquisition.

(1) Wind vane measuring accuracy

According to wind vane manufacturer specifications [49], the wind vane used in this study to represent measuring instrument has an accuracy of ±1°. Under the assumptions of the Guide to the Expression of Uncertainty in Measurement (GUM) [50], where the distribution of measurements is assumed to be approximately normal, an expanded uncertainty of ±1° with a coverage factor $k = 2$ corresponds to a standard uncertainty of:

$$u = \frac{1^\circ}{2} = 0.5^\circ$$

(2) Wind vane sensor resolution

The resolution of the wind vane is specified as 0.1° [49]. However, in practice, resolution uncertainty contributes negligibly compared to other dominant sources such as measuring accuracy or alignment error. According to GUM [50], when an uncertainty component is significantly smaller than others and has minimal influence on the combined standard uncertainty, it may be omitted. Therefore, the uncertainty due to sensor resolution is not included in the final uncertainty budget in this study.

(3) Mast alignment error

Notably, manual orientation during mast installation is still commonly performed using hand compasses (as shown in appendix A). Therefore, the two main contributors to alignment error is hand compass accuracy and operator error. According to the Handbook [51, p. 650-1.39], the typical accuracy of a hand compass under optimal conditions is ±2°. In the absence of more detailed information on the distribution of compass errors, a rectangular distribution is assumed. Per GUM [50], the corresponding standard uncertainty is:

$$u = 2^\circ/\sqrt{3} \approx 1.15^\circ$$

Operator-induced error, such as misreading the compass or improper setup, introduces additional uncertainty [52]. While the exact value is difficult to quantify universally, a combined mast alignment of 2° is the suggested number reported in industry studies (e.g. ECN [48]). Thus, a conservative standard uncertainty of $u = 2^\circ$ is used to account for total alignment error.

(4) Flow distortion
Wind flow distortion caused by mast structure and mounting effects is typically estimated to contribute less than 1.5° to the directional uncertainty. Following the ECN industry study [48], a conservative estimate of $u = 1.5^\circ$ is adopted in this study.

(5) Data acquisition
The data logger used in the met mast system has a digital resolution of 0.1° [49]. However, due to the signal digitisation, this uncertainty is considered negligible in the total budget and is thus assigned $u = 0^\circ$.

Since all the above uncertainties are assumed to be independent components, the total standard uncertainty in wind measurements is given by the root-sum-square:

$$\sigma_{\text{measurement, total}} = \sqrt{0.5^2 + 2.0^2 + 1.5^2} \approx 2.54^\circ \approx 3^\circ \tag{3.1}$$

The numerical values of estimate uncertainty should not have excessive digits. As the uncertainty estimation is conservative in this study and from limited information, there is an uncertainty in the estimated uncertainty. Hence, over-precision is misleading. The total wind direction measurement is rounded up to **3 degrees** in the subsequent analysis.

3.1.2 Propagation of Directional Uncertainty to AEP

Measurement uncertainty in wind direction primarily originates from setup errors which introduce systematic offsets in directional readings. To quantify how this directional measurement uncertainty propagates to uncertainty in AEP, a sensitivity analysis is performed by artificially shifting the entire wind rose.

In this analysis, the wind rose is rotated by $\pm 6^\circ$ in increments of 2°, both clockwise and counterclockwise. The step size of 2° is chosen to match the directional bin width used in the Weibull sector fitting in Section 2.3.1, ensuring that each bin contains sufficient data for a reliable statistical fit. The total shift range of $\pm 6^\circ$ correspond to two standard deviations (2σ) based on a directional measurement uncertainty of 3° from Equation (3.1). This covers approximately 95% of expected deviations when assuming normally distributed errors.

For each shifted wind rose, the AEP is recalculated using the wake model setup introduced earlier. The resulting relationship between directional shift and AEP change is evaluated using linear regression. As shown in Section 4.1, the dependence is well-approximated by a linear model, with coefficients of determination (R^2) of 0.993 for *Site A* and 0.985 for *Site B*, confirming the linearity of the response in the tested range.

Based on standard uncertainty propagation principles from [50] and [53], when the input-output relationship is approximately linear, the uncertainty in AEP due to directional measurement errors can be estimated as:

$$\sigma_{AEP} = \left|\frac{dAEP}{d\theta}\right| * \sigma_\theta \tag{3.2}$$

where σ_θ is the standard uncertainty in the directional measurement and $dAEP/d\theta$ is the sensitivity obtained from the regression slope. The final AEP uncertainty is computed using a directional uncertainty σ_θ of 3 degrees, calculated in Equation (3.1).

3.2 LTWRD Uncertainty

As mentioned in Section 2.2.2, Long-Term Wind Rose Difference (LTWRD) refers to the observed discrepancy between the wind direction distributions derived from short-term measurements and those from long-term reanalysis datasets. Unlike wind speed, wind direction is typically not corrected or calibrated during long-term extrapolation procedures. As a result, differences between the short-term and long-term wind roses can introduce additional uncertainty in directional analysis and AEP estimation.

For *Site A*, the concurrent period between the mast measurements and reanalysis data spans only 5.5 years. This limited duration restricts the number of independent samples that can be used for uncertainty quantification. Non-overlapping sampling permits only five distinct AEP estimates, while overlapping subsampling introduces high similarity among the samples, leading to clustering and potentially biased estimates, as detailed in Appendix B.1. Therefore, resampling strategy is necessary to properly represent the variability. This study employs a Monte Carlo approach, integrated with Moving Block Bootstrapping (MBB), to estimate the LTWRD uncertainty. Directly applying MBB to the original wind speed time series leads to synthetic samples with seasonally imbalanced representations, resulting in an artificial inflation of AEP error distribution, as discussed in Appendix B.2. To preserve the seasonal characteristics of the data while allowing temporal resampling, MBB is instead applied to the residuals, as further described in Section 3.2.2. The LTWRD uncertainty is then quantified as the standard deviation of the AEP error across synthetic samples generated from bootstrapped time windows.

3.2.1 Spectral Analysis

Prior to implementing any resampling methods, a spectral analysis is conducted to identify dominant temporal correlations in wind speed time series. In wind energy applications, spectral analysis provides a critical insight into the scales of variability. The power spectral density (PSD) describes how variance is distributed across frequency components in the time series.

In this study, the logarithmic form of the PSD, $fS(f)$, is used to better visualise the variance contribution from each frequency decade. The variance of a process can be computed as [54]:

$$\sigma^2 = \int_{-\infty}^{\infty} fS(f)\ d\,\log f = 2\int_{0}^{\infty} fS(f)\ d\,\log f$$

Hence, when plotted against logarithmic frequency, the area under the $fS(f)$ curve corresponds to half the total variance. This representation allows for easier identification of which time scales dominate the variability.

Figure 3.1 presents the log-frequency spectral density of 10-minute wind speed data from *Site A*. Distinct peaks are visible at daily and semi-daily timescales, consistent with diurnal and semi-diurnal cycles. While a flattening platform is observed near the 1-week mark, no sharp peak is present. This suggests a transition zone between low-frequency trends (e.g., seasonal variation) and high-frequency fluctuations (e.g., turbulence), rather than a dominant weekly cycle. Figure 3.2 shows the corresponding $S(f)$ plot, which helps confirm the interpretation. Although a peak at the annual scale is theoretically expected, it is not prominent due to the limited 5.5-year dataset.

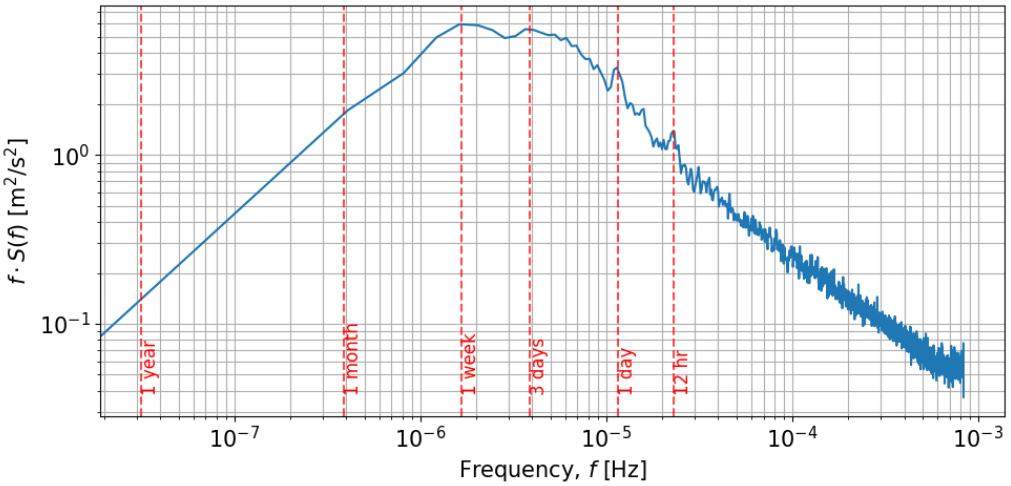

Figure 3.1: Log-frequency spectral density ($fS(f)$) of wind speed at *Site A*, showing the relative contribution of each frequency decade to total variance.

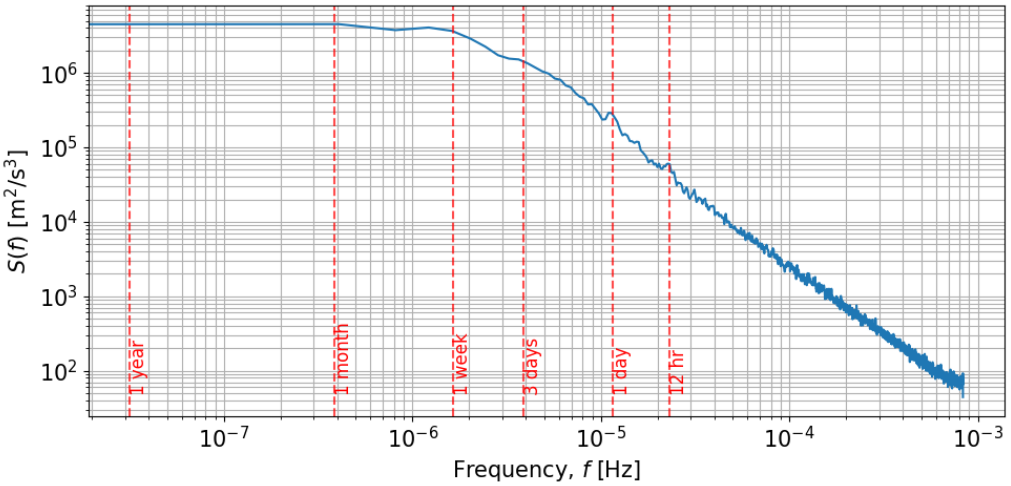

Figure 3.2: Frequency spectral density ($S(f)$) of wind speed at *Site A*.

3.2.2 MBB apply to residuals

Traditional bootstrapping assumes that data points are independently and identically distributed, which is not valid for wind time series. To preserve the inherent temporal structures such as autocorrelation and seasonal variability, the Moving Block Bootstrapping (MBB) method [55] is used in this study. MBB resamples the data in blocks of contiguous time intervals, preserving local dependencies and allowing the resampled series to maintain realistic temporal structure.

The selection of an appropriate block size is critical. If the blocks are too short, they may fail to capture autocorrelated structures; if too long, they reduce the number of independent resamples, increasing estimation bias and reducing variability. Since there is no universally optimal block size for wind time series, several approaches can guide this choice.

One approach is based on spectral analysis, which identifies dominant timescales of variability. In the spectral density plot for *Site A*, peaks appear at 12-hour and 1-day intervals,

corresponding to diurnal and semi-diurnal cycles. Although no distinct peak is observed at longer timescales, synoptic variability (3–7 days) is commonly found in wind data. Consequently, block size selection should aim to preserve periods longer than the diurnal peak variance.

Another common method uses the autocorrelation function (ACF), with the block size determined by the lag at which the ACF falls below a threshold (e.g., 0.1) to indicate a loss of temporal dependence. While an ACF plot was not generated for this study due to time constraints, previous research by Monahan [56] shows that ACFs of surface wind vectors in midlatitude ocean regions typically decay below 0.1 after approximately 4 days, while tropical regions retain higher autocorrelation (e.g., 0.3 after 4 days). Since the site in this study is located in the midlatitudes, a conservative block size of 7 days is adopted to ensure the preservation of synoptic-scale structures while allowing sufficient resampling variability.

The wind time series is modelled as a stochastic process composed of deterministic and random components. Following the decomposition from Climate analysis textbook[57], the wind variable $X(T)$ is represented as:

$$X(T) = X_{\text{trend}}(T) + X_{\text{out}}(T) + S(T) \cdot X_{\text{residual}}(T) \tag{3.3}$$

where T denotes continuous time, $X_{\text{trend}}(T)$ represents seasonality signal by applying seasonal-trend decomposition, $X_{\text{out}}(T)$ accounts for rare outlier events, $S(T)$ is a time-dependent variability function, and $X_{\text{residual}}(T)$ is a stationary noise process with zero mean and unit variance. Given the circular nature of wind direction, outliers are not easily defined and thus $X_{\text{out}}(T)$ is omitted. The time series is then discretised as:

$$X(i) = X_{\text{trend}}(i) + S(i) \cdot X_{\text{residual}}(i) \tag{3.4}$$

where i denotes discrete timestamps. To ensure continuity and validity of this model, missing timestamps in the data are filled using nearest-neighbour interpolation for wind speed and circular interpolation (”wrap”) for wind direction.

The components of Equation (3.4) are derived separately for wind speed and wind direction (see Section 3.2.2.1 and Section 3.2.2.2, respectively).

3.2.2.1 Wind Speed

As described in Equation (3.4), the wind speed time series is first decomposed into trend, variability, and residual components before applying the MBB method to residuals. This decomposition enables the generation of synthetic time series that preserve the temporal structure of the original data.

The log-frequency plot of the power spectral density (Figure 3.1) revealed prominent peaks at 12-hour and 24-hour (daily), indicating the presence of both diurnal variability in the wind speed signal. To accurately capture these multi-scale temporal patterns, the Multiple Seasonal-Trend Decomposition using LOESS (MSTL) is employed [58, 59]. This method allows for the inclusion of multiple seasonal components and provides decomposition based on those components. Unlike a simple moving average or low-pass filter, which can introduce distortions when trends shift overtime, MSTL can separate daily and annual patterns while preserving non-stationary behaviour in the trend component.

In this context, the term ”trend” in MSTL refers to the long-term variation over the full time series, whereas in the MBB decomposition Equation (3.4), the term ”trend” (denoted $X_{trend}(i)$) refers to the seasonal component that is used as the deterministic based onto which residuals are resampled.

The MSTL decomposition is configured with two seasonalities: daily and yearly cycles. Although the 12-hour peak is visible in Figure 3.1, the 24-hour cycle is the dominant diurnal pattern. Additionally, when decomposing using a 24-hour seasonal component, the MSTL method also implicitly capture higher-order harmonics. On the other hand, while the annual signal is less dominant in the spectral plot, it is included to ensure the resulting synthetic time series retains the annual variability necessary for AEP estimation.

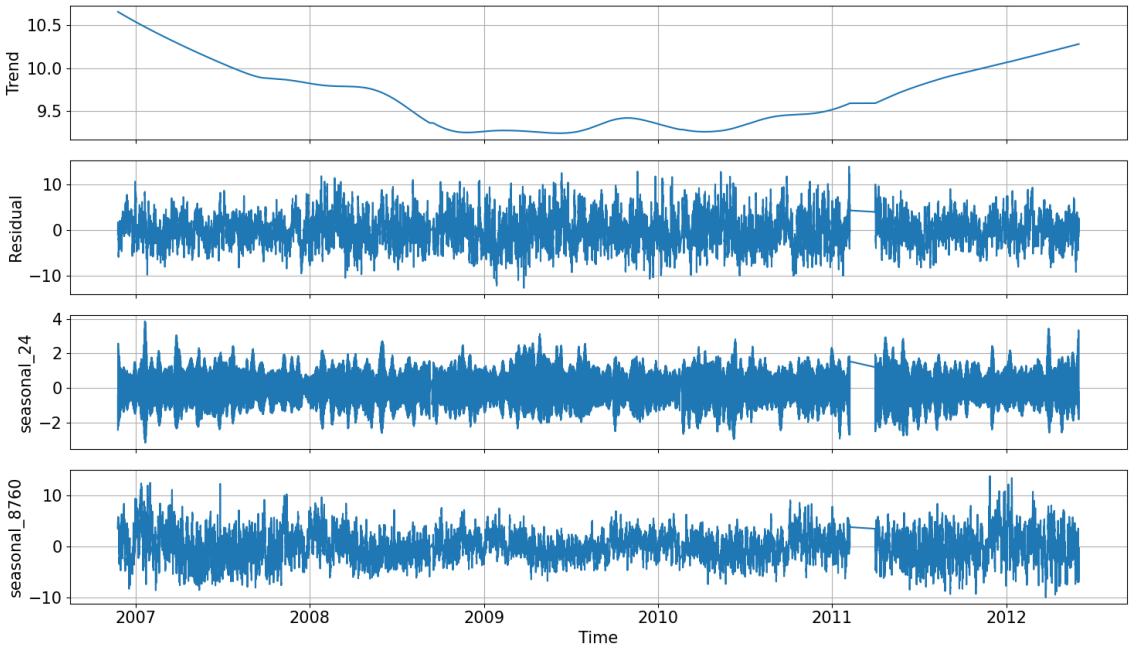

Figure 3.3: MSTL decomposition of the wind speed time series at the *Site A* mast. Each subplot shows the trend, residual, and seasonal components (daily and yearly).

The reconstructed seasonal base signal $WSP_{trend}(i)$ is obtained by summing the extracted seasonal components from MSTL. This is illustrated in Figure 3.4, where the seasonal base is plotted alongside with the original wind speed data. This based is then used as the deterministic foundation in the MBB framework, and residuals are resampled in 1-week blocks to preserve short-term autocorrelation.

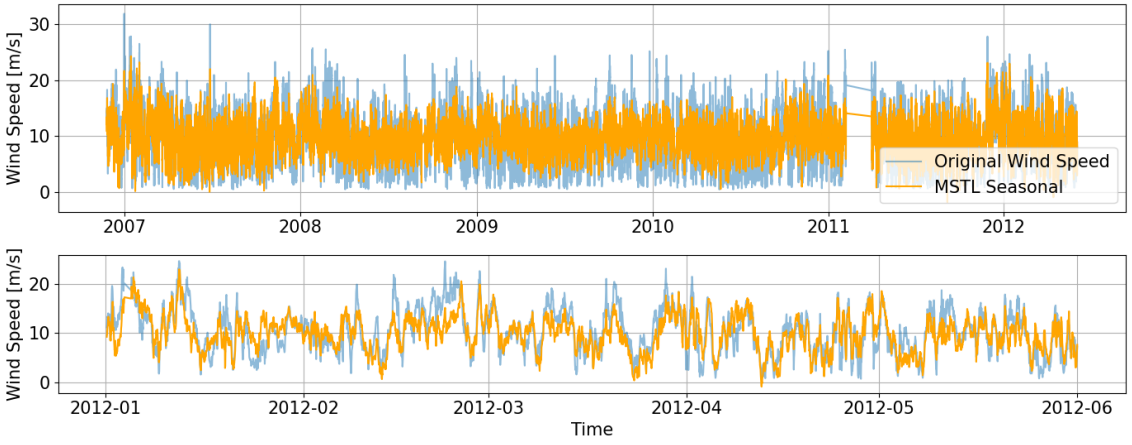

Figure 3.4: Synthetic wind speed using MSTL seasonal base: full record (top) and zoomed view of 2012 (bottom).

To characterise the short-term variability of wind speed $S_{\text{wsp}}(i)$, a local standard deviation is computed using a centred rolling window of one week (block size = 168):

$$S_{\text{wsp}}(i) = \text{std}\left(WSP(i \pm \tfrac{\text{block size}}{2})\right) \tag{3.5}$$

This scaling function accounts for time-varying fluctuations in the amplitude of wind speed variability.

Following the decomposition model introduced in Equation (3.4), the residual (or normalised noise) component of the wind speed is computed as:

$$WSP_{\text{residual}}(i) = \frac{WSP(i) - WSP_{\text{trend}}(i)}{S_{\text{wsp}}(i)} \tag{3.6}$$

3.2.2.2 Wind Direction

Wind direction is a circular variable, which means that direct arithmetic operations can lead to incorrect results near the $0°$ and $360°$ boundary. To address this, all preprocessing is performed in vector space by converting directional angles into unit vectors. Each wind direction $\theta(i)$ is mapped to Cartesian coordinates via: $(u(i), v(i)) = (\cos(\theta(i)),\ \sin(\theta(i)))$. To extract the trend in wind direction, MSTL is applied separately to the u and v components using the same seasonality cycles as used for wind speed, including daily and yearly cycles. The extracted vectors $(u^*(i), v^*(i))$ may deviate from the unit circle due to filtering. Before converting back to angular space, the vectors are renormalised:

$$\text{norm}(i) = \sqrt{u^*(i)^2 + v^*(i)^2}, \quad u^*(i) \leftarrow \frac{u^*(i)}{\text{norm}(i)}, \quad v^*(i) \leftarrow \frac{v^*(i)}{\text{norm}(i)} \tag{3.7}$$

This guarantees that the angle recovered using the function $atan2(v^*, u^*)$ lies on the unit circle. The final trend in wind direction, $WD_{\text{trend}}(i)$, is then:

$$WD_{\text{trend}}(i) = \left(\text{atan2}(v^*(i), u^*(i)) \times \frac{180°}{\pi}\right) \bmod 360° \tag{3.8}$$

Figure 3.5 shows the MTSL wind direction trend for the full dataset and a zoomed-in period from 2012.

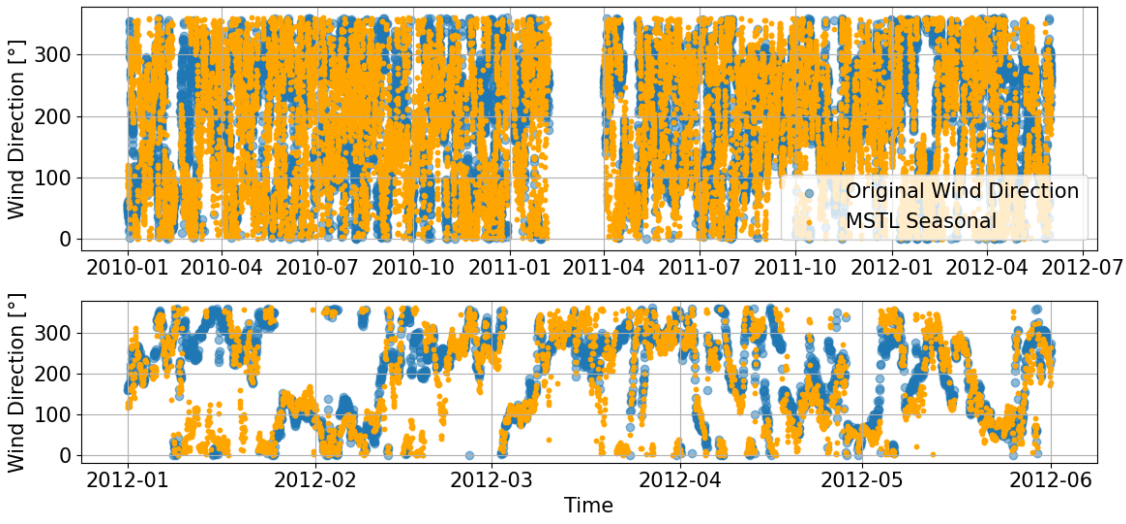

Figure 3.5: Synthetic wind direction using MSTL seasonal base: full record (top) and zoomed view of 2012 (bottom).

The short-term variability of wind direction $S_{\text{wd}}(i)$ is quantified using the circular standard deviation. Similar to wind speed, a rolling window approach is used; however, since wind direction is processed as unit vectors on the unit circle, the resultant vector length $R(i)$ is first computed and then converted to a circular standard deviation [60]:

$$R(i) = \sqrt{\left(\frac{1}{N}\sum_{j=1}^{N} u_j\right)^2 + \left(\frac{1}{N}\sum_{j=1}^{N} v_j\right)^2} \tag{3.9}$$

$$S_{\text{wd}}(i) = \sqrt{-2 \ln R(i)} \tag{3.10}$$

To compute the residual (i.e., noise) component, direct subtraction of angular values is avoided. Instead, the smallest signed angular difference is computed via Euler's identity ($e^{j\theta} = cos\theta + jsin\theta$):

$$WD_{\text{residual}}(i) = \frac{1}{S_{\text{wd}}(i)} \cdot \text{deg}\left(e^{j(WD(i) - WD_{\text{trend}}(i))}\right) \tag{3.11}$$

This method ensures the correct sign and magnitude of the angular difference, avoiding errors such as interpreting a -5° offset as 355°.

Once the residuals are calculated using Equation (3.6) and Equation (3.11), MBB is applied using overlapping 1-week blocks to generate synthetic residual samples. These resampled residuals are then combined with the seasonal component and the average trend from the MSTL decomposition to reconstruct synthetic time series that preserve both the seasonal structure and the short-term statistical variability. By applying MBB to the residuals and extract the trend from the raw data, this approach effectively removes the influence of long-term trends and interannual variability, isolating the uncertainty attributable to short-term fluctuations.

To isolate the effect of wind direction uncertainty, wind speed is held constant across simulations. Synthesised wind speed from the mast measurements is used for both simulations, while wind direction inputs are drawn separately from the reanalysis and mast datasets. A total of 200 Monte Carlo simulations are performed, each corresponding to a randomly resampled time window. The convergence analysis for the number of samples is provided in Appendix B.3. The LTWRD uncertainty is finally defined as the standard deviation of the resulting 200 AEP differences between the two directional inputs.

3.2.3 Disjunct Sampling

As previously mentioned, the measured data are available at 10-minute resolution, whereas the reanalysis data have a 1-hour temporal resolution. Li [61] demonstrated that increasing the sampling interval can underestimate average wind speeds and reduce statistical representativeness of the wind regime. Although that study focuses on extreme wind speeds for structural design, the same principle applies to AEP estimation, where temporal mismatch can distort directional distributions and sectoral frequency content. To enable a consistent comparison and estimation of LTWRD uncertainty between the two datasets, and to quantify the uncertainty introduced by temporal sampling resolution, a disjunct sampling approach was applied to the 10-minute measured data. In this method, only the observations recorded at the start of each hour were retained to construct an effective hourly time series, thereby mimicking a lower-frequency measurement regime

without applying block averaging. This disjunct dataset was compared against the full 10-minute dataset using 200 randomly selected 1-year windows, with AEP estimated separately from both using a sector-wise Weibull fitting approach followed by wake modelling with PyWake.

The rationale for using disjunct rather than block-averaged sampling lies in the frequency-domain characteristics of the wind signal. Spectral analysis indicates that the dominant energy content in the wind variability occurs at timescales longer than one hour, as shown in Figure 3.1. Consequently, no substantial difference is expected between AEP estimates obtained using disjunct sampling and those based on block-averaged data. This expectation is further supported by the negligible difference observed between AEP estimates derived from disjunct sampling and those obtained using conventional block-averaging method, as discussed in Appendix B.4. Thus, disjunct sampling provides a practical and computationally efficient way to assess the sensitivity of AEP estimation to temporal resolution.

To further examine the representativeness of the disjunct dataset, histograms of wind speed and direction for both the 10-minute and 1-hour disjunct datasets are plotted side-by-side in Figure 3.6. As expected from the spectral analysis, the visual difference between the two distributions is minimal, indicating that disjunct sampling does not significantly distort the underlying wind regime at *Site A*. Nonetheless, even small discrepancies can lead to measurable differences in AEP, particularly when propagated through nonlinear wake effects and directional weighting. Therefore, the impact of disjunct sampling must still be treated as a source of uncertainty and quantified accordingly.

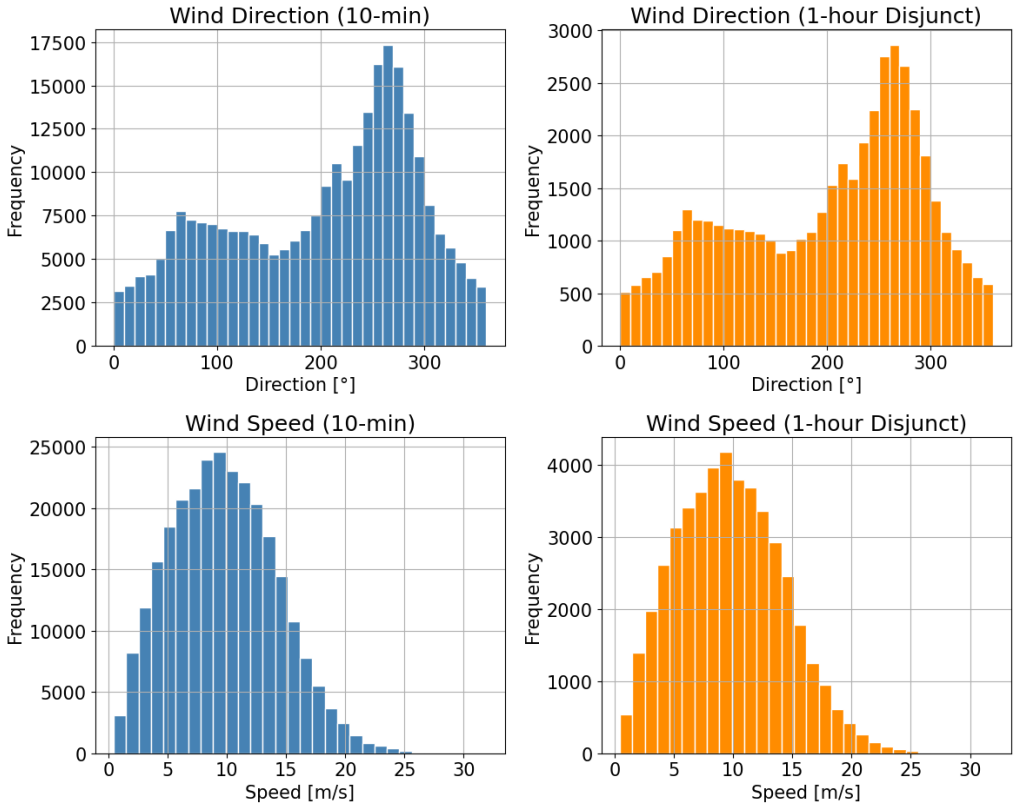

Figure 3.6: Comparison of wind direction and wind speed distributions between 10-minute and 1-hour (disjunct) datasets at *Site A*. Top row: wind direction; bottom row: wind speed. Left: 10-minute data; right: disjunct 1-hour data.

Since disjunct sampling represents an independent source of uncertainty, it is treated as a separate component and can be combined with other LTWRD uncertainty contributions through a root-sum-square approach.

3.2.4 Overview of LTWRD Uncertainty Estimation Procedure

This subsection provides a concise summary of the methodology used to estimate LTWRD uncertainty. As the theoretical framework and technical implementation are detailed in the previous subsections, this overview aims to clarify the practical workflow without overwhelming the reader with methodological depth.

To quantify LTWRD uncertainty, the following steps are performed:

1. **Decompose the measured and reanalysis wind time series** into three components: trend (i.e., seasonality), variability, and residual, using Multiple Seasonal-Trend Decomposition using LOESS (MSTL). Wind speed and wind direction are treated separately, with direction represented in vector space to preserve circular properties.
2. **Apply Moving Block Bootstrapping (MBB)** on residuals using 1-week blocks. This produces synthetic samples that retain the autocorrelation and seasonal characteristics of the original data.
3. **Reconstruct 1-year synthetic time series** by combining resampled residuals with the original seasonal base. Wind speed is held fixed (to ensure consistent sectoral Weibull fits across measurements and reanalysis), while wind direction is drawn separately from either the measured or reanalysis datasets, allowing the isolation of directional effects.
4. **Estimate AEP** for each synthetic 1-year time series (200 Monte Carlo samples in total) using PyWake's wake modelling framework. LTWRD data bias uncertainty is then defined as the standard deviation of AEP differences between simulations using measured vs. reanalysis-based synthetic wind direction inputs.
5. **Quantify disjunct sampling uncertainty** by comparing AEP estimates from the full 10-minute measurement dataset and a 1-hour disjunct subset, using the same methodology.
6. **Combine both uncertainty components** using root-sum-square method to yield the total LTWRD uncertainty.

The total LTWRD uncertainty thus represents the combined impact of temporal resolution and climatological directional mismatch.

3.3 IAV Uncertainty

While the interannual variability (IAV) of wind speed has been widely studied and well-documented, much less attention has been paid to IAV in wind direction. Existing studies on wind direction variability primarily focus on frequency changes across directional sectors [62, 63, 64], rather than attempting to quantify directional variation in degrees. This is largely due to the circular nature of wind direction data, which makes conventional statistical metrics (such as the mean or standard deviation) problematic and meaningless. As Keevallik et al. [65] emphasise, long-term averages of wind direction are particularly misleading when distributions are multimodal or highly dispersed.

To test this assertion, a trial application of the vector averaging and Yamartino method [18] was carried out to estimate the standard deviation of wind direction and assess whether

the synthesised directional data could represent the actual distribution. The results confirmed that the vector averaged and Yamartino-derived standard deviation did not adequately reproduce the characteristics of the original dataset, supporting Keevallik et al.'s conclusion. Details of this test are provided in Appendix C.

Some researchers (e.g., [66]) have used average wind direction over short periods, but this approach is not reliable for long-term climatological assessments. Instead, wind roses (i.e., sectoral frequency distributions) are typically used to represent directional characteristics. However, simply comparing the differences between two wind rose profiles is insufficient to characterise AEP impact, since directional changes may have highly asymmetric effects depending on the wind farm layout. For instance, if the dominant wind direction is shifted from south to west or to east could result in similar wind rose deviation metrics but yield different AEP outcomes.

To demonstrate that sectoral frequency distribution may not be a reliable predictor of AEP variation, a sensitivity test was conducted using reanalysis data (Figure 3.7). In this experiment, only the wind direction frequency (i.e., the shape of the wind rose) was varied across years, while the sectoral Weibull parameters (i.e., shape and scale in each sector) were held constant. This isolates the impact of directional variability on AEP. The y-axis of Figure 3.7 shows the Normalised Wind Rose Difference (WRD), a metric that quantifies the deviation in directional frequency between a specific year and the multi-year reference distribution. It is defined as:

$$WR^{(j)}_{\text{diff,norm}} = (\Sigma^{all}_{\text{sec},i=0} | f^{(j)}_{\text{sec},i} - f^{(\text{ref})}_{\text{sec},i} |)/2 * 100 \tag{3.12}$$

where $f^{(j)}_{\text{sec},i}$ denotes the frequency in sector i of the wind rose for the j-th year, and $f^{(\text{ref})}_{\text{sec},i}$ is the corresponding frequency from the full-period (multi-year) wind rose. This metric, first introduced by Kjeller Vindteknikk [22], ranges from 0% (identical wind roses) to 100% (no overlap in any sector), and is normalised by a factor of 2 to reflect the maximum possible absolute difference between two valid probability distributions.

Figure 3.7 shows that AEP variation correlates strongly with the mean wind speed, while showing no consistent linear relationship with WRD when wind speed is held constant. This is partly because the mean wind speed is itself calculated as a directionally weighted value, meaning that changes in the wind rose distribution can indirectly influence AEP via altered sectoral weighting. However, WRD only measures the magnitude of the difference between two wind roses without capturing how those differences are distributed across sectors. Two wind roses with the same WRD can yield very different sectoral weightings, and thus different weighted mean wind speeds and AEP outcomes. This makes WRD a poor standalone predictor of AEP variation.

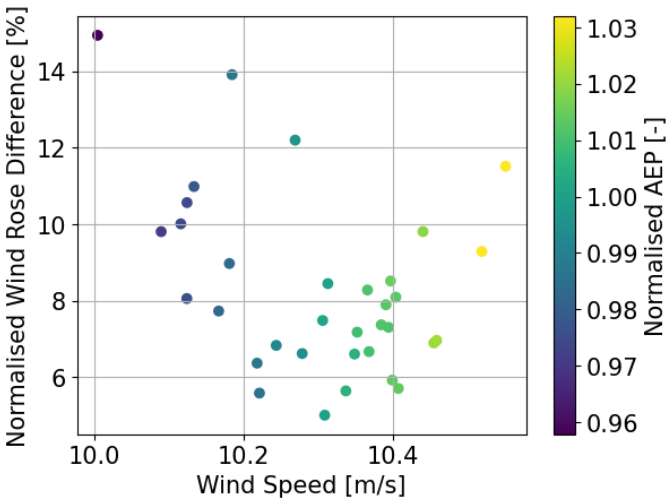

Figure 3.7: Sensitivity of AEP to variations in mean wind speed and directional frequency. Each point represents a synthetic year with fixed Weibull parameters and varied wind rose shape.

Given this, it becomes impractical to quantify IAV uncertainty in wind direction using angular offsets or even WRD metrics alone. Instead, this study adopts an approach analogous to the wind speed case: calculating the standard deviation of simulated AEPs over multiple years while fixing the wind speed distribution and varying only the directional frequencies (i.e., wind rose). Specifically, the Weibull A and k parameters for each wind direction sector are calibrated using the full period reanalysis dataset. Then, the interannual variability in direction is introduced by changing only the wind rose annually, holding all other inputs constant.

In this way, the IAV of AEP driven by wind direction is quantified directly as the standard deviation of annual AEP estimates, reflecting true yield variability rather than relying on ambiguous directional metrics. This variability is expressed as a normalised standard uncertainty in AEP due to direction-driven IAV, defined as:

$$\text{IAV} = \frac{\sigma_{\text{annual AEP}}}{\text{AEP}_{\text{overall}}} \cdot \frac{1}{\sqrt{N}} \tag{3.13}$$

where $\sigma_{\text{annual AEP}}$ is the standard deviation of simulated AEP across N years, and $\text{AEP}_{\text{overall}}$ is the mean AEP over the full period. The $\frac{1}{\sqrt{N}}$ factor accounts for the fact that the uncertainty is expressed relative to the estimate of the mean AEP. This approach allows IAV uncertainty to be incorporated consistently with other uncertainty components in a wind resource assessment framework.

4 Results

4.1 Measurement Uncertainty

As introduced in Equation (3.2), the influence of measurement uncertainty on AEP is assessed by shifting the wind rose direction by a fixed angular offset and recalculating the resulting AEP. Figure 4.1 shows the relationship between the applied directional shift and the corresponding change in AEP at *Site A* and *Site B*. The observed trend closely follows a linear relationship, indicating that the AEP change is approximately proportional to the magnitude of directional bias.

Given this linearity, the uncertainty in AEP resulting from wind direction measurement bias can be estimated as the product of the absolute slope of the fitter line and the directional measurement uncertainty, σ_θ, introduced in Section 2.2.1. For *Site A*, the slope of the fitted regression line is approximately -0.093. With a total measurement uncertainty of $\sigma_\theta = 3°$, the propagated AEP uncertainty is calculated as:

$$\sigma_{AEP} = |-0.093| * \sigma_\theta = 0.093 * 3 \approx 0.28\% \tag{4.1}$$

For *Site B*, the slope of the fitted regression line is approximately -0.230. With a total measurement uncertainty of $\sigma_\theta = 3°$, the propagated AEP uncertainty is calculated as:

$$\sigma_{AEP} = |-0.230| * \sigma_\theta = 0.230 * 3 \approx 0.69\% \tag{4.2}$$

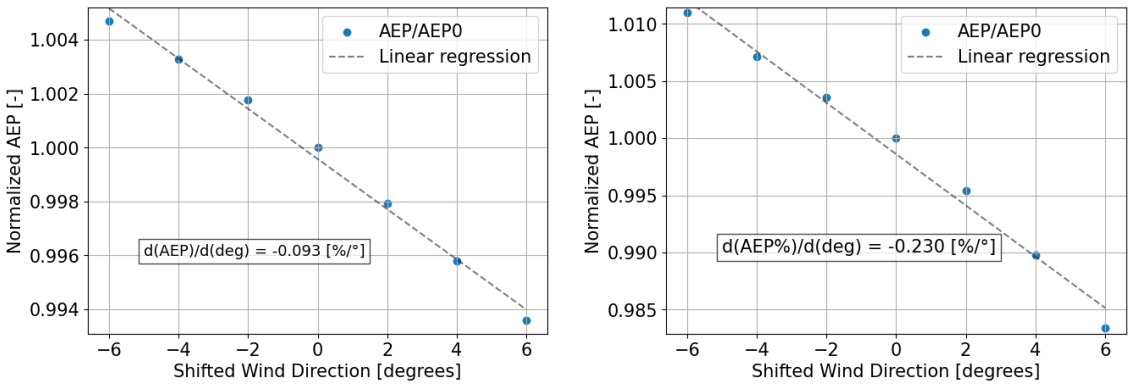

Figure 4.1: Relationship between directional shift and corresponding change in normalised AEP at *Site A* (left) and *Site B* (right).

This result quantifies the influence of wind direction measurement bias on AEP and confirms the suitability of a linear uncertainty propagation approach within the tested range of directional shifts ($\pm 6°$). Although a $3°$ uncertainty is adopted in this study, higher thresholds, up to $5°$ or even $10°$, are recognised in industry practice. For example, Foussekis [67] reports a acceptance criterion of $< 10°$ for directional offset for offshore systems during early deployment phases, and tighter tolerance of $< 5°$ after full calibration. Notably, the linear relationship between directional shift and AEP deviation remains valid even at higher uncertainty level. The linear propagation model described in Equation (3.2) remains applicable for estimating AEP uncertainty, if the directional uncertainty does not

exceed approximately $6°$. For larger uncertainties, additional simulations and updated sensitivity curves would be required to ensure validation.

4.2 LTWRD Uncertainty

Data Bias Uncertainty

The LTWRD uncertainty was quantified by evaluating the distribution of normalised AEP errors between measured and reanalysis-based wind resource assessments at the target site. As shown in Figure 4.2, the distribution of errors at *Site A*, derived from 200 Monte Carlo samples using 1-year windows over a 5.5-year concurrent period, follows a near normal shape with a mean of 0.14% and a standard deviation of 0.35%. This standard deviation is taken as *Site A*'s LTWRD uncertainty under the assumption that 1 year is the relevant time-averaging basis and that each sample captures intrannual seasonal variability.

At *Site B*, the distribution of normalised AEP errors over 1-year windows exhibits a much wider spread, with a standard deviation of 4.27%, as depicted in Figure 4.2 (right). The higher LTWRD uncertainty at *Site B* suggests that site-specific factors can amplify the challenges of correcting reanalysis data, reinforcing the need for careful site-based uncertainty evaluation. Notably, the distribution is not centred at 0% but has a mean bias of -2.61%, suggesting a systematic difference between the reanalysis and measured datasets. This may reflect a climatological mismatch, where the long-term downscale datasets does not adequately capture the directional inflow experienced during the measurement period. However, this deviation could also arise from the short duration of measured data (only 1.1 years), which might have coincided with an anomalous directional regime.

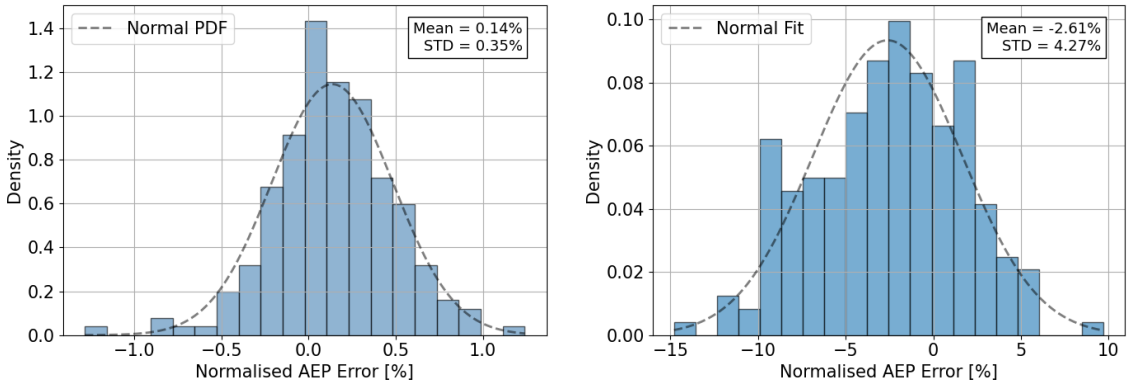

Figure 4.2: Distribution of normalised AEP error between measured and reanalysis datasets over 1-year windows at *Site A* (left) and *Site B* (right).

It is important to note that LTWRD uncertainty is sensitive to the length of the concurrent measurement period. A sensitivity test at *Site A* demonstrated that restricting the concurrent period to just the first 3 years resulted in a reduced standard deviation of 0.33%, potentially underestimating the true LTWRD uncertainty. This finding underscores the importance of using the longest possible concurrent period to ensure robust estimates. In this study, the full available concurrent datasets were used: 5.5 years for *Site A* and 1.1 years for *Site B*. However, the limited duration at *Site B* likely increases the risk of underrepresenting sample variability, which may lead to an underestimation of LTWRD uncertainty. This limitation and its implications are further discussed in the sensitivity analysis in Section 5.3.

Disjunct Sampling Uncertainty

To account for the impact of temporal resolution on AEP estimation, disjunct sampling was applied to the 10-minute measured data to simulate a 1-hour measurement regime, as introduced in Section 3.2.3. As shown in Figure 4.3 (left), the disjunct sampling uncertainty at *Site A*, which quantifies the error introduced by reduced measurement frequency, is calculated to be 0.37%. Similarly, in Figure 4.3 (right), the uncertainty introduced by disjunct sampling at *Site B* is calculated as 0.47%.

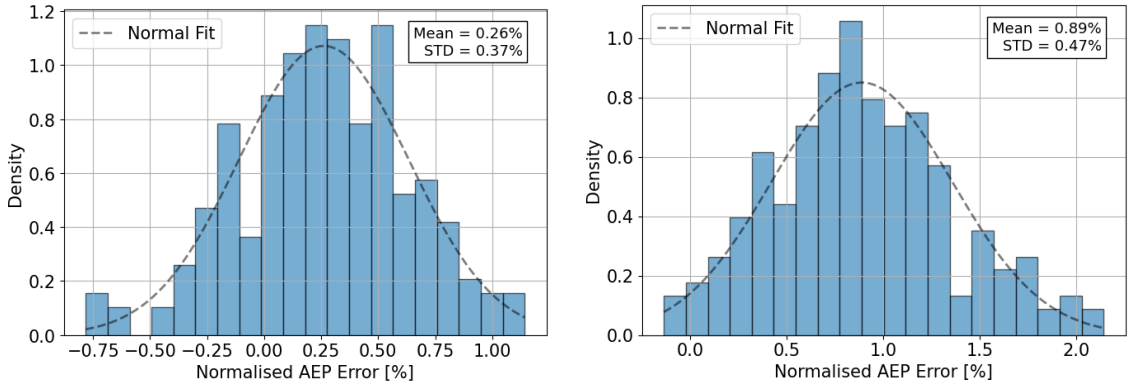

Figure 4.3: Distribution of normalised AEP error between 10-minute measured data and 1-hour disjunct-sampled datasets over 1-year windows at *Site A* (left) and *Site B* (right).

The difference in disjunct sampling uncertainty between the two sites is primarily attributable to site-specific characteristics, such as the wind rose distribution and wind farm layout. Site A, which exhibits a broader and more uniform wind direction distribution, is less sensitive to short-term fluctuations. Even if both sites experience a similar level of turbulent fluctuation, Site A is more likely to smooth out these variations, resulting in a lower propagation of uncertainty into AEP. In contrast, Site B is dominated by a narrow wind rose and has a more closely spaced, heavily waked layout. This makes it more sensitive to small directional shifts, meaning the same magnitude of wind direction fluctuation can lead to a larger deviation in wind rose sectors and a more pronounced impact on AEP. Thus, the higher disjunct sampling uncertainty at Site B is consistent with its greater directional sensitivity and layout-induced wake effects.

This finding highlights a key consideration when assessing the effect of reduced temporal resolution: sites with different directional characteristics and layouts can respond differently to the same magnitude of wind fluctuation. Disjunct sampling uncertainty should therefore be evaluated on a site-by-site basis through AEP simulations, rather than assumed to have a uniform effect across projects.

Total LTWRD Uncertainty

The total LTWRD uncertainty reflects the compounded effect of two independent sources: temporal sampling resolution and data discrepancy between measurements and reanalysis. These two effects arise from fundamentally distinct mechanisms, disjunct sampling stems from reduced temporal resolution, while LTWRD data bias uncertainty originates from discrepancies between short-term measurements and long-term reanalysis. Therefore, they are combined using the root-sum-square approach:

$$\sigma_{\text{LTWRD,total}} = \sqrt{\sigma^2_{\text{disjunct}} + \sigma^2_{\text{LTWRD, data bias}}} \tag{4.3}$$

where σ_{disjunct} is the standard deviation of normalised AEP error introduced by 1-hour disjunct sampling, and $\sigma_{\text{LTWRD, data bias}}$ is the uncertainty due to the discrepancy between measured and reanalysis directional distributions. The resulting total LTWRD uncertainties are:

$$\sigma_{\text{LTWRD,total, Site A}} = \sqrt{0.37^2 + 0.35^2} \approx 0.51\% \tag{4.4}$$

$$\sigma_{\text{LTWRD,total, Site B}} = \sqrt{0.47^2 + 4.27^2} \approx 4.30\% \tag{4.5}$$

These results underscore the dominant role of directional mismatch in the total LTWRD uncertainty, particularly for directionally sensitive sites like *Site B*.

LTWRD uncertainty arises from the assumption that the long-term directional distribution, typically derived from downscaled reanalysis data, accurately reflects the site's climatological inflow. If this assumption is violated, the resulting wind rose deviation introduces a systematic bias into all future AEP estimates.

While some LTC techniques include direction correction (e.g., mean-sectoral offsets or FFT-based interpolation) [68], such corrections are not standard in reanalysis workflows. In practice, directional biases often remain uncorrected, undermining representativeness in directional analyses. At *Site B*, the reanalysis wind rose is taken directly from the downscaling model without post-processing to align it with site-specific measurements.

Figure 4.4 compares the directional frequency distributions from measurements and reanalysis during the concurrent period. Although the two wind roses appear visually similar and the normalised wind rose difference (Equation (3.12)) is 10.3%, their energy roses differ substantially. As Langreder and Højstrup [69] emphasise, visual similarity can be misleading, since small differences in directional frequency can have a disproportionate impact on AEP. This occurs because directional frequency scales each sector's contribution to total energy production, and not all sectors contribute equally due to wake effects and layout asymmetry.

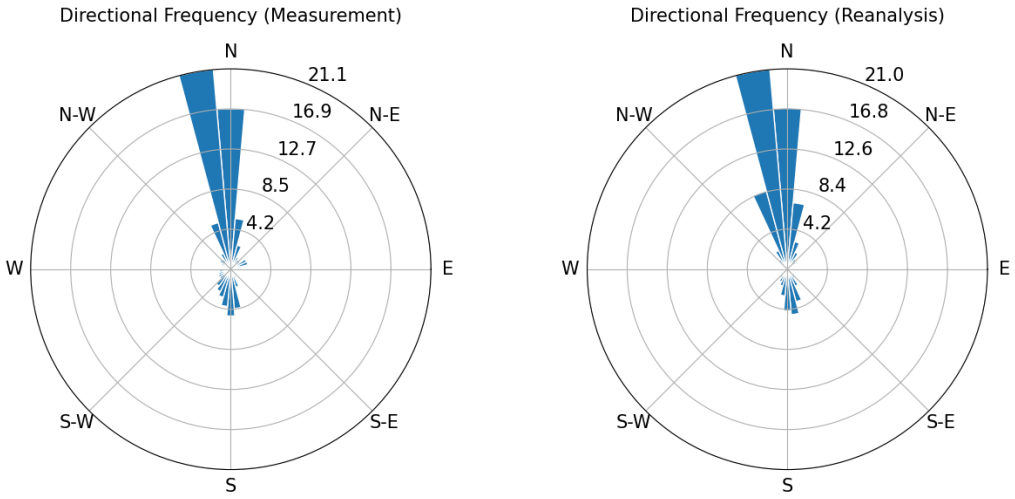

Figure 4.4: Comparison of directional frequency distributions from measurement and reanalysis data during the concurrent period at *Site B*.

Without robust correction, such directional mismatches propagate through the wake model and bias long-term AEP estimates. The *Site B* results illustrate that directional uncer-

tainty is often a structural problem in the long-term climatology. To mitigate this, correction method should be applied and should focus on aligning the energy rose, rather than merely matching wind rose frequencies, to ensure that wind direction is accurately represented in long-term projections.

4.3 IAV Uncertainty

To quantify the IAV of direction-driven AEP uncertainty, annual AEP estimates were computed from reanalysis data by varying the wind rose for each year while holding sector-wise Weibull parameters fixed. The resulting standard deviation of normalised AEP reflects the directional variability from year to year.

At *Site A*, using the full 33-year reanalysis period, the standard deviations of annual AEP was found to be 1.83% , as shown in Figure 4.5 (left), corresponding to an IAV standard uncertainty of $\text{IAV}_{\text{Site A, full}} = \frac{1.83\%}{\sqrt{33}} \approx 0.32\%$. While a longer reanalysis record is available, several studies recommend using a reference period of approximately 15-20 years for long-term representativeness [70]. When only the most recent 20 years were considered, the standard deviation slightly decreased to 1.77% , as shown in Figure 4.5 (right). This results in a higher standard uncertainty:

$$\text{IAV}_{\text{Site A, last 20}} = \frac{1.77\%}{\sqrt{20}} \approx 0.40\% \tag{4.6}$$

The standard uncertainty is higher in 20 years with lower standard deviation of annual AEP. This counter-intuitive result occurs because standard uncertainty reflects not only the variability but also the confidence in the estimate. Although the variability is slightly lower, the shorter reference period (20 years) provides fewer samples, leading to higher uncertainty in the estimated mean.

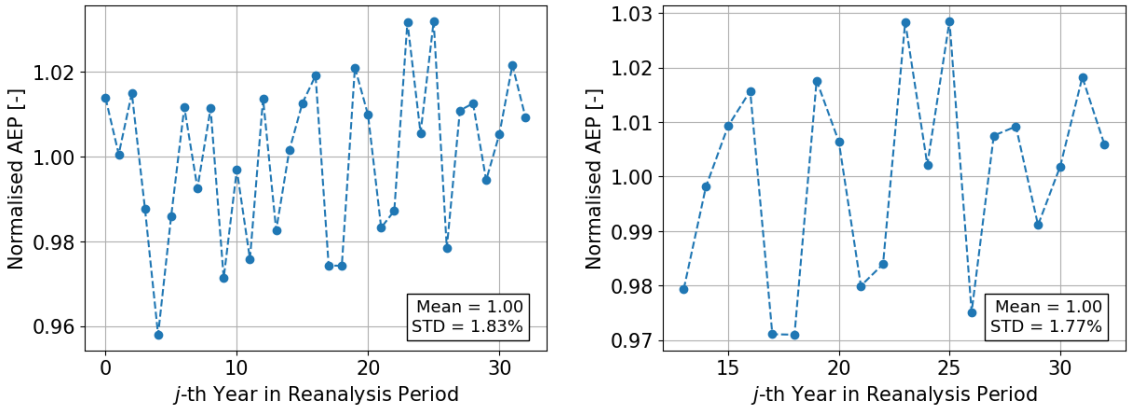

Figure 4.5: Normalised annual AEP across 33 years (left) and the most recent 20 years (right) of reanalysis data at *Site A*.

For *Site B*, which also used 20 years of reanalysis data, the observed standard deviation was higher at 2.65% (Figure 4.6), yielding a directional IAV uncertainty:

$$\text{IAV}_{\text{Site B}} = \frac{2.65\%}{\sqrt{20}} \approx 0.59\% \tag{4.7}$$

The higher IAV at *Site B* could be attributed to its more uni-directional and heavy-waked wind farm layout, making the site more sensitive to year-to-year variations in wind direction, as discussed in the previous result section.

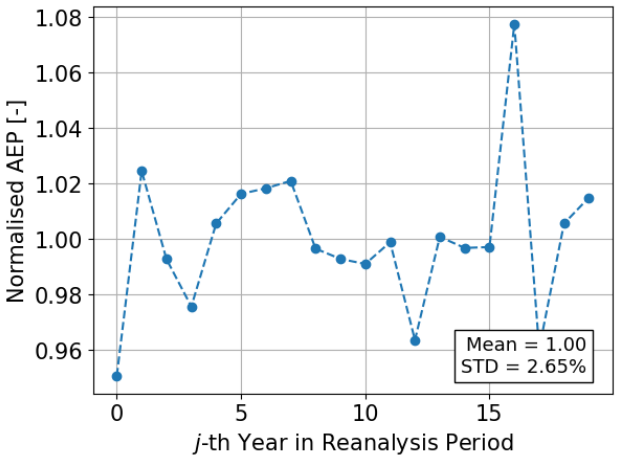

Figure 4.6: Normalised annual AEP for *Site B* over a 20-year reanalysis period.

These results confirm that the length of the reanalysis record has a direct influence on the estimated IAV uncertainty, and that regional wind direction climatology also plays a significant role in the magnitude of interannual variability.

4.4 Summary of Directional Uncertainty

The three main components of directional uncertainty were evaluated separately for *Site A* and *Site B* using site-specific measured data and reanalysis data. The resulting standard uncertainties, expressed as percentages of AEP, are summarised in Table 4.1.

Table 4.1: Comparison of directional uncertainty components on AEP at *Site A* and *Site B*.

Site	Measurement [%]	LTWRD [%]	IAV [%]	Total Uncertainty [%]
Site A	0.28	0.51	0.40	0.71
Site B	0.69	4.30	0.59	4.39

The total directional uncertainty, calculated via the root-sum-square (RSS) of the three contributions, reflects the combined impact of all directional effects on AEP. At *Site A*, the total uncertainty is moderate (0.71%), with each component contributing on a similar scale. In contrast, *Site B* shows a much larger total uncertainty (4.39%), almost entirely driven by the LTWRD component (4.30%).

The relative importance of each component can be expressed as its fractional variance contribution:

$$\text{Contribution of } \sigma_{\text{measurement}} = \frac{\sigma^2_{\text{measurement}}}{\sigma^2_{\text{measurement}} + \sigma^2_{\text{LTWRD}} + \sigma^2_{\text{IAV}}} \times 100 \tag{4.8}$$

By this measure, measurement uncertainty accounts for about 16% of the total variance at *Site A*, but only 2.5% at *Site B*, where LTWRD overwhelmingly dominates. This indicates that while improving measurement accuracy is beneficial, the representativeness and accuracy of long-term directional reference data could exert a far greater influence on total directional uncertainty.

For context, comparable studies of typical wind projects report total AEP wind data uncertainties, considering only measurement uncertainty and long-term wind resource variability, of approximately 2.1–3.7% for array-averaged wind speed [71]. These values refer to wind speed uncertainty, which generally has a stronger effect on AEP than directional uncertainty. The 0.71% observed at *Site A* is therefore well below these typical ranges, whereas the 4.39% at *Site B* is unusually high, even compared to typical wind speed uncertainty. This suggests that neglecting the influence of wind direction on AEP can lead to a significant underestimation of total uncertainty.

Overall, the results demonstrate that directional uncertainty is highly site-specific. Robust AEP assessments therefore require site-by-site evaluation of wind rose uncertainty, with particular attention to layouts with strong directional sensitivity or sites with highly unidirectional inflow.

5 Discussion

5.1 Significance of Directional Uncertainty in the Context of Total AEP Uncertainty

In the wind industry, reported total AEP uncertainties vary depending on site conditions, data availability, and the specific uncertainty categories included. For example, Lee and Fields [6] report total AEP uncertainties ranging from 5% to 15% (average 10%), Mortensen and Ejsing Jørgensen [72] report an average of 10% with a range of 6%–21%, and the Wind Resource Assessment Handbook [71] reports a range of 4.1%–7.5%. While the numbers vary, typical projects generally fall in the range of 8–10%, with very low-uncertainty projects at 4–7% and high-uncertainty projects exceeding 10%.

It is important to note that these totals usually include several uncertainty categories that are excluded from this study, such as future wind resource changes, wind shear and veer (model extrapolation), wake modelling, and plant/turbine performance. Within this context, the total directional wind data uncertainty at *Site B* (4.39%) is striking: it already represents a substantial fraction of the total uncertainty budget for a typical project, even without accounting for other categories. Moreover, the LTWRD component alone at *Site B* (4.27%) is of a similar magnitude to the entire AEP uncertainty for low-uncertainty projects.

This finding underscores that in certain cases (particularly for sites with highly uni-directional inflow and strong wake sensitivity, as observed at *Site B*), the uncertainty contribution from wind direction alone can approach or even match the total AEP uncertainty normally expected in project assessments. Neglecting this factor risks significant underestimation of project risk. Moreover, the magnitude of the LTWRD term at *Site B* suggests that improving long-term wind rose representativeness should be treated as a priority in the uncertainty mitigation strategy.

5.2 Mast Alignment Deviation

In practice, the uncertainty associated with mast alignment deviation, as discussed in Section 2.2.1, is likely to be significantly greater than the conservative estimate of 2° used in this study. Various sources can contribute to this deviation, including magnetic declination, human-induced errors, and general inaccuracies in compass readings. Common mistakes such as misinterpreting directions, improper handling of the compass, and a lack of understanding of how the compass operates can all introduce additional alignment errors.

Due to the absence of a reliable method to quantitatively assess human error, this study adopts a conservative estimate of 2° to represent mast alignment deviation. However, it is important to acknowledge that the actual value could be significantly higher, which would result in greater measurement uncertainty.

The results presented in Section 4.1 offer a useful basis for understanding the relationship between measurement uncertainty and AEP uncertainty. As illustrated in Figure 4.1, the relationship between wind direction bias and AEP remains approximately linear for deviations up to ±6°. This indicates that the mapping function derived here remains valid even if the actual mast alignment error exceeds 2°, as long as the total measurement uncertainty remains within the ±6° range. For measurement uncertainties beyond this range, additional simulations would be required to establish the relationship.

5.3 Sensitivity of LTWRD to Concurrent Period Length

When applying the Moving Block Bootstrap (MBB) to estimate LTWRD uncertainty, the length of the concurrent period between measured and reanalysis data directly affects the variability of the resulting AEP error distribution. A shorter concurrent period provides fewer residuals for resampling, which limits the diversity of bootstrapped time series. Consequently, shorter periods often yield lower standard deviations, leading to an underestimation of LTWRD uncertainty.

This effect was evident in the *Site A* case study, where reducing the concurrent period from 5.5 years to the first 3 years decreased the standard deviation of the normalised AEP error from 0.37% to 0.35% (Figure 5.1). Although this change appears small, it demonstrates the sensitivity of LTWRD uncertainty to concurrent period length and suggests a risk of underprediction if the available period is too short. This reinforces the importance of using the longest possible concurrent period to better capture the “true” variability.

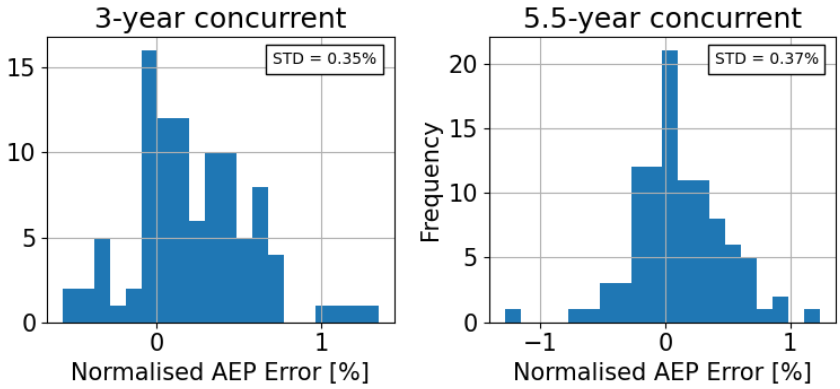

Figure 5.1: Effect of concurrent period length on LTWRD uncertainty (*Site A*).

Due to the limited measurement period at *Site A*, it is not possible to explore this sensitivity beyond 5.5 years. To investigate the relationship more thoroughly, the same method was applied to the Lindenberg datasets in the Tall Tower dataset [73], which offers a longer time series of both measured data and reanalysis data. In this test, the site-specific turbine and layout for Lindenberg were not used, as the objective was not to produce a locally representative AEP value, but rather to examine the general trend. Specifically, the aim was to demonstrate that the standard deviation increases with concurrent period length before converging toward a stable value, regardless of the absolute magnitude of the results.

The results in Figure 5.2 confirm the hypothesised trend: as the concurrent period length increases, the residual pool available for MBB becomes more diverse, which increases the variability in bootstrapped AEP estimates and thus the estimated LTWRD uncertainty. The curve show a gradual rise in standard deviation up to around nine years, beyond which the values appear to plateau, suggesting a convergence point where adding more years of data has minimal impact on the results.

Although the absolute change may appear small, the standard deviation increases from 0.89% for a 3-year period to 1.08% for a 9+ year period (a difference of 0.19 percentage points), this change could be impactful when placed in context. For instance, measurement uncertainty at *Site A* is only 0.28%, indicating that concurrent period length alone could contribute a meaningful share of the total uncertainty budget.

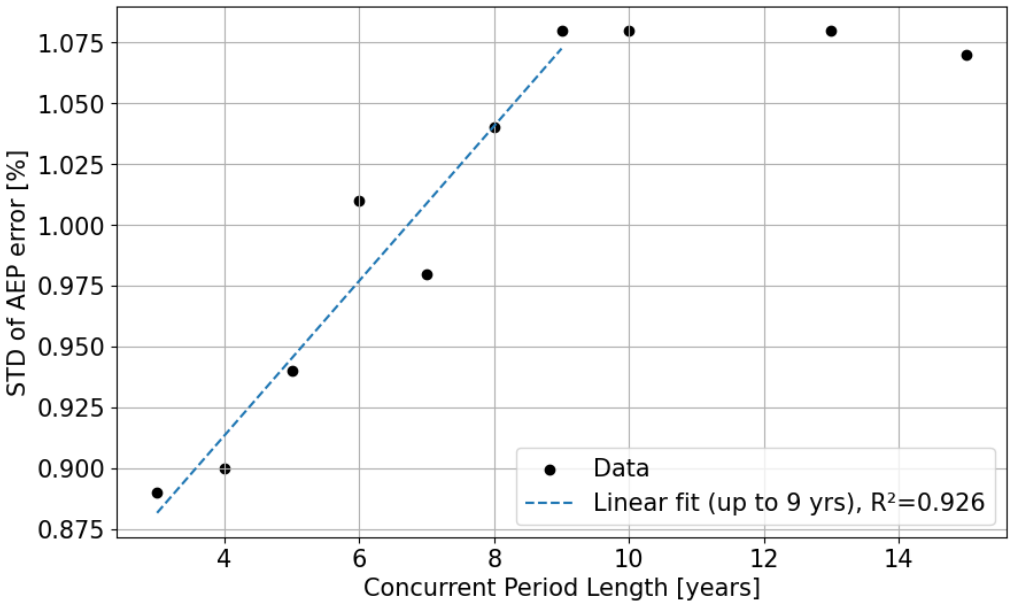

Figure 5.2: Sensitivity of LTWRD uncertainty to concurrent period length based on the Tall Tower dataset [73].

It should be noted that these results are derived from a single dataset. The observed convergence and magnitude of sensitivity may vary depending on site characteristics and data quality. Further studies using multiple sites are recommended before generalising these findings.

5.4 Wake Model Uncertainty

The following sections investigate two main sources of wake model uncertainty. First, in Section 5.4.1, a model comparison is performed using three widely adopted wake models to evaluate differences in AEP predictions and directional sensitivity across sites. Then, in Section 5.4.2, a parameter sensitivity analysis is conducted on *TurboGaussianDeficit* to assess how variations in internal model parameters (such as the turbulence intensity coefficients c_1, c_2, and the wake expansion factor A) influence AEP estimates under directional shifts.

5.4.1 Wake Model Comparison

This section discusses the uncertainty introduced by the selection of wake models and their parameter settings. As discussed previously, the baseline model employed in this study is the adjusted version of Nygaard et al (2022)[42], implemented in PyWake as *TurboGaussianDeficit*, commonly associated with the industry-accepted wake model ”TurbOPark”. While the Nygaard 2022 model has seen widespread adoption for offshore wind energy applications, alternative models remain in use, each with its own treatment of wake recovery, turbulence propagation, and blockage effects.

To explore the directional sensitivity associated with different wake modelling approaches, a comparison was made using three representative wake deficit models in PyWake [74]. These models differ primarily in how they describe the velocity deficit profile behind turbines and the treatment of wake recovery.

- **TurboNOJDeficit (Nygaard 2020)** is a variant of the traditional NOJ (Nielsen-Olesen-Jensen) engineering model, where the wake is represented as a uniform ”top-hat”

profile with a sharp edge. It assumes a fixed wake expansion rate influenced by local turbulence intensity but does not account for the gradual radial decay of the deficit.

- **TurboGaussianDeficit (Nygaard 2022)** improves upon the NOJ formulation by replacing the top-hat deficit with a Gaussian distribution. This allows for a more realistic radial wake profile that tapers smoothly from the wake center to the free stream, improving the physical fidelity of wake interactions. It is the wake model used in PyWake's implementation of TurbOPark.
- **FugaDeficit** is a physics-based model derived from a linearised Reynolds-averaged Navier-Stokes (RANS) equations, solved in the frequency domain. In this study, the model is applied under the assumption of homogeneous terrain and neutral atmospheric stability, with wake deficits computed in spectral space. This allows FUGA to capture linear wake interactions and turbulence propagation more rigorously than empirical models. The model setup uses a roughness length of 0.0001 m and a boundary layer height of 400 m. These parameters align with the recommendations for neutral conditions of Ott et al. [47] and the course material on ideal atmospheric boundary layer conditions in the Baltic Sea region [75]. Unlike empirical wake models, FUGA avoids pre-defined deficit shapes and instead computes the wake dynamically from turbine thrust and ambient conditions.

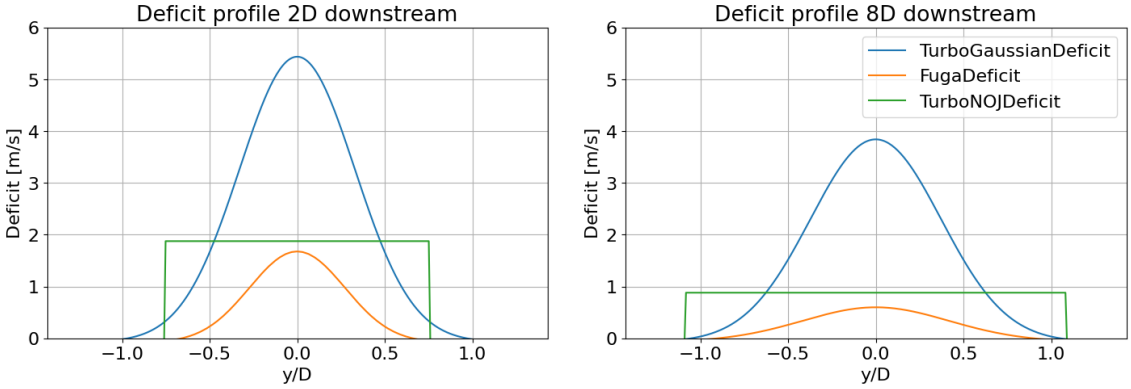

Figure 5.3: Comparison of downstream wake deficit at *Site A*.

As shown in Table 5.1 and Table 5.2, AEP was calculated using all three models with directional shifts of $+6^\circ$ and -6°, and compared against the unshifted baseline. The changes in AEP are reported as percentages normalised by the original AEP and shown in the last column of each table. All models produce different absolute AEP values and different sensitivities to wind direction shifts. *Site B* shows high sensitivity to $\pm 6^\circ$ directional shifts across all models, particularly with the *TurboGaussianDeficit*, which exhibits a 4% spread in AEP. This sensitivity is likely due to the site's narrow wind rose and heavily waked layout. In contrast, *Site A* shows much lower directional sensitivity (within $\pm 1\%$), consistent with its broader wind rose and more symmetric, widely spread layout.

Across both sites, the directional sensitivity follows the same ranking: *TurboGaussianDeficit* yields the largest AEP deviation, followed by *TurboNOJDeficit*, and then *FugaDeficit*. This trend may stem from differences in how wake centre-line deficits and wake expansion are modelled, see Figure 5.3. The *TurboNOJDeficit* model uses a top-hat wake profile, while *TurboGaussianDeficit* employs a Gaussian wake profile that concentrates

deficits more sharply. Both models use fixed-direction averaging, which can underestimate wake width and overestimate velocity deficits, thereby increasing directional sensitivity. In contrast, FUGA applies a spectral-domain approach with simple and Gaussian directional averaging methods [47], which effectively filter out high-frequency directional noise. This leads to reduced wake deficits and flatter wake profiles, inherently smoothing the response to directional shifts. As a result, FUGA tends to show lower sensitivity to small angular changes in inflow direction. While this may lead to underestimation of directional effects, engineering models like *TurboGaussianDeficit* may conservatively overestimate the impact.

Table 5.1: Directional sensitivity of AEP under different wake models at *Site A*.

Wake Model	AEP (Original) [GWh]	AEP (+6°) [GWh]	AEP (–6°) [GWh]	$\lvert\Delta AEP_{+6^\circ - -6^\circ}\rvert$ [%]
TurboNOJDeficit	1.6813e+06	1.6805e+06 (–0.050%)	1.6819e+06 (+0.038%)	0.088
TurboGaussianDeficit	1.4910e+06	1.4892e+06 (–0.122%)	1.4913e+06 (+0.018%)	0.140
FugaDeficit	1.5797e+06	1.5794e+06 (–0.015%)	1.5799e+06 (+0.018%)	0.033

Table 5.2: Directional sensitivity of AEP under different wake models at *Site B*.

Wake Model	AEP (Original) [GWh]	AEP (+6°) [GWh]	AEP (–6°) [GWh]	$\lvert\Delta AEP_{+6^\circ - -6^\circ}\rvert$ [%]
TurboNOJDeficit	6314.1	6269.6 (–0.705%)	6371.2 (+0.905%)	1.610
TurboGaussianDeficit	5739.8	5597.5 (–2.479%)	5823.9 (+1.464%)	3.943
FugaDeficit	6496.2	6445.3 (–0.784%)	6538.6 (+0.653%)	1.437

5.4.2 Parameter Sensitivity of TurboGaussianDeficit

In addition to model-to-model variation, internal parameter settings of the *TurboGaussianDeficit* model were also tested (e.g., c_1, c_2, wake expansion factor A). TurbOPark models the wake deficit using a Gaussian-shaped velocity deficit profile derived from the Bastankhah and Porté-Agel formulation [76], with wake expansion modelled as a function of turbulence intensity and thrust coefficient.

The characteristic wake width $\sigma_{w,i}(x_i)$ as a function of downstream distance x_i is given by [42]:

$$\frac{\sigma_{w,i}(x_i)}{D_i} = \epsilon_i + \frac{AI_0}{\beta}\left(\sqrt{(\alpha+\beta x_i/D_i)^2+1} - \sqrt{1+\alpha^2} - \ln\left[\frac{(\sqrt{(\alpha+\beta x_i/D_i)^2+1}+1)\alpha}{(\sqrt{1+\alpha^2}+1)(\alpha+\beta x_i/D_i)}\right]\right) \tag{5.1}$$

In this expression, ϵ_i represents the initial wake width, and D_i is the rotor diameter. The constant A is a wake expansion calibration parameter, with a recommended default value of 0.04 based on validation from Ørsted's 19 offshore wind farm [42]. I_0 is the ambient turbulence intensity, while $\alpha = c_1 I_0$ and $\beta = c_2 I_0/\sqrt{C_T}$, where c_1 and c_2 are model constants typically set to 1.5 and 0.8, respectively, as proposed in the Frandsen turbulence model [44].

To assess the model's sensitivity, a parametric sweep was conducted by varying $c1$, $c2$ and A by $\pm 33\%$ around their default values. The AEP response was then evaluated under angular wind direction shifts of $\pm 6^\circ$.

The results, shown in Figure 5.4 and Figure 5.5, highlight the non-negligible influence of model parameters on directional sensitivity. The black cross in each figure indicates the default parameter values ($c_1 = 1.5$, $c_2 = 0.8$, and $A = 0.04$). For *Site B*, depending on the parameter combination, AEP difference ($\Delta AEP_{+6^\circ - -6^\circ}$) ranged from 3.7% to

4.2%, resulting in a 0.5% spread. Although smaller than the variation observed between different wake model formulations, this spread could still be significant, when comparing uncertainties across multiple sources.

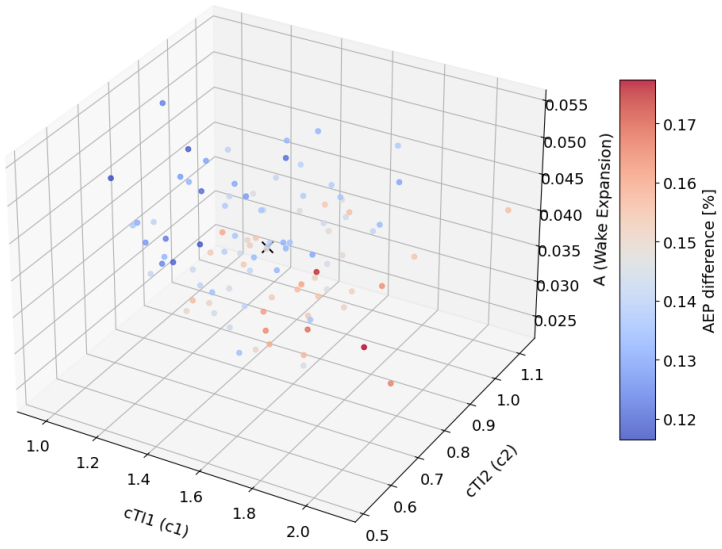

Figure 5.4: Sensitivity of AEP to TurbOPark parameter variations under $\pm 6°$ wind direction shift at *Site A*.

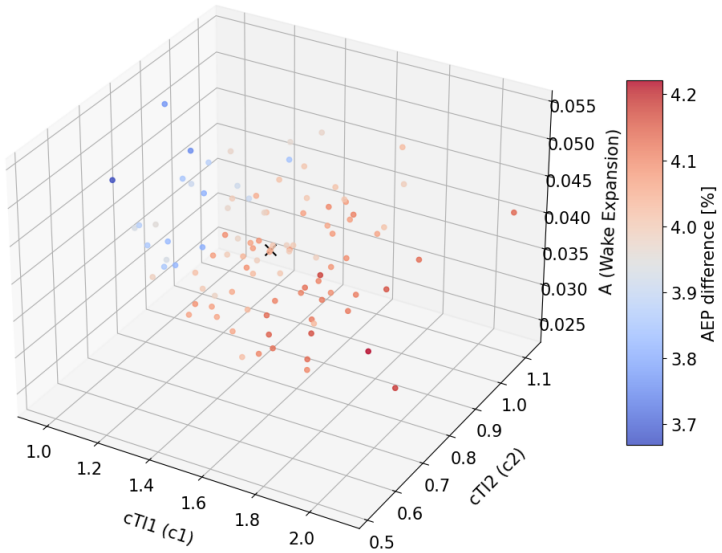

Figure 5.5: Sensitivity of AEP to TurbOPark parameter variations under $\pm 6°$ wind direction shift at *Site B*.

Given that model validation is beyond the scope of this study, the TurbOPark model with its default parameter set was used. This default configuration has been previously validated by Ørsted on multiple offshore wind farms and serves as a practical and industry-accepted

baseline. However, it is important to note that both the model and its parameters remain tunable. Future work focusing on model validation and site-specific calibration will be necessary to improve modelling accuracy and ensure more reliable AEP predictions under directional uncertainty.

6 Conclusion

This thesis set out to address a gap in wind resource assessment: the limited treatment of wind direction uncertainty in AEP estimation. By isolating directional uncertainty from wind data, the analysis decomposed it into three independent components: measurement uncertainty, long-term wind rose difference (LTWRD), and interannual variability (IAV). Two offshore sites with contrasting layouts, inflow regimes, and measurement records were studied to test the robustness of the methodology and explore site-specific behaviour.

The results show that the composition of total directional uncertainty is highly site dependent. At the more symmetric, broadly distributed inflow site (Site A), total directional uncertainty was modest (0.71% of AEP), with contributions from all three components on a similar scale. In contrast, the directionally sensitive site with a more compact, asymmetric layout (Site B) exhibited a much higher total directional uncertainty of 4.39%, driven almost entirely by LTWRD (4.27%). This magnitude is comparable to, or even exceeds, the total uncertainty reported for low-uncertainty projects (typically range from 4% to 7%), emphasising the potential impact of uncorrected directional bias.

From these findings, three key conclusions could be drawn:

1. No universal composition: The relative contribution of each directional uncertainty component varies significantly across sites; general assumptions of composition are unreliable.
2. LTWRD can dominate: In directionally sensitive sites, long-term directional bias can be the single largest uncertainty source, even surpassing wind speed uncertainty in relative importance.
3. Correction must be energy-based: Aligning long-term and short-term wind direction by energy rose, rather than wind rose frequency, offers a more meaningful correction for AEP estimation.

In answering the research questions, two main insights emerge. First, measurement uncertainty can be improved or reduced with higher-accuracy instruments, but it is rarely the dominant source of directional uncertainty. LTWRD generally plays a far greater role, particularly at directionally sensitive sites. Second, none of the proposed methods for quantifying directional uncertainty proved to be a reliable standalone predictor of AEP variation. The Yamartino method fails to represent actual directional variation in multimodal or broad distributions, leading to underestimation of its impact on AEP; while wind rose deviation reflects percentage differences in frequency without considering shape and shows no direct correlation to AEP changes.

For industry practice, this work highlights that while improving measurement accuracy can reduce the measurement component, LTWRD is often the dominant source of directional uncertainty. Therefore, it is essential to: (1) explicitly include wind direction in long-term correction frameworks, (2) conduct site-by-site directional uncertainty assessments, and (3) consider directional uncertainty alongside wind speed in the total uncertainty budget to avoid underestimating project risk.

Future research could expand this framework to include spatial extrapolation effects, long-term directional trends under climate change, increased measurement uncertainty from installation errors, and detailed wake model sensitivity studies. By expanding the scope

to these areas, a more comprehensive and realistic representation of directional effects in AEP uncertainty can be achieved, improving the reliability of wind project assessments.

Bibliography

[1] WindEurope. *Wind Energy in Europe – 2024 Statistics and the Outlook for 2025-2030*. https://windeurope.org/intelligence-platform/product/wind-energy-in-europe-2024-statistics-and-the-outlook-for-2025-2030/. Accessed: 13 August 2025. 2024.

[2] Germán Martínez Montes and Enrique Prados Martín. "Profitability of wind energy: Short-term risk factors and possible improvements". In: *Renewable and Sustainable Energy Reviews* 11.9 (2007), pp. 2191–2200. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2006.03.009. URL: https://www.sciencedirect.com/science/article/pii/S1364032106000578.

[3] EMD International A/S. *Loss and Uncertainty in WindPRO 2.9*. https://help.emd.dk/knowledgebase/content/WindPRO2.9/12-UK_WindPRO2.9_LOSSUNCERTAINTY.pdf. 2010.

[4] Kerry S. Klemmer, Emily P. Condon, and Michael F. Howland. "Evaluation of wind resource uncertainty on energy production estimates for offshore wind farms". In: *Journal of Renewable and Sustainable Energy* 16.1 (Jan. 2024), p. 013302. ISSN: 1941-7012. DOI: 10.1063/5.0166830. eprint: https://pubs.aip.org/aip/jrse/article-pdf/doi/10.1063/5.0166830/18283884/013302_1_5.0166830.pdf. URL: https://doi.org/10.1063/5.0166830.

[5] N. Bodini and M. Optis. "Operational-based annual energy production uncertainty: are its components actually uncorrelated?" In: *Wind Energy Science* 5.4 (2020), pp. 1435–1448. DOI: 10.5194/wes-5-1435-2020. URL: https://wes.copernicus.org/articles/5/1435/2020/.

[6] J. C. Y. Lee and M. J. Fields. "An overview of wind-energy-production prediction bias, losses, and uncertainties". In: *Wind Energy Science* 6.2 (2021), pp. 311–365. DOI: 10.5194/wes-6-311-2021. URL: https://wes.copernicus.org/articles/6/311/2021/.

[7] Duo Fang, Mingzhe Zou, and Sasa Djokic. "Probabilistic OPF incorporating uncertainties in wind power outputs and line thermal ratings". In: *2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*. IEEE. 2018, pp. 1–6.

[8] S.Z. Djokic et al. "Evaluation of wind turbine power outputs with and without uncertainties in input wind speed and wind direction data". In: *IET Renewable Power Generation* 14.15 (2020), pp. 2801–2809. DOI: https://doi.org/10.1049/iet-rpg.2020.0113. eprint: https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/iet-rpg.2020.0113. URL: https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-rpg.2020.0113.

[9] Anna von Brandis et al. "An investigation of spatial wind direction variability and its consideration in engineering models". In: *Wind Energy Science* 8 (Apr. 2023), pp. 589–606. DOI: 10.5194/wes-8-589-2023.

[10] M. Dörenkämper et al. "The Making of the New European Wind Atlas – Part 2: Production and evaluation". In: *Geoscientific Model Development* 13.10 (2020), pp. 5079–5102. DOI: 10.5194/gmd-13-5079-2020. URL: https://gmd.copernicus.org/articles/13/5079/2020/.

[11] Fernando Porté-Agel, Yu-Ting Wu, and Chang-Hung Chen. "A Numerical Study of the Effects of Wind Direction on Turbine Wakes and Power Losses in a Large Wind Farm". In: *Energies* 6.10 (2013), pp. 5297–5313. ISSN: 1996-1073. DOI: 10.3390/en6105297. URL: https://www.mdpi.com/1996-1073/6/10/5297.

[12] M. Gaumond et al. “Evaluation of the wind direction uncertainty and its impact on wake modeling at the Horns Rev offshore wind farm”. In: *Wind Energy* 17.8 (2014), pp. 1169–1178. DOI: https://doi.org/10.1002/we.1625. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.1625. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1625.

[13] Julian Quick et al. “Optimization under uncertainty for wake steering strategies”. In: *Journal of physics: Conference series*. Vol. 854. 1. IOP Publishing. 2017, p. 012036.

[14] M. Paul van der Laan et al. “The k-□-fP model applied to wind farms”. In: *Wind Energy* 18.12 (2015), pp. 2065–2084. DOI: https://doi.org/10.1002/we.1804. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.1804. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1804.

[15] A. Rott et al. “Robust active wake control in consideration of wind direction variability and uncertainty”. In: *Wind Energy Science* 3.2 (2018), pp. 869–882. DOI: 10.5194/wes-3-869-2018. URL: https://wes.copernicus.org/articles/3/869/2018/.

[16] O.M. Essenwanger. *Elements of Statistical Analysis*. General climatology. Elsevier, 1986. ISBN: 9780444424266. URL: https://books.google.dk/books?id=P6Y4vgAACAAJ.

[17] Nicholas Fisher and Toby Lewis. “Estimating the Common Mean Direction of Several Circular or Spherical Distributions with Differing Dispersions”. In: *Biometrika* 71 (Dec. 1984). DOI: 10.2307/2336584.

[18] R. J. Yamartino. “A Comparison of Several “Single-Pass” Estimators of the Standard Deviation of Wind Direction”. In: *Journal of Applied Meteorology and Climatology* 23.9 (1984), pp. 1362–1366. DOI: 10.1175/1520-0450(1984)023<1362:ACOSPE>2.0.CO;2. URL: https://journals.ametsoc.org/view/journals/apme/23/9/1520-0450_1984_023_1362_acospe_2_0_co_2.xml.

[19] D. Bruce Turner. “Comparison of Three Methods for Calculating the Standard Deviation of the Wind Direction”. In: *Journal of Applied Meteorology and Climatology* 25.5 (1986), pp. 703–707. DOI: 10.1175/1520-0450(1986)025<0703:COTMFC>2.0.CO;2. URL: https://journals.ametsoc.org/view/journals/apme/25/5/1520-0450_1986_025_0703_cotmfc_2_0_co_2.xml.

[20] Rudolf O. Weber. “Estimators for the Standard Deviation of Horizontal Wind Direction”. In: *Journal of Applied Meteorology* 36.10 (1997), pp. 1403–1415. DOI: 10.1175/1520-0450(1997)036<1403:EFTSDO>2.0.CO;2. URL: https://journals.ametsoc.org/view/journals/apme/36/10/1520-0450_1997_036_1403_eftsdo_2.0.co_2.xml.

[21] Nicholas I. Fisher. “Comment on ‘A Method for Estimating the Standard Deviation of Wind Directions’.” In: *Journal of Applied Meteorology* 22.11 (Nov. 1983), pp. 1971–1971. DOI: 10.1175/1520-0450(1983)022<1971:COMFET>2.0.CO;2.

[22] Sónia Liléo. *Wind Resource Assessment: How to Measure Representativeness?* Presentation, EWEA Workshop, Dublin, June 25, 2013. KJELLER VINDTEKNIKK, 2013. URL: https://www.ewea.org/events/workshops/wp-content/uploads/2013/06/EWEA-RA2013-Dublin-1-4-Sonia-Lileo-Kjeller-Vindteknikk.pdf.

[23] Scott Applequist. “Wind Rose Bias Correction”. In: *Journal of Applied Meteorology and Climatology* 51.7 (2012), pp. 1305–1309. DOI: 10.1175/JAMC-D-11-0193.1. URL: https://journals.ametsoc.org/view/journals/apme/51/7/jamc-d-11-0193.1.xml.

[24] James G Droppo and Bruce A Napier. “Wind direction bias in generating wind roses and conducting sector-based air dispersion modeling”. In: *Journal of the Air & Waste Management Association* 58.7 (2008), pp. 913–918.

[25] H. Hersbach et al. *ERA5 hourly data on single levels from 1940 to present*. 2023. DOI: 10.24381/cds.adbb2d47. URL: https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels.

[26] Hans Hersbach et al. “The ERA5 global reanalysis”. In: *Quarterly Journal of the Royal Meteorological Society* 146.730 (2020), pp. 1999–2049. DOI: https://doi.org/10.1002/qj.3803. eprint: https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.3803. URL: https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803.

[27] James M. Wilczak et al. “Evaluation and Bias Correction of the ERA5 Reanalysis over the United States for Wind and Solar Energy Applications”. In: *Energies* 17.7 (2024). ISSN: 1996-1073. DOI: 10.3390/en17071667. URL: https://www.mdpi.com/1996-1073/17/7/1667.

[28] Jon Olauson. “ERA5: The new champion of wind power modelling?” In: *Renewable Energy* 126 (2018), pp. 322–331. ISSN: 0960-1481. DOI: https://doi.org/10.1016/j.renene.2018.03.056. URL: https://www.sciencedirect.com/science/article/pii/S0960148118303677.

[29] Ed Sharp et al. “Evaluating the accuracy of CFSR reanalysis hourly wind speed forecasts for the UK, using in situ measurements and geographical information”. In: *Renewable Energy* 77 (2015), pp. 527–538. ISSN: 0960-1481. DOI: https://doi.org/10.1016/j.renene.2014.12.025. URL: https://www.sciencedirect.com/science/article/pii/S0960148114008520.

[30] Thomas Duc et al. “Local turbulence parameterization improves the Jensen wake model and its implementation for power optimization of an operating wind farm”. In: *Wind Energy Science* 4 (May 2019), pp. 287–302. DOI: 10.5194/wes-4-287-2019.

[31] Eunkuk Son et al. “Characteristics of turbine spacing in a wind farm using an optimal design process”. In: *Renewable Energy* 65 (2014). SI:AFORE 2012, pp. 245–249. ISSN: 0960-1481. DOI: https://doi.org/10.1016/j.renene.2013.09.022. URL: https://www.sciencedirect.com/science/article/pii/S0960148113004898.

[32] Johan Meyers and Charles Meneveau. “Optimal turbine spacing in fully developed wind farm boundary layers”. In: *Wind Energy* 15.2 (2012), pp. 305–317. DOI: https://doi.org/10.1002/we.469. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.469. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/we.469.

[33] Sandra Schwegmann et al. “Enabling Virtual Met Masts for wind energy applications through machine learning-methods”. In: *Energy and AI* 11 (Oct. 2022), p. 100209. DOI: 10.1016/j.egyai.2022.100209.

[34] Alex Clerc et al. “A systematic method for quantifying wind flow modelling uncertainty in wind resource assessment”. In: *Journal of Wind Engineering and Industrial Aerodynamics* 111 (2012), pp. 85–94. ISSN: 0167-6105. DOI: https://doi.org/10.1016/j.jweia.2012.08.006. URL: https://www.sciencedirect.com/science/article/pii/S0167610512002358.

[35] Ib Troen et al. “Complex terrain wind resource estimation with the wind-atlas method: Prediction errors using linearized and nonlinear CFD micro-scale models”. In: *European wind energy conference & exhibition 2014*. European Wind Energy Association (EWEA). 2014.

[36] Mark Kelly et al. “Uncertainty in vertical extrapolation of measured wind speed via shear”. In: (2019).

[37] U.S. Department of Energy. *Wind Resource Assessment Handbook: Fundamentals for Conducting a Successful Monitoring Program*. Tech. rep. Golden, CO: Office of Energy Efficiency and Renewable Energy, 1997. URL: https://www.inscc.utah.edu/~krueger/5270/wind_resource_handbook.pdf.

[38] N Bodini, J K Lundquist, and A Kirincich. “Offshore Wind Turbines Will Encounter Very Low Atmospheric Turbulence”. In: *Journal of Physics: Conference Series* 1452.1 (Jan. 2020). Chapter 3.2: Wind Veer, p. 012023. DOI: 10.1088/1742-6596/1452/1/012023. URL: https://dx.doi.org/10.1088/1742-6596/1452/1/012023.

[39] Mads M. Pedersen et al. *PyWake 2.5.0: An open-source wind farm simulation tool*. Tech. rep. DTU Wind, Technical University of Denmark, 2023. URL: https://gitlab.windenergy.dtu.dk/TOPFARM/PyWake.

[40] A. Nabil and A. Nofal. "Estimation of Wind Speed Distribution Using Maximum Likelihood Method". In: *Mathematics and Statistics* 10.2 (2022), pp. 202–208. DOI: 10.13189/ms.2022.100201.

[41] DTU Wind Energy. *PyWake Overview Notebook*. https://topfarm.pages.windenergy.dtu.dk/PyWake/notebooks/Overview.html. Accessed: 2025-07-03. 2018.

[42] J. G. Pedersen et al. "Turbulence Optimized Park model with Gaussian wake profile". In: *Journal of Physics: Conference Series*. Vol. 2265. 2. IOP Publishing, 2022, p. 022063. DOI: 10.1088/1742-6596/2265/2/022063.

[43] Nicolai Gayle Nygaard et al. "Modelling cluster wakes and wind farm blockage". In: *Journal of Physics: Conference Series* 1618.6 (Sept. 2020), p. 062072. DOI: 10.1088/1742-6596/1618/6/062072. URL: https://dx.doi.org/10.1088/1742-6596/1618/6/062072.

[44] Sten Tronæs Frandsen. "Turbulence and turbulence-generated structural loading in wind turbine clusters". English. Risø-R-1188(EN). PhD thesis. DTU, 2007. ISBN: 87-550-3458-6.

[45] Alexander R. Meyer Forsting, Niels Troldborg, and Mac Gaunaa. "The flow upstream of a row of aligned wind turbine rotors and its effect on power production". In: *Wind Energy* 20.1 (2017), pp. 63–77. DOI: 10.1002/we.1991. URL: https://doi.org/10.1002/we.1991.

[46] N. O. Jensen. *A note on wind generator interaction*. Tech. rep. Risø-M No. 2411. Risø National Laboratory, 1983.

[47] S. Ott and M. Nielsen. *Developments of the offshore wind turbine wake model Fuga*. Tech. rep. Technical University of Denmark, DTU Wind Energy, 2014.

[48] J.P. Verhoef and G. Bergman. *Vaisala Triton® Wind Profiler 609 validation against IEC compliant meteorological mast MM2 at the ECN test site Wieringermeer*. Tech. rep. ECN-X–15-086. ECN, 2015, p. 40. URL: https://www.vaisala.com/sites/default/files/documents/ECN_Triton_609_Validation_x15086rev1.pdf.

[49] Thies Clima. *Wind Direction Transmitter First Class: Instruction for Use (4.3151.xx.40x)*. Document No. 021887/08/22. 2022, p. 31. URL: https://www.thiesclima.com/en/db/dnl/4.3151.xx.40x_wr-geber-firstclass_eng.pdf.

[50] *Evaluation of measurement data – Guide to the expression of uncertainty in measurement*. DOI: 10.59161/JCGM100-2008E, Sections 6.2 and 4.3. JCGM (Joint Committee for Guides in Metrology). 2008.

[51] United States Department of Agriculture (USDA). *Part 650 Engineering Field Handbook - Chapter 1: Surveying*. 210-650-H, 2nd Ed., Amend. 1, Jan 2022. 2022. URL: https://directives.nrcs.usda.gov/sites/default/files2/1712930814/33252.pdf.

[52] R.C. Brinker et al. *The Surveying Handbook*. Chapter 20: Compass Surveying. Springer Science+Business Media New York, 1987.

[53] Stony Brook University Physics Department. *Error Analysis and Significant Figures*. Accessed July 2025. 2006. URL: http://phylabs1.physics.sunysb.edu/introlabs/ReferenceDocs/ErrorAnalysis.pdf.

[54] Jacob Berg et al. *DTU 46100: Introduction to Micrometeorology for Wind Energy*. DTU Wind Energy E No. E-0232, Chapter 4.3.2: Spectra from measured time-series. DTU Wind Energy, 2022.

[55] Amy Shallcross. "The moving blocks bootstrap parametric time series models". In: *Water Resources* 32 (June 1996), pp. 1875–1882. DOI: 10.1029/96WR00928.

[56] Adam H. Monahan. “The Temporal Autocorrelation Structure of Sea Surface Winds”. In: *Journal of Climate* 25.19 (2012), pp. 6684–6700. DOI: 10.1175/JCLI-D-11-00698.1. URL: https://journals.ametsoc.org/view/journals/clim/25/19/jcli-d-11-00698.1.xml.

[57] Manfred Mudelsee. *Climate Time Series Analysis: Classical Statistical and Bootstrap Methods*. Vol. 42. Atmospheric and Oceanographic Sciences Library. Chapter 3.3.1 Nonparametric: moving block bootstrap. Springer, 2014. ISBN: 978-3-319-04449-1. DOI: 10.1007/978-3-319-04450-7.

[58] Kosala Bandara, Christoph Bergmeir, and Rob J. Hyndman. “mstl: Seasonal Decomposition of Time Series by Loess”. In: *arXiv preprint arXiv:2107.13462* (2021). URL: https://arxiv.org/pdf/2107.13462.

[59] statsmodels developers. *MSTL decomposition example*. https://www.statsmodels.org/dev/examples/notebooks/generated/mstl_decomposition.html. Accessed: 2025-07-09. 2025.

[60] Philipp Berens. “CircStat: A MATLAB Toolbox for Circular Statistics”. In: *Journal of Statistical Software* 31.10 (2009). URL: https://doi.org/10.18637/jss.v031.i10.

[61] S.H. Li. “Effect of disjunct sampling on calibration of design wind speed”. In: *Journal of Wind Engineering and Industrial Aerodynamics* 183 (2018), pp. 283–294. DOI: 10.1016/j.jweia.2018.11.016.

[62] Sirje Keevallik. “Wind speed and velocity at three Estonian coastal stations 1969–1992”. In: *Estonian Journal of Engineering* 57 (Jan. 2008). DOI: 10.3176/eng.2008.3.02.

[63] Jaak Jaagus. “Long-term changes in frequencies of wind directions on the western coast of Estonia”. In: *Tallinna Ülikooli Ökoloogia Instituut. Publikatsioonid* 11 (Jan. 2009), pp. 11–24.

[64] Jaak Jaagus and Ain Kull. “Changes in surface wind directions in Estonia during 1966-2008 and their relationships with large-scale atmospheric circulation”. In: *Estonian Journal of Earth Sciences* 60 (Dec. 2011), pp. 220–231. DOI: 10.3176/earth.2011.4.03.

[65] Sirje Keevallik et al. “Outlook for wind measurement at Estonian automatic weather stations”. In: *Estonian Journal of Engineering* 13 (Jan. 2007), p. 234. DOI: 10.3176/eng.2007.3.05.

[66] A. Gadian et al. “Directional persistence of low wind speed observations”. In: *Journal of Wind Engineering and Industrial Aerodynamics* 92.12 (2004), pp. 1061–1074. ISSN: 0167-6105. DOI: https://doi.org/10.1016/j.jweia.2004.05.007. URL: https://www.sciencedirect.com/science/article/pii/S0167610504000947.

[67] Dimitri Foussekis and Fragiskos Mouzakis. “Wind resource assessment uncertainty for a TLP-based met mast”. In: *Journal of Physics: Conference Series* 2018 (Sept. 2021), p. 012018. DOI: 10.1088/1742-6596/2018/1/012018.

[68] Morten Nielsen. *Long-term correction of wind observations by diffusion-based transformation*. English. DTU Wind Energy E 0183. Denmark: DTU Wind Energy, 2019. ISBN: 978-87-93549-51-7.

[69] Wiebke Langreder and Jørgen Højstrup. *Long-term correction: Facts and Fiction*. https://www.ewea.org/events/workshops/wp-content/uploads/2015/05/Tech15a-PO-003.pdf. Wind Solutions, Denmark, presented at EWEA Workshop. 2015.

[70] Sónia Liléo et al. *Long-Term Correction of Wind Measurements*. Tech. rep. Elforsk Report 2013:18. Kjeller Vindteknikk, Jan. 2013. URL: https://energiforskmedia.blob.core.windows.net/media/19814/long-term-correction-of-wind-measurements-elforskrapport-2013-18.pdf.

[71] New York State Energy Research and Development Authority. *Wind Resource Assessment Handbook: Final Report*. Tech. rep. Final Report 10-30. Albany, NY, USA: New York State Energy Research and Development Authority, Oct. 2010. URL: https://www.inscc.utah.edu/~krueger/5270/wind_resource_handbook.pdf.

[72] {Niels Gylling} Mortensen, Morten Nielsen, and Hans {Ejsing Jørgensen}. *First Offshore Comparative Resource and Energy Yield Assessment Procedures (CREYAP)*. English. EWEA Offshore 2013 ; Conference date: 19-11-2013 Through 21-11-2013. 2013. URL: http://www.ewea.org/offshore2013/.

[73] J. Ramon et al. "The Tall Tower Dataset: a unique initiative to boost wind energy research". In: *Earth System Science Data* 12.1 (2020), pp. 429–439. DOI: 10.5194/essd-12-429-2020. URL: https://essd.copernicus.org/articles/12/429/2020/.

[74] DTU Wind Energy. *WakeDeficitModels — PyWake Documentation*. Accessed: 2025-08-02. 2018. URL: https://topfarm.pages.windenergy.dtu.dk/PyWake/notebooks/WakeDeficitModels.html.

[75] {Søren Ejling} Larsen. *The atmospheric boundary layer over land and sea: Focus on the off-shore Southern Baltic and Southern North Sea region*. English. Denmark: DTU Wind Energy, 2013.

[76] Majid Bastankhah and Fernando Porté-Agel. "A new analytical model for wind-turbine wakes". In: *Renewable Energy* 70 (2014). Special issue on aerodynamics of offshore wind energy systems and wakes, pp. 116–123. ISSN: 0960-1481. DOI: https://doi.org/10.1016/j.renene.2014.01.002. URL: https://www.sciencedirect.com/science/article/pii/S0960148114000317.

[77] M. Kelly and H. E. Jørgensen. "Statistical characterization of roughness uncertainty and impact on wind resource estimation". In: *Wind Energy Science* 2.1 (2017), pp. 189–209. DOI: 10.5194/wes-2-189-2017. URL: https://wes.copernicus.org/articles/2/189/2017/.

[78] Jacob Berg et al. *DTU 46100: Introduction to Micrometeorology for Wind Energy*. English. DTU Wind Energy E E-0232. Denmark: DTU Wind Energy, 2022.

A Hand Compass

Figure A.1: Installer using hand compass in orientation stage.

B LTWRD Uncertainty Quantification

B.1 Limitations of the Sub-sampling Approach

As discussed in Section 3.2, the sub-sampling method was not selected for estimating LTWRD uncertainty due to two key limitations observed in this study.

First, the distribution of normalised AEP errors (computed between measured and reanalysis datasets) was found to be distinctly non-Gaussian, even when increasing the number of windows (Figure B.1). The errors tended to cluster into separate groups, a behaviour likely caused by persistent biases within certain periods of the relatively short 5.5-year concurrent dataset. Such clustering can result in an overestimation of the standard deviation. By contrast, combining MBB with MSTL decomposition mitigates the dominance of specific biased periods by: (1) isolating and removing trend and seasonal components, and (2) resampling continuous blocks randomly.

Second, the short concurrent period means there are few truly independent 1-year samples available. Sub-sampling without replacement therefore yields samples that are highly similar, which risks underestimating variability. MBB addresses this by allowing overlapping and repeated sampling of time blocks, thereby creating a more diverse ensemble of synthetic 1-year realisations while still preserving key temporal dependencies such as autocorrelation and seasonality.

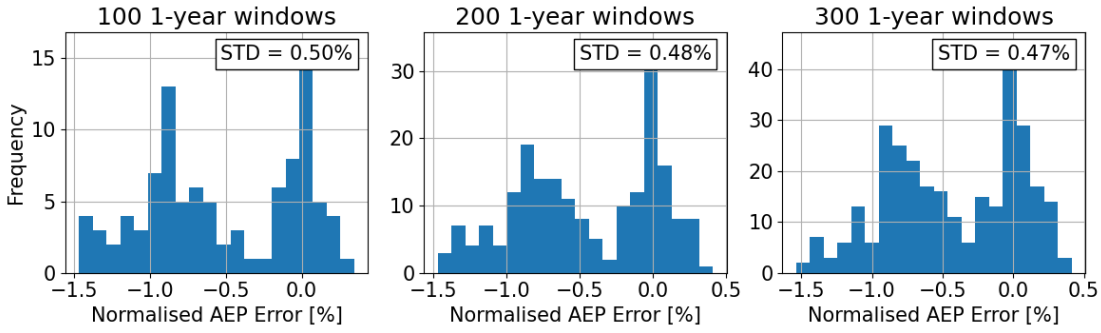

Figure B.1: Effect of increasing the number of windows on the distribution of normalised AEP errors in the sub-sampling approach.

B.2 Avoiding Seasonal Bias in MBB Sampling

Applying MBB directly to the original wind time series can distort the seasonal composition of resampled 1-year windows. Ideally, each month should appear with a frequency of $1/12 \approx 0.083$. However, empirical results from six random 1-year MBB samples (Table B.1) show large deviations from this ideal. Some months occur more than twice as often, while others are entirely absent, leading to an over- or under-representation of specific seasons and biasing the seasonal wind characteristics.

Table B.1: Monthly Representation in Six Resampled 1-Year MBB Windows

Index	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1	0.144	0.059	0.077	0.092	0.088	0.093	0.061	0.070	0.069	0.118	0.052	0.076
2	0.059	0.093	0.097	0.102	0.147	0.058	0.096	0	0.038	0.058	0.172	0.079
3	0.038	0.171	0.033	0.095	0.096	0.096	0.109	0.076	0.019	0.071	0.033	0.164
4	0.048	0.077	0.181	0.078	0.046	0.099	0.063	0.102	0.036	0.091	0.085	0.096
5	0.079	0.102	0.182	0.153	0.043	0.080	0.030	0.057	0.040	0.113	0.039	0.081
6	0.018	0.087	0.123	0.086	0.092	0.026	0.056	0.085	0.028	0.172	0.096	0.130

Since wind direction influences AEP indirectly by shaping the frequency-weighted wind speed distribution, seasonal shifts in directional frequency can bias energy estimates. At a given site, wind speed characteristics vary across directional sectors (Figure 2.4). Thus, an over-representation of certain months in resampled series changes the weighted wind speed distribution, introducing a seasonal bias in energy estimation.

To quantify this, the wind rose deviation (WRD) between mast and reanalysis data was calculated for each month using an velocity-weighting exponent $p = 2$, following Kelly's [77] formulation:

$$WRD_m = \sum \left(|f_{\text{synth, i}} - f_{\text{mast, i}}| \cdot A^2_{\text{mast, i}} \right) \tag{B.1}$$

where $f_{\text{synth, i}}$ and $f_{\text{mast, i}}$ represent the wind direction frequency distributions from synthetic and mast data respectively in sector i, and $A_{\text{mast, i}}$ is the Weibull scale parameter derived from mast measurements in sector i.

Results for six MBB realisations (Table B.2) show large variability between months, with deviations tending to peak in winter. Averaging across 100 random 1-year realisations (Table B.3) confirms a clear seasonal pattern: winter months show systematically higher disagreement between mast and reanalysis wind roses.

Table B.2: Monthly WRD Between Mast and Reanalysis Across Six Resampled 1-Year MBB Windows for *Site A*.

Index	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1	48.03	45.27	47.31	35.92	46.20	50.96	27.96	31.80	51.57	48.91	67.75	34.00
2	56.25	69.71	56.09	66.96	49.73	25.71	39.64	37.24	41.19	37.22	34.48	66.07
3	73.41	50.81	44.59	47.75	35.87	21.82	25.99	28.44	49.69	44.17	49.49	40.64
4	48.48	45.27	47.31	35.92	46.20	50.96	27.96	31.80	51.57	48.91	62.19	38.45
5	48.03	45.27	44.59	35.92	47.20	57.81	28.43	32.82	53.16	48.17	49.49	58.50
6	70.38	77.24	74.83	28.61	24.96	57.81	23.28	39.82	35.33	30.34	58.34	60.76

Table B.3: Averaged WRD per Month Between Mast and Reanalysis over 100 Random 1-Year MBB Windows for *Site A*.

Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
56.20	97.18	56.09	45.08	39.50	38.04	28.23	32.55	47.66	42.95	56.41	50.84

To avoid conflating seasonal effects with random variability, MBB is applied to residuals after removing trend and seasonality. This ensures that bootstrap samples preserve short-term autocorrelation structures without introducing artificial seasonal imbalances, leading to a more unbiased estimate of LTWRD uncertainty.

B.3 Sample Size Convergence Analysis

To ensure that the Monte Carlo resampling approach provides statistically reliable estimates, a convergence analysis is performed by varying the number of resampled 1-year windows: 100, 200 , and 300. The resulting distributions of normalised AEP error is shown in Figure B.2.

As the number of samples increase, the distribution becomes progressively smoother and more symmetric, approaching the shape of a Gaussian distribution. The standard deviation of the normalised AEP error decreases slightly from 0.37% (100 samples) to 0.35% (200 samples) and remains stable at 0.35% for 300 samples. The numbers in the standard deviation indicates that the resampling process has converged with 200 samples.

Although the distribution does not perfectly match a normal distribution, it exhibits no significant skewness or multimodal characteristics. Therefore, the use of 200 samples is considered statistically sufficient for quantifying the LTWRD uncertainty in this study.

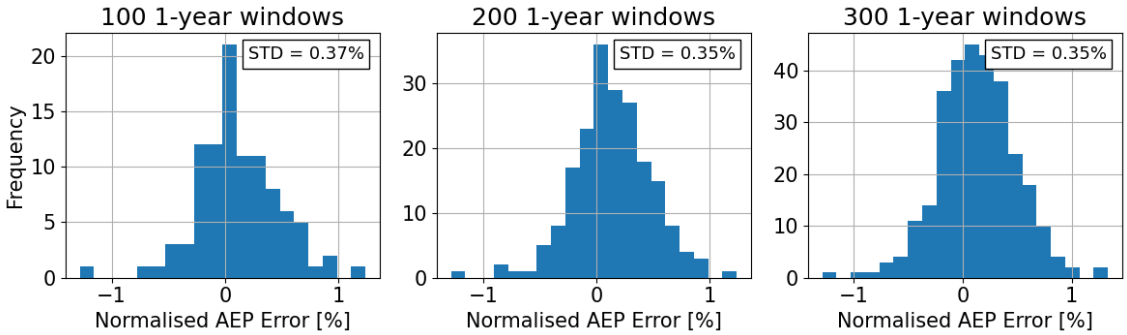

Figure B.2: Distribution of normalised AEP error across different Monte Carlo sample counts: 100, 200, 300 one-year windows.

B.4 Disjunct vs. Block-Averaging Sampling

As described in Section 3.2.3, the effect of temporal resolution on AEP estimation was evaluated by comparing disjunct sampling with block-averaging applied to the 10-minute measured data. In the disjunct approach, only the first observation of each hour is retained, simulating 1-hour resolution without smoothing. In contrast, block-averaging computes hourly values by averaging all six 10-minute observations within each hour.

Figure B.3 shows the distributions of normalised AEP errors from 200 Monte Carlo simulations for both methods. The results are nearly identical: mean errors of 0.26% (disjunct) and 0.33% (block-averaged), with both having a standard deviation of 0.37%. These negligible differences confirm the theoretical expectation that sub-hourly variability has little effect on annual energy estimates.

Given its simplicity, computational efficiency, and demonstrated equivalence to block-averaging in this context, the disjunct sampling method was adopted for all subsequent analyses in this study.

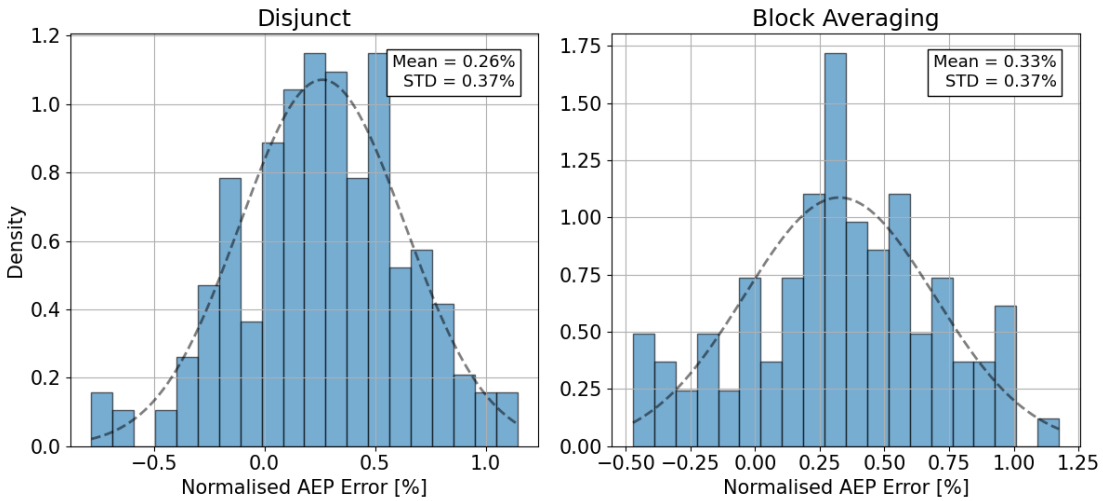

Figure B.3: Comparison of normalised AEP error distributions using disjunct sampling (left) and block-averaging (right) on 10-minute measured data.

C Quantify the Variability in Wind Direction

C.1 Simple directional statistics, mean and variance of wind directions

To characterise wind conditions, both wind speed and wind direction are described in a planar geographic coordinate system, with wind direction θ defined between 0° and 360° (where 0° is equivalent to 360°). The cyclic nature of wind direction complicates conventional averaging: for example, two equal wind speeds at 1° and 359° have an arithmetic mean of 180°, which is physically meaningless. Correct averaging must account for the circular domain, which is achieved using vector-based methods.

Following the generalised vector averaging approach [78], each wind direction is decomposed into horizontal components. For N wind vectors, each with speed U_n and direction θ_n, the component-wise averages are:

$$s_{va} = \overline{U^p sin\theta} \equiv \frac{1}{N}\Sigma_{n=1}^{N} {U_n}^p sin\theta_n \tag{C.1}$$

$$c_{va} = \overline{U^p cos\theta} \equiv \frac{1}{N}\Sigma_{n=1}^{N} {U_n}^p cos\theta_n \tag{C.2}$$

Where p is a user-defined power exponent. In this study, $p = 2$ is used to approximate an energy-weighted mean, following the framework of Kelly (2017) [77]. Kelly defines the effective wind-power exponent p through the relationship $AEP \propto U^p$, where $p = ln(AEP)/ln(U)$, with assumptions of a canonical power curve with cubic dependence below rated speed and Rayleigh-distributed wind speeds. For typical turbine sites, where the mean wind speed at hub height is approximately 60-80% of the rated wind speed, hence the effective exponent is consistently close to 2.0. The average wind direction is then given by [78]:

$$\overline{\theta} = \frac{180^o}{\pi} mod\left[tan^{-1}\left(\frac{s_{va}}{c_{va}}\right),\ 2\pi\right] \tag{C.3}$$

To quantify the spread of wind directions, the Yamartino method [18] is commonly used, as it avoids errors from angular wrap-around, i.e. the direction abruptly crossing from 359^o to 0^o. It employs a vector average with $p = 0$, meaning all directions are treated equally regardless of wind speed, which is appropriate for assessing distribution shape rather than energy content (which calculate with $p = 2$). The standard deviation of wind direction is then:

$$\epsilon = \sqrt{1 - s_{va,p=0}^2 - c_{va,p=0}^2} \tag{C.4}$$

$$\sigma_\theta = sin^{-1}(\epsilon)\left[1 + \left(\frac{2}{\sqrt{3}} - 1\right)\epsilon^3\right] \tag{C.5}$$

The original aim was to assess interannual variability (IAV) in wind direction by computing the annual mean wind direction for each year (interpreted as the dominant inflow direction) and assuming that the annual spread of wind directions remained relatively constant. IAV would then be quantified as the standard deviation of these annual means.

However, applying the Yamartino method to annual data from *Site A* and *Site B* yielded yearly σ_θ values close to $70°$, indicating that the wind direction distributions are broad, multimodal, or nearly uniform. In such cases, the concept of a single "dominant" direction is not meaningful. Nevertheless, the standard deviation of annual means over the full reanalysis period was still computed to test whether it could provide a representative measure of IAV, as discussed in the following section.

C.2 Limitations of Using Yamartino-Derived Standard Deviation for IAV Representation

The IAV uncertainty calculated from the original dataset for *Site A* is 1.83%. To test whether this could be reproduced using a simplified statistical representation, the standard deviation of the yearly mean wind direction was computed using vector averaging followed by the Yamartino method. This yielded an interannual variation σ_θ of $11.2°$, interpreted here as "how much the dominant inflow direction varies from year to year."

To test its suitability for modelling IAV, σ_θ was used as the standard deviation for a Monte Carlo simulation. A normal distribution with 200 random samples (mean=$0°$, standard deviation=$11.2°$) was generated, each value representing a random directional shift θ_{shift} applied to the original wind rose WR_{original}:

$$WR_{\text{shifted}} = (WR_{\text{original}} + \theta_{\text{shift}}) \bmod 360 \tag{C.6}$$

For each shifted wind rose WR_{shifted}, AEP was recalculated. The resulting standard deviation across the 200 AEP values was only 0.08%, more than an order of magnitude smaller than the 1.83% obtained from the real dataset. This severe underestimation confirms that using annual mean wind direction variability as a surrogate for IAV does not represent the true directional uncertainty in AEP.

The underlying reason lies in the limitations of vector averaging and Yamartino-based standard deviation for wind direction:

1. Loss of distribution shape: The mean direction capture only dominant inflow angle and discards information about secondary peaks, distribution width, or asymmetry.
2. Bimodality bias: In bimodal years, the mean can fall in a low-frequency direction between two peaks, yielding misleading variability metrics.
3. Ignores shape evolution: Interannual variability often manifest as changes in the shape of the wind rose (e.g., shifts in lobe frequency or spread), not just its centre.

At *Site A*, the wind direction distribution is inherently asymmetric and frequently multimodal. These characteristic mean that a single scalar metric such as Yamartino σ_θ cannot capture the year-to-year evolution of directional patterns that influence energy yield. Therefore, simulating IAV using only the variability of the annual mean wind direction is not representative.