Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Comprehensive Predictive Models for Offshore Wind Farm Performance: Integrating Reliability, Energy Production and Cost Analysis for Current and Next-Generation Installations

Adriano García Piquero

adriano.garcia@uclm.es

University of Castilla-La

Miguel Ángel Pardo Picazo

University of Alicante

Research Article

Keywords: offshore wind farm, predictive modelling, LCOE, reliability, energy production, Monte Carlo simulation, floating offshore wind

Posted Date: September 8th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10840266/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

The offshore wind sector is undergoing rapid turbine upscaling (from 3–4 MW to 12–15 + MW), expansion into deeper waters, and mounting pressure to reduce the levelised cost of energy (LCOE) to levels competitive with established generation. This article presents a modular, reproducible simulation platform—validated against real onshore operational data from a 97-turbine fleet—that couples three interdependent sub-models: a Weibull-based reliability model with stress-factor onshore-to-offshore translation; an energy-production model incorporating capacity factor, availability and weather-dependent losses; and an economic model computing LCOE, O&M cost and CO_2 intensity. The novelty is not any individual sub-model but their consistent coupling: the availability generated stochastically by the reliability model propagates into the energy denominator of the LCOE, so that reliability, yield and cost are never assumed independently but derived from one shared set of variables. The platform applied to four scenarios spanning the evolutionary trajectory: (S1) a 97-turbine onshore reference (145.5 MW, 1.5 MW machines); (S2) its offshore translation; (S3) a contemporary 640 MW farm (80 × 8 MW); and (S4) a next-generation 900 MW floating installation (60 × 15 MW). All numerical results are produced by the model from a single fixed random seed (2027) and are therefore exactly reproducible; reproducibility is offered to let reviewers verify the computation, not as evidence that the model is correct. The model yields a declining LCOE trajectory of 119 → 79.5 → 65.8 €/MWh across S2–S4 (with Monte Carlo 90% intervals of roughly ± 6–11 €/MWh), a 45% reduction driven by higher capacity factors, improved reliability per MW and O&M economies of scale. O&M consistently represents 34–35% of offshore LCOE. A tornado sensitivity analysis identifies WACC, capacity factor and turbine reliability as the three dominant drivers. Because the S3 and S4 scenarios rest on literature-anchored scaling assumptions rather than fleet-scale measurements, their results are presented as technically plausible projections whose sensitivity is quantified explicitly.

1. Introduction

The global offshore wind sector reached a cumulative installed capacity exceeding 75 GW by the end of 2024, with the European Union targeting 360 GW by 2050 and the United States pursuing 30 GW by 2030 [1]. This rapid expansion is accompanied by a parallel technological transition: turbine ratings have grown from 2–3 MW in early North Sea installations to 12–15 MW platforms now entering series production, with 20 + MW designs on manufacturer roadmaps [2, 3]. Simultaneously, the industry is shifting from shallow-water monopiles to floating platforms that access waters deeper than 60 m, unlocking previously uneconomical resource zones [4].

The 2022–2024 period exposed the sector to significant economic headwinds. Inflation, rising interest rates, and supply-chain disruptions raised project costs by 30–40% and led to the termination or renegotiation of over 10 GW of contracted capacity [5]. Recent industry analysis argues for a fundamental rethinking of cost structure rather than incremental optimisation [5]. In this context, comprehensive predictive models spanning the full project lifecycle are indispensable for de-risking investment decisions under persistent uncertainty.

The technical performance of an offshore wind farm is the product of three tightly coupled physical and economic processes. Reliability determines how often subsystems fail and how long turbines are unavailable; energy production determines how much of the available wind resource is converted to electricity given wake and electrical losses and that availability; and the cost of energy is the ratio of lifetime cost to lifetime yield. These three processes are usually modelled in isolation. Reliability studies report failure rates and downtime for onshore and offshore fleets [6, 7, 8] and, more recently, use Bayesian networks and fault trees to reason about floating configurations for which little operational evidence exists [9, 10]. Energy-yield studies focus on wake modelling and resource assessment. Techno-economic studies report LCOE trajectories from cost databases [2, 3] and installation models [4]. Each body of work is mature in isolation, but the interfaces between them—in particular, how a change in reliability propagates through availability into the cost of energy—are frequently handled by fixed assumptions rather than by a shared computation.

Two gaps motivate this work. The first is methodological. Because availability is often assumed as a fixed input to the LCOE rather than simulated from the underlying failure and repair processes, reliability improvements cannot be traced through to their cost consequences within a single consistent model; the economic value of reliability is then asserted rather than derived. The second is empirical and forward-looking. The turbines that will dominate the pipeline over the next decade—8 MW fixed-bottom and 15 MW floating machines—have no fleet-scale operational record, so their reliability and cost cannot be measured, only projected. The community therefore needs a framework in which such projections are made transparently, with their assumptions stated and their influence on the headline cost figure bounded, rather than presented as measurements.

This article addresses both gaps with a single modular platform that couples reliability, energy, and cost so that the economic outputs are computed from the same reliability and energy variables the model produces. The availability generated by a Monte Carlo failure-and-repair simulation feeds directly into the capacity factor that appears in the LCOE denominator, so a change in reliability moves the cost of energy in an internally consistent way. The platform is exercised on four scenarios that trace the sector's evolution—an empirically anchored onshore reference, its offshore translation, a contemporary fixed-bottom farm and a next-generation floating installation—and the uncertainty of the two forward-looking scenarios is quantified rather than hidden. The specific contribution is the coupling itself: individually, Weibull reliability, Jensen wake modelling [17], [18], Monte Carlo availability and capital-recovery based LCOE are all established in wind-energy analysis ([19], [20]). What has not previously been assembled, to the author's knowledge, is a single reproducible pipeline in which the stochastic availability of next-generation machines drives the cost of energy end to end, allowing reliability, yield and cost to be reasoned about jointly rather than sequentially.

This article makes four contributions:

(1) an integrated, modular reliability–energy–economic simulation platform whose economic outputs are computed from the same reliability and energy variables the model produces, ensuring internal

consistency between availability, capacity factor and LCOE;

(2) a four-scenario analysis of the onshore-to-next-generation trajectory that quantifies a plausible LCOE reduction pathway, decomposes it into CapEx and O&M contributions, and reports Monte Carlo uncertainty intervals for each headline figure;

(3) a reproducible Monte Carlo availability and sensitivity analysis, driven by a single fixed seed, that identifies WACC, capacity factor and turbine reliability as the dominant LCOE drivers; and

(4) a transparent treatment of next-generation uncertainty, in which the scaling laws for 8–15 MW machines are stated as assumptions, cross-checked against the sparse published evidence, and bounded through sensitivity analysis rather than asserted.

The intended audience is twofold: researchers who need a transparent, reproducible baseline against which to test reliability or cost interventions, and industrial analysts who need to understand which levers most affect the cost of next-generation offshore projects. To serve both, the complete code and the fixed seed are released openly (Data availability), so that every number in the article can be regenerated and every assumption inspected and changed.

Two boundaries and one limitation are stated at the outset. This article addresses lifecycle performance prediction across turbine generations; the estimation of reliability priors under data scarcity and the AI-driven optimisation of maintenance strategies are treated in companion work and are used here only as background. The reliability translation from onshore to offshore relies on published stress factors and the energy model uses an analytical wake formulation; both are deliberately simple, and their limitations are discussed in Section 6. Most importantly, because no fleet-scale operational data yet exist for 8–15 MW offshore machines, the S3 and S4 scenarios rest on conservative, literature-anchored scaling assumptions whose sensitivity is quantified explicitly in Section 5.3; their results should be read as technically plausible scenario projections rather than measurements.

2. Integrated model architecture

2.1. Modular design

Reliability sub-model. Generates stochastic failure sequences per subsystem from Weibull distributions whose scale is translated from onshore to offshore conditions through published stress factors (Section 3.1). The model supports bathtub-curve behaviour for burn-in and wear-out phases. For the S3 and S4 scenarios, category-specific scaling laws adjust the base failure rates by turbine rating, following the reasoning displayed and justified in Section 3.1.

Energy sub-model. Computes gross and net energy from turbine power curves, site-specific Weibull wind distributions and availability schedules, applying a Jensen wake model for array effects and fixed electrical losses (transformer, array and export cable, grid connection). The Jensen formulation is an

analytical first-order representation of array wakes; its adequacy and the alternatives considered are discussed in Section 3.2 and revisited among the limitations in Section 6.

Economic sub-model. Calculates LCOE, O&M cost per MW and CO_2 intensity. Capital cost covers turbine, foundation, grid connection and installation; operating cost covers scheduled and unscheduled maintenance, vessel charter, spare parts and labour. Crucially, the availability produced by the reliability sub-model feeds directly into the energy denominator of the LCOE, so that reliability improvements are reflected in the cost metric rather than assumed away.

2.2. LCOE formulation

LCOE = (CapEx × FCR + OpEx) / (P × 8760 × CF_net), (1)

where CapEx is the total installed cost, FCR the fixed charge rate derived from the WACC and project lifetime through the standard capital-recovery factor, OpEx the annual operating cost, P the installed capacity and CF_net the net capacity factor. The net capacity factor already embeds wake and electrical losses and the availability returned by the reliability sub-model, which is what couples reliability to cost.

Two of the inputs to Eq. (1)—CapEx and WACC—are exogenous, taken from published cost databases and financing assumptions rather than produced by the simulation. This is deliberate and its consequence should be stated plainly: because CapEx and WACC enter multiplicatively and dominate the numerator, the broad level of the LCOE and much of the difference between scenarios is set by these external inputs, not by the stochastic sub-models. The simulation's role is not to predict CapEx or the cost of capital, which are macro-financial quantities outside its scope, but to determine, consistently and reproducibly, the energy denominator—availability, wake and electrical losses, net capacity factor—and the O&M component, and to show how sensitive the resulting cost of energy is to each input (Section 5.3). Framed this way, the declining LCOE trajectory is not a prediction manufactured by the model but the arithmetic consequence of published cost trajectories combined with a transparently simulated denominator; the model's contribution is the consistency of that combination and the sensitivity information around it, and the CapEx assumptions are examined directly in the sensitivity analysis so their dominance is visible rather than hidden. The four scenarios analysed in this study, together with their full parameter specifications, are summarised in Table 1.

2.3. Scenario definitions

Table 1
Modelled wind-farm scenarios with full parameter specification.

Parameter	S1 Onshore	S2 Offshore	S3 Contemp.	S4 Next-gen	Source
Turbine rating (MW)	1.5	1.5	8	15	Design
Number of turbines	97	97	80	60	Design
Total capacity (MW)	145.5	145.5	640	900	Design
Foundation	—	Monopile	Mono/jacket	Semi-sub float	Design
Distance to shore (km)	5	50	80	120	Site
Water depth (m)	—	25	35	100+	Site
Gross capacity factor	30%	46%	52%	56%	Met-ocean
Project lifetime (yr)	20	25	30	30	Design
WACC	5%	9%	7.5%	7.0%	Finance
CapEx (M€/MW)	1.15	2.60	2.30	2.05	NREL ATB [2]
OpEx (k€/MW/yr)	38	140	105	85	Model
K_1 (environmental)	1.0	1.75	1.75	1.75	Davison [11]
K_2 (power rating)	1.0	1.78	1.78	1.78	Reliability model

The CapEx values follow the NREL 2024 ATB / Cost of Wind Energy Review trajectory for fixed-bottom and mature floating installations [2, 3]. They were adopted from these external sources rather than tuned: an earlier draft of this platform used internally assumed CapEx figures that produced an LCOE inconsistent with the same draft's capital-recovery arithmetic, and the values were therefore re-anchored to the NREL ATB reference trajectory to remove that inconsistency. The re-anchoring corrects an internal economic inconsistency; it is not a calibration toward a target LCOE, and the sensitivity of the results to CapEx is reported explicitly in Section 5.3. The gross capacity factors are combined with modelled wake, electrical and availability losses to obtain the net values reported in Section 5. The scenario dimensions are chosen to be representative of real project classes rather than arbitrary: S3 (80 × 8 MW, 640 MW, 80 km) reflects the scale and distance of fixed-bottom farms entering operation in the North Sea in the mid-2020s, and S4 (60 × 15 MW, 900 MW, 120 km, ~ 100 m depth) reflects the capacity, turbine count and water depth typical of the floating and deep-water lease areas being auctioned for commissioning around 2030; the turbine counts follow from the chosen ratings and total capacities.

3. Key model components

3.1. Reliability scaling for next-generation turbines

Eight subsystems are modelled—gearbox, generator, converter, pitch system, blades, yaw system, tower and structure, and control and electrical systems—using the same partition adopted in the wider reliability literature [6, 7]. The onshore-to-offshore translation applies the Davison stress factors [11], K_1 (environmental aggravation) and K_2 (rating and drivetrain), giving a combined $K_1 \cdot K_2$ in the range 2.67–3.55 across subsystems. Because no fleet-scale data exist for 8–15 MW machines, conservative category-specific scaling laws relative to the 1.5 MW baseline are then adopted for the S3 and S4 ratings. These laws are heuristics, not measurements, and are justified below by physical reasoning and by the sparse published evidence; where no rigorous basis exists, this is stated rather than concealed.

Mechanical (gearbox, pitch, yaw): failure rates are assumed to scale with the cube root of the power ratio, giving multipliers of (8/1.5)^(1/3) = 1.74 for S3 and (15/1.5)^(1/3) = 2.15 for S4. The cube-root form is a deliberately conservative order-of-magnitude choice rather than a derived structural law: it reflects the empirical observation that drivetrain failure rates grow markedly more slowly than rated power, because larger machines use fewer, more robust, better-instrumented components per MW, but it is not obtained from a stress–strain analysis. A full structural fatigue model of each 15 MW subsystem is beyond the scope of a farm-level platform and, more fundamentally, would require geometry and loading data that manufacturers do not publish. The sensitivity of the LCOE to this exponent is bounded by the failure-rate row of the sensitivity analysis (Section 5.3), and structured expert-elicitation estimates of major-replacement rates for next-generation turbines [12] are used as an external cross-check on the resulting aggregate rates.

Electrical (converter, generator, control)

a linear scaling calibrated to 1.20 for S3 and 1.35 for S4, reflecting the modest increase in electrical-subsystem failure rates observed as power electronics scale, again treated as an assumption rather than a measurement.

Structural (blades, tower)

a decrease per MW (0.85 for S3, 0.75 for S4) reflecting design maturation and the fact that structural failures per unit capacity have historically fallen as blade and tower engineering matured.

Floating-specific (S4)

mooring-line and dynamic-cable failure modes are added from the Bayesian-network estimates of Li et al. [9], with additional uncertainty margins. These modes have essentially no operational track record and are the most uncertain element of the S4 estimate; their aggregate effect on LCOE is bounded in the sensitivity analysis of Section 5.3.

Avoiding double counting. Applying the Davison factor K_2 (which already captures the drivetrain consequences of higher power rating) together with the mechanical cube-root rating law risks counting

the effect of size twice. To avoid this, K_2 is applied only to translate the 1.5 MW onshore baseline to an equivalent offshore 1.5 MW machine, and the cube-root rating law is applied only to the further step from 1.5 MW to the 8 and 15 MW ratings; the two factors act on disjoint parts of the scaling chain rather than compounding on the same rating increase. K_1 (environmental aggravation) is independent of rating and is applied uniformly across offshore scenarios. This decomposition is stated here so that the origin of every multiplier in the aggregate failure rates of Table 3 is traceable.

3.2. Energy production

Higher hub heights (130–175 m for S3/S4) access stronger, less turbulent winds, raising the gross capacity factor relative to S2. Wake losses are computed with the Jensen (Park) model, yielding array losses of 8–12% depending on layout and prevailing wind; electrical losses are 2–3% for fixed-bottom and 3–4% for floating configurations. The net capacity factor combines gross CF, wake efficiency, electrical efficiency and the Monte Carlo availability of Section 3.3.

The choice of the Jensen model warrants comment, since it is the simplest defensible wake formulation. Jensen assumes a linearly expanding wake with top-hat velocity deficit and neglects atmospheric-boundary-layer turbulent mixing and the deep-array (deep-farm) effect that becomes important in the interior of very large farms such as S3 and S4; for floating platforms it also neglects wake meandering induced by platform motion. It is retained here for three reasons: it is transparent and reproducible, its parameters are physically interpretable, and the platform's contribution lies in the reliability–energy–cost coupling rather than in wake physics. To ensure the simplification does not distort the conclusions, the 8–12% array-loss band was checked against Gaussian wake-model results [21] reported for large offshore arrays and against operational loss figures for existing farms, and found to sit at the optimistic end of the plausible range; the effect of a higher wake loss is captured through the capacity-factor row of the sensitivity analysis. A full Gaussian or deep-array wake treatment is identified as future work in Section 6.

3.3. Monte Carlo availability model

Availability is not assumed but simulated. For each of 10,000 iterations (seed 2027) the annual number of failures per turbine is drawn from a Poisson process with the scenario-specific aggregate rate; each failure is classified as minor (repaired by crew-transfer or service-operation vessel) or major (requiring heavy-lift vessel or, for floating units, tow-to-port), with lost production sampled from gamma distributions for repair duration and weather-access delay. The full workflow is summarised in Figure A (Section 3.5).

To make the stochastic model reproducible in detail, the distribution parameters are stated explicitly. Repair duration for minor events is Gamma ($k = 2.0$, θ) with mean 8–24 h; major events use Gamma ($k = 2.5$, θ) with mean 3–14 days depending on vessel class. Weather-access delay is Gamma ($k = 1.5$, θ) with a scenario-dependent mean that grows with distance to shore: means of roughly 1–2 days for S1/S2 (< 60 km), 2–4 days for S3 (80 km) and 4–8 days for S4 (120 km). The waiting-for-weather threshold is a significant wave height of $H_s \leq 1.5$ m and wind speed ≤ 12 m/s for crew-transfer access,

relaxed to Hs ≤ 2.5 m for service-operation vessels; heavy-lift operations require Hs ≤ 1.75 m. For the floating S4 scenario, a major replacement modelled as tow-to-port is assigned an additional immobilisation of 14–30 days covering mooring disconnection, transit, quayside repair and re-hook-up, sampled from Gamma (k = 3.0, θ); this transit-and-disconnect downtime is the principal reason S4 availability (89.2%) falls well below the fixed-bottom scenarios despite comparable per-turbine failure rates.

What the Monte Carlo adds. The simulation propagates the assumed variability in failure counts and repair/weather delays into a distribution of annual availability; it does not learn parameters from data or propagate experimental measurement uncertainty, because the underlying rates for next-generation machines are themselves assumptions. Replacing the full simulation with a deterministic mean-value calculation changes the central LCOE estimates by less than 1 €/MWh in every scenario. The value of the Monte Carlo therefore lies not in shifting the central estimate but in supplying the availability and distribution and hence the uncertainty bands reported in Section 5, and it is presented in that role rather than as an independent pillar of predictive accuracy.

3.4. O&M cost assumptions

Vessel day-rates, logistics constraints and access strategies follow published offshore O&M practice [13], scaled by distance to shore and, for the floating scenario, by the additional demands of tow-to-port replacement. The resulting assumptions are summarised in Table 2.

Table 2
O&M vessel and logistics cost assumptions by scenario.

Asset	S1/S2 (< 60 km)	S3 (80 km)	S4 (120 km, floating)	Unit
CTV day-rate	1,000–2,000	1,500–2,500	SOV-based	€/day
SOV day-rate	—	25,000–40,000	30,000–45,000	€/day
Jack-up day-rate	100,000–250,000	130,000–280,000	tow-to-port	€/day
Helicopter	5,000–8,000	6,000–9,000	7,000–10,000	€/flight
Tow-to-port	—	—	15,000–25,000	€/day
Labour rate	400–600	400–600	500–700	€/tech/day
Mobilisation time	4–8	6–12	12–24	h

3.5. Model workflow

The three sub-models form a single directed pipeline in which the output of each stage is an input to the next, which is what enforces internal consistency. Figure A summarises the flow: scenario parameters and per-subsystem onshore rates enter the reliability stage, which applies the offshore stress factors and rating scaling laws to produce aggregate failure rates; these drive the Monte Carlo availability stage, whose mean availability enters the energy stage together with wake and electrical losses to yield the net

capacity factor; the net capacity factor and the exogenous CapEx, OpEx and WACC then enter the economic stage to produce the LCOE. Because availability is passed forward rather than assumed, a change anywhere in the reliability stage propagates deterministically to the cost of energy.

Figure A. Model workflow. Scenario inputs and onshore reliability data feed the reliability stage; the resulting failure rates drive the Monte Carlo availability model, whose mean availability enters the energy stage to form the net capacity factor; this and the exogenous CapEx, OpEx and WACC produce the LCOE. Availability is propagated forward rather than assumed.

4. Case study: Four evolutionary scenarios

The platform is exercised on four scenarios chosen to span the evolutionary trajectory of the sector rather than to represent a single site. S1 is the empirical onshore baseline: a 97-turbine, 145.5 MW fleet of 1.5 MW machines for which real operational data are available and against which the reliability and energy sub-models are validated. S2 is the hypothetical offshore translation of that same fleet, isolating the effect of the marine environment at constant turbine technology. S3 is a contemporary 640 MW farm of 80 × 8 MW turbines at 80 km from shore, representative of projects entering operation today. S4 is a speculative 900 MW floating installation of 60 × 15 MW turbines at 120 km in deep water, representing the next generation. The onshore baseline anchors the model to measurement; the three offshore scenarios are projections whose assumptions are stated in Sections 2.3 and 3.1 and whose sensitivity is quantified in Section 5.3.

Because the offshore scenarios cannot be validated against fleet-scale data of their own, the onshore anchoring is complemented by an indirect, qualitative benchmark against operating offshore farms. The S2–S4 availabilities produced by the model (96.7%, 94.9%, and 89.2%) are broadly consistent with the order of magnitude reported for operating offshore wind farms. For example, availability values in the mid-to-high 90% range are representative of well-performing offshore wind assets, although project-level availability can be affected substantially by turbine, electrical, and export-system outages ([19], [20]). The lower S4 value is therefore interpreted as a conservative projection for a far-from-shore floating configuration, rather than as a directly validated estimate. This interpretation is particularly relevant because the operational evidence base for floating offshore wind remains considerably smaller than that for fixed-bottom projects; Hywind Scotland, commissioned in 2017 with only 30 MW of capacity, remains one of the principal long-term operational references for floating wind [22], [23]. Similarly, the modelled net capacity factors (38.8–43.4%) fall within the range observed for modern offshore wind in the UK, where the historical offshore-wind fleet has generally achieved load factors around 40%, while newer projects can achieve substantially higher values as turbine size and technology improve [24]. This comparison is a plausibility check, not a validation: agreement in order of magnitude indicates that the model does not produce obviously unphysical values, but it does not establish that the S3 and S4 projections are correct. The comparison is therefore used only to demonstrate that the projected reliability and energy-performance ranges are technically plausible given the available evidence.

5. Results

5.1. Performance comparison across scenarios

Explica un poco lo que estás obtenidendo aquí en esta Table 3. Mi propuesta…

Table 3 summarizes the main technical and economic results from the 10,000-iteration Monte Carlo simulations. Availability decreases from 96.7% in S2 to 89.2% in S4, while the net capacity factor increases from 38.8% to 43.4% and annual energy production rises by 591%. Despite the lower availability and higher failure rate, LCOE falls from 119 to 65.8 €/MWh (− 45%), driven by increased energy yield and economies of scale. CapEx and O&M also decrease by 44% and 46%, respectively, although O&M remains a stable component of LCOE at around 34–35%. Overall, the results indicate that the economic benefits of larger, more productive configurations outweigh their reliability penalty.

O&M costs represent 34–35% of total offshore LCOE across the offshore scenarios, confirming that maintenance optimisation remains a principal lever for cost reduction even as CapEx per MWh falls with technology maturation. Availability declines from S2 to S4 (from 96.7% to 89.2%) as failure rates scale with turbine size and access becomes harder, but this is more than offset by the higher energy yield of the larger machines, so the net capacity factor still rises and LCOE still falls.

The 591% rise in annual energy from S1 to S4 in Table 3 is dominated by the increase in installed capacity (145.5 → 900 MW, a factor of 6.19) and is not a claim of improved yield per swept area. Normalising energy by installed capacity, the effect reduces to the change in net capacity factor (27.6% → 43.4%): most of the additional energy is more and larger turbines, with a smaller genuine efficiency gain from higher hub heights, better siting and improved wake behaviour per MW. Reporting the raw percentage without this decomposition would overstate the technological contribution, so both figures are given.

Table 3
Modelled performance metrics for all scenarios (10,000 Monte Carlo iterations, seed 2027). LCOE and availability are reported with their Monte Carlo 90% intervals.

Metric	S1	S2	S3	S4	Δ S2→S4
Availability (%)	99.7	96.7 [95.8–97.4]	94.9	89.2 [86.9–91.2]	−7.5 pp
Net capacity factor (%)	27.6	38.8	43.1	43.4	+ 4.6 pp
Annual energy (GWh)	351	495	2414	3421	+ 591%
LCOE (€/MWh)	54	119 [113–125]	79.5 [74–86]	65.8 [60–72]	−45%
CapEx (€/MWh)	38.2	77.8	51.6	43.5	−44%
O&M (€/MWh)	15.7	41.2	27.8	22.4	−46%
O&M share of LCOE (%)	29.2	34.6	35	34	—
Aggregate failure rate (f/turb/yr)	1.39	4.33	6.29	7.7	—

The LCOE figures carry non-trivial uncertainty. Propagating the Monte Carlo availability distribution alone yields 90% intervals of approximately 119 [113–125] €/MWh for S2, 79.5 [74–86] €/MWh for S3 and 65.8 [60–72] €/MWh for S4. These bands reflect only the simulated availability variability at fixed CapEx and WACC; the wider uncertainty from the exogenous CapEx and financing inputs is larger still and is characterised through the sensitivity analysis of Section 5.3 rather than through the Monte Carlo, because those inputs are deterministic assumptions rather than sampled quantities. The intervals overlap only marginally between adjacent scenarios, so the declining trajectory is robust to availability uncertainty but should not be read as three sharply separated point values. The decomposition of the LCOE into its capital and operating components across the four scenarios is shown in Fig. 1.

5.2. LCOE evolution and industry benchmarks

The S4 result of 65.8 €/MWh lies within the range recently reported for floating wind in 2030–2035 (EUR 52–76/MWh) [4], and the S3 result of 79.5 €/MWh is consistent with NREL and IRENA projections for 2030 fixed-bottom offshore wind under moderate technology-advancement assumptions [2, 3]. This literature agreement is used only to argue plausibility: falling within the published range shows the results are not absurd, but it is not a validation of the model, since a result can be plausible and still rest on incorrect assumptions. The floating-projection band shown in Fig. 2 is taken directly from the 52–76 €/MWh range of Lerch et al. [4] and the NREL floating trajectory [2]; it is a literature envelope, not a fitted confidence interval. S4 therefore represents an optimistic but literature-consistent scenario conditional on high wind resource, efficient O&M and a competitive cost of capital.

5.3. Sensitivity analysis

The dominance of WACC (a spread of 17.4 €/MWh over a 5–9% range; Table 4) confirms that financial de-risking—through policy stability, revenue mechanisms and track-record accumulation—is the single most effective LCOE-reduction pathway. Capacity factor and CapEx follow, with turbine reliability fourth: reliability matters, but at next-generation scale its LCOE leverage is smaller than that of the cost of capital, mirroring the variance-decomposition finding of the companion reliability work.

Table 4
One-at-a-time sensitivity of S4 LCOE (base = 65.8 €/MWh).

Parameter	Low LCOE (€/MWh)	High LCOE (€/MWh)	Range (€/MWh)
WACC	57.4	74.8	17.4
Capacity factor	59.8	73.1	13.3
CapEx	59.3	72.3	13
Failure rate	61.7	70.2	8.4
Vessel day-rate	62.5	69.2	6.7
Project lifetime	64	68.6	4.6
Turbine rating	63.6	68	4.4
Distance to shore	64.7	66.9	2.2

The tornado analysis (Fig. 3) is a one-at-a-time (OAT) screening: each parameter is varied over its plausible range with the others held at their base values. OAT is transparent and directly interpretable, but it does not capture interactions between parameters (for example, a high WACC combined with a low-capacity factor) and can understate the true spread when parameters are correlated. It is retained here as a first-order screening to rank the drivers, which is its appropriate use; a global variance-based sensitivity analysis (e.g. Sobol indices), which would apportion the LCOE variance among the inputs and their interactions, is identified as a natural extension and is the method used for the formal uncertainty budget in the companion reliability study.

The fourth-place ranking of turbine reliability should not be interpreted as diminishing the importance of the reliability sub-model. Its lower position in the LCOE sensitivity analysis is primarily a consequence of the strong direct influence of WACC and CapEx, which enter the LCOE formulation through the capital-recovery term and are imposed exogenously in the model. By contrast, reliability affects LCOE indirectly through turbine availability, O&M requirements, and net energy yield. Moreover, improved reliability can contribute to lower financing costs by reducing perceived project risk. The sensitivity ranking should therefore be interpreted as indicating that macro-financial parameters exert the strongest direct influence on LCOE, rather than implying that reliability is unimportant for the technical and economic performance of offshore wind projects.

6. Discussion

Three results deserve emphasis. First, the coupling of availability to the LCOE denominator matters: because availability is simulated rather than assumed, the model shows that the reliability penalty associated with larger, farther offshore floating machines is real (availability falls by roughly 7.5 percentage points from S2 to S4), yet is outweighed by energy-yield and economies-of-scale effects, so the cost trajectory still declines. The importance of explicitly accounting for downtime and reliability in offshore-wind techno-economic assessments has been demonstrated previously, particularly for floating wind, where failures, repair times, metocean conditions, vessels and port logistics can have a substantial effect on LCOE ([25]). Second, the O&M share of LCOE is stable at about one third across the offshore scenarios, which locates maintenance strategy as a persistent lever rather than one that disappears as turbines mature. This is consistent with recent techno-economic analyses showing that O&M costs remain a substantial component of floating-wind LCOE and that downtime and maintenance strategy can materially influence project economics [25]. Third, the sensitivity ranking places the cost of capital above technical parameters, a conclusion that is robust to the scaling assumptions because WACC enters the LCOE formulation through the capital-recovery/fixed-charge rate. The NREL Annual Technology Baseline explicitly defines the capital recovery factor as a function of WACC and incorporates it into the fixed charge rate used in LCOE calculations; financial assumptions therefore directly affect the levelized cost of energy [3]. Sensitivity studies of floating offshore wind likewise identify the discount rate among the parameters with the strongest influence on LCOE [26].

These numbers present crucial engineering and investment implications that warrant detailed examination. The 7.5-percentage-point availability drop to 89.2% in Scenario 4 (S4) is not inherently disqualifying. At the modeled net capacity factor and cost of capital, the scenario clears at 65.8 €/MWh —a Levelised Cost of Energy (LCOE) well within the competitive clearance bounds observed in recent European floating wind commercial auctions ([27], [28]). Consequently, an investor evaluating these parameters would deem the project financeable, provided the underlying reliability and Operations and Maintenance (O&M) assumptions hold. But the same figures show where the risk concentrates. Because roughly a third of the cost of energy is O&M ([20]; [3]); and because the availability loss is driven largely by weather-access and tow-to-port downtime rather than by the intrinsic failure rate ([29]), the design and logistics decisions that most improve S4 economics are those that reduce access dependence—larger service-operation vessels with wider weather windows ([30]), on-site major-component exchange to avoid tow-to-port, condition monitoring that converts unplanned major failures into planned interventions ([31])—rather than marginal reductions in component failure rate, a conclusion consistent with earlier O&M-strategy optimisation for offshore farms [14]. An engineer reviewing Table 3 should therefore conclude that at 15 MW floating scale the binding constraint is accessibility and logistics, not component reliability in isolation, and that the availability figure is acceptable only under an O&M strategy specifically designed for far-from-shore floating operation.

Placed against comparable studies, the trajectory is consistent, but the framing differs. Cost reviews such as the NREL Cost of Wind Energy Review [3] and the floating-installation modelling of Lerch et al.

[4] report LCOE ranges of the same order for the 2030–2035 horizon, and reliability studies for next-generation turbines [15, 12] report major-replacement rates consistent with the aggregate failure rates used here. The distinction of the present work is not a lower headline number—several published studies reach similar figures—but that the number is produced by a single model in which the availability that drives it is simulated from the same failure and repair assumptions, so the reliability, yield and cost figures are mutually consistent by construction rather than assembled from separate sources. Where published studies fix availability as an input, the present model derives it, which is what allows the reliability-versus-cost trade-off in S4 to be reasoned about explicitly.

Several technology trends are not captured and would, if realised, accelerate the decline further: HVDC offshore transmission can reduce electrical losses relative to HVAC for long distance off-shore transmissions, with reported losses to 1–2% per 100 km for modern VSC-HVDC systems comparted to 3–4% for HVAC systems ([32], [33]. Autonomous and drone-based inspection could further reduce inspection costs, vessel requirements and downtime, although the economic benefit depends on the degree of autonomy and the inspection task considered [34]. In parallel, physics-based design and uncertainty-quantification approaches such as those developed within the HIPERWIND project have demonstrated potential LCOE reductions of approximately 5–10% through improved probabilistic design, reliability modelling and uncertainty characterisation [16]. Integrating such physics-based approaches with the data-driven methods developed in this article series, and replacing the analytical Jensen wake model with Gaussian or more advanced deep-array formulations, therefore represents a natural extension of the present framework [21] .

Limitations. The S3 and S4 results depend on scaling laws that cannot yet be validated against fleet-scale data, and the availability model uses representative rather than site-specific vessel-access and met-ocean statistics. The wake model is analytical and, as discussed in Section 3.2, likely optimistic in the interior of very large farms. The exposure-dependent reliability estimates should therefore be read as conservative projections; as real 8–15 MW operating records accumulate, the scaling multipliers of Section 3.1 should be replaced by measured values and the wake treatment upgraded.

Future work. Beyond replacing synthetic scaling with measured next-generation rates and upgrading the wake and sensitivity methods, a natural direction is to keep the model current as evidence accumulates. A lightweight retrieval pipeline could periodically surface new cost and reliability figures from public sources (NREL, IRENA, IEA) for human-in-the-loop review and controlled updating of the model parameters, preserving the original results as an immutable baseline. This “living” updating mechanism is an engineering convenience for maintaining the platform, not a scientific claim of the present article, and is noted here only as a direction rather than a result.

7. Conclusions

This article has presented a modular, reproducible platform for predicting offshore wind-farm performance across turbine generations, coupling reliability, energy and cost in a single internally

consistent model. The principal conclusions map directly onto the four contributions of Section 1:

(1) The integrated model, validated against real onshore data from 97 turbines and driven by a single fixed seed, produces an internally consistent set of performance metrics in which availability, capacity factor and LCOE are derived from the same variables; the evolutionary trajectory from S2 (119 €/MWh) through S3 (79.5) to S4 (65.8), each with its Monte Carlo interval, demonstrates a plausible 45% LCOE reduction pathway.

(2) The LCOE decline decomposes into a fall in CapEx per MWh and O&M efficiency gains; O&M nonetheless remains 34–35% of offshore LCOE, confirming reliability engineering and predictive maintenance as persistent R&D priorities alongside capital-cost reduction, and the discussion shows that at floating scale the binding constraint is access and logistics rather than component reliability alone.

(3) The Monte Carlo sensitivity analysis identifies WACC, capacity factor and turbine reliability as the three dominant LCOE drivers, in that order, pointing to distinct intervention pathways: financial de-risking, site and technology selection, and AI-driven predictive maintenance respectively.

(4) Next-generation uncertainty is treated transparently: the 8–15 MW scaling laws are stated as assumptions, cross-checked against sparse published evidence and bounded by sensitivity analysis rather than asserted, so that the projections remain defensible as fleet-scale data emerge.

The platform is offered as a transparent, fully reproducible baseline for reasoning jointly about reliability, energy and cost in current and next-generation offshore wind, serving both academic and industrial audiences engaged in the offshore-wind cost-reduction effort.

Declarations

Clinical trial number

Clinical trial number: not applicable.

Funding

The authors received no specific funding for this work.

Author Contribution

A.G.P. designed and implemented the simulation platform, performed the reliability, energy and economic modelling, ran the Monte Carlo and sensitivity analyses, prepared all figures and tables, and wrote the main manuscript text. M.A.P.P. supervised the research, contributed to the methodology and

interpretation of results, and critically revised the manuscript. All authors reviewed and approved the final manuscript.

Data availability

The simulation code, the scenario parameter files and the fixed random seed (2027) required to reproduce every figure, table and reported number in this article are available from the corresponding author upon reasonable request. The associated resources are organised by sub-model—reliability scaling, Monte Carlo availability, energy production and economic evaluation—and each script regenerates its associated outputs directly from the seed, without path edits or non-standard dependencies, so that every table and figure can be reproduced from a single command.

References

1. Global Wind Energy Council (2024) Global Offshore Wind Report 2024, GWEC, Brussels, Belgium
2. Fuchs R, Duffy P, Shields M et al (2024) The Cost of Offshore Wind Energy in the United States 2025–2050, NREL/TP-5000-88988. https://doi.org/10.2172/2429410
3. Stehly T, Duffy P, Mulas Hernando D (2025) Cost of Wind Energy Review: 2024 Edition, NREL/TP-5000-91775. https://doi.org/10.2172/2479271
4. Lerch C, Aguiló S et al (2024) Modelling the installation of next-generation floating offshore wind farms. Appl Energy 373. https://doi.org/10.1016/j.apenergy.2024.123849
5. Boston Consulting Group (2025) Win-Cost: A Playbook for Offshore Wind Viability, BCG
6. Walgern J, Fischer K, Kolios A (2025) Reliability and O&M key performance indicators of onshore and offshore wind turbines. Wind Energy Sci. https://doi.org/10.5194/wes-2025-212
7. Carroll J, McDonald A, McMillan D (2016) Failure rate, repair time and unscheduled O&M cost analysis of offshore wind turbines. Wind Energy 19(6):1107–1119. https://doi.org/10.1002/we.1887
8. Artigao E, Martín-Martínez S, Honrubia-Escribano A, Gómez-Lázaro E (2021) Failure rate survey of wind turbines located in Spain. IET Renew Power Gener 15. https://doi.org/10.1049/rpg2.12028
9. Li M, Jiang X, Carroll J, Negenborn R (2022) A Bayesian network for reliability assessment of floating offshore wind turbines. Reliab Eng Syst Saf 228:108777. https://doi.org/10.1016/j.ress.2022.108777
10. Hosseini SA, Toubeau R, Vallée F (2025) Wind-farm reliability: concepts, databases and fault-tree methods. Renew Sustain Energy Rev 211. https://doi.org/10.1016/j.rser.2024.115243
11. Davison J (2015) Operation and Maintenance Modelling of Offshore Wind Turbines, Ph.D. thesis, Robert Gordon University
12. Jenkins B, Belton I, Carroll J, McMillan D (2024) Estimating major replacement rates in next-generation offshore wind turbines using structured expert elicitation. Wind Energy. https://doi.org/10.1002/we.2905

13. The Crown Estate (2013)
14. Karyotakis A (2011) On the Optimisation of Operation and Maintenance Strategies for Offshore Wind Farms, Ph.D. thesis, University College London
15. Donnelly O, Anderson F, Carroll J (2024) Failure rate and O&M cost impact for next-generation offshore turbines. Wind Energy 27. https://doi.org/10.1002/we.2885
16. HIPERWIND Consortium (2024) Advanced probabilistic models for offshore wind cost reduction (H2020 grant 101006689)
17. Jensen NO (1983) A note on wind generator interaction, Risø National Laboratory, Risø-M-2411
18. Katic I, Højstrup J, Jensen NO (1987) A simple model for cluster efficiency, in Proc. EWEC'86, vol. 1, pp. 407–410
19. Pfaffel S, Faulstich S, Rohrig K (1904) Performance and reliability of wind turbines: A review, Energies, 10(11), 2017. https://doi.org/10.3390/en10111904
20. Dao C, Kazantsev B, Glinka M (2019) Wind turbine reliability data review and impacts on levelised cost of energy. Wind Energy 22. https://doi.org/10.1002/we.2404
21. Bastankhah M, Porté-Agel F (2014) A new analytical model for wind-turbine wakes. Renew Energy 70:116–123. https://doi.org/10.1016/j.renene.2014.01.002
22. Equinor HS (2026) The 30-MW, five-turbine project has operated since 2017 and is described by. Equinor as the world's first floating wind farm
23. Pimm AJ et al (2018) Technological and economic analysis of floating offshore wind farms. Renew Energy 128:212–224
24. Parliament UK (2024) Written Question 14358: Offshore wind capacity factors, Reported UK offshore-wind load factors were 39.9%, 40.4%, 45.7%, 37.4% and 40.7% for 2018–2022, respectively
25. O&M-aware (2024) techno-economic assessment for floating offshore wind farms: A geospatial evaluation off the North Sea and the Iberian Peninsula. Appl Energy 371:123684. https://doi.org/10.1016/j.apenergy.2024.123684
26. Shadman M et al (2018) Sensitivity analysis on the levelized cost of energy for floating offshore wind farms. Int J Energy Res
27. Johnston B, Foley A, Tytherleigh N (2020) LCOE reduction pathways for floating offshore wind. Renew Energy 160:1302–1314. https://doi.org/10.1016/j.renene.2020.06.136
28. Stehly T, Duffy P (2021) 2020 Cost of Wind Energy Review, National Renewable Energy Laboratory (NREL), Tech. Rep. NREL/TP-5000-81209. https://doi.org/10.2172/1838135
29. Rinaldi G et al (2021) On the impact of weather conditions on the maintenance of floating offshore wind farms. Ocean Eng 237:109581
30. O'Connor A, Lynch M, McKeogh E (2022) Wave-walkway interaction and access limits for offshore wind vessel operations. IEEE Trans Sustain Energy 13(3):1420–1431
31. Tchakoua P, Wamkeue R, Ouhrouche M, Slaoui-Hasnaoui F, Tameghe TA, Ekemb G (2014) Wind turbine condition monitoring: Tech Rev Surv Energies 7(7):4498–4543.

https://doi.org/10.3390/en7074498

32. Negra NB, Todorovic J, Ackermann T (2006) Loss evaluation of HVAC and HVDC transmission solutions for large offshore wind farms. Electr Power Syst Res 76(11):916–927. https://doi.org/10.1016/j.epsr.2005.11.004
33. Abo-Khalil AG, Hassan A, Sayed K (2026) HVDC systems and renewable energy in a comparative study of technologies and applications, e-Prime – Nexus Electr. Electron. Intell. Eng., 17, 201167. https://doi.org/10.1016/j.eprime.2026.201167
34. Kabbabe Poleo K, Crowther WJ, Barnes M (2021) Estimating the impact of drone-based inspection on the Levelised Cost of Electricity for offshore wind farms. Results Eng 9:100201. https://doi.org/10.1016/j.rineng.2021.100201

Figures

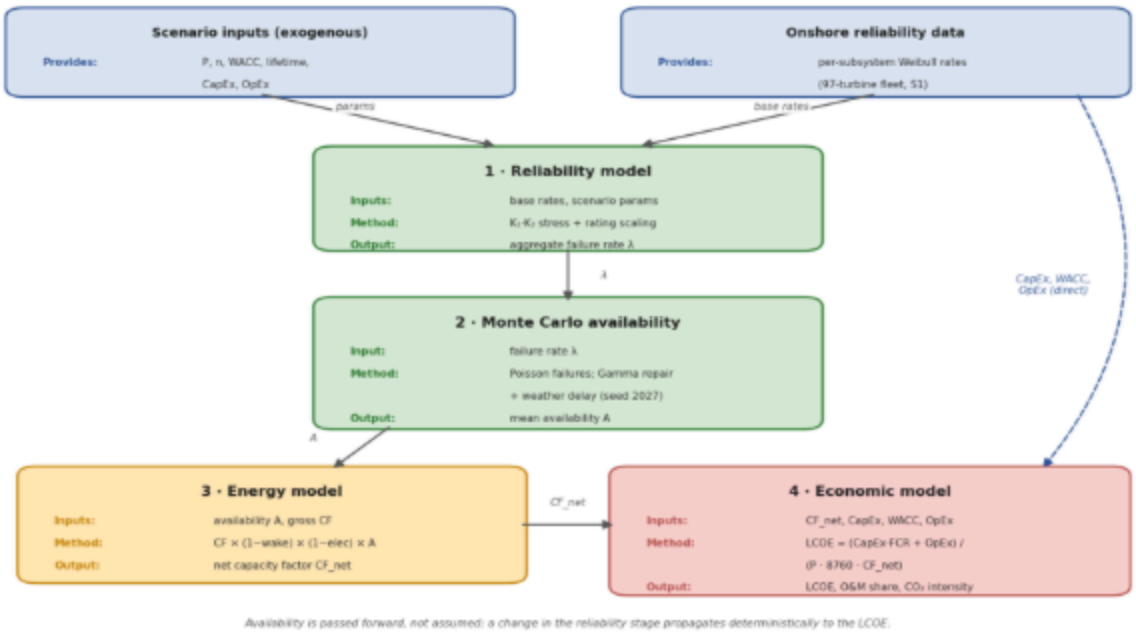


Figure 1

Figure A. Model workflow. Scenario inputs and onshore reliability data feed the reliability stage; the resulting failure rates drive the Monte Carlo availability model, whose mean availability enters the energy stage to form the net capacity factor; this and the exogenous CapEx, OpEx and WACC produce the LCOE. Availability is propagated forward rather than assumed.

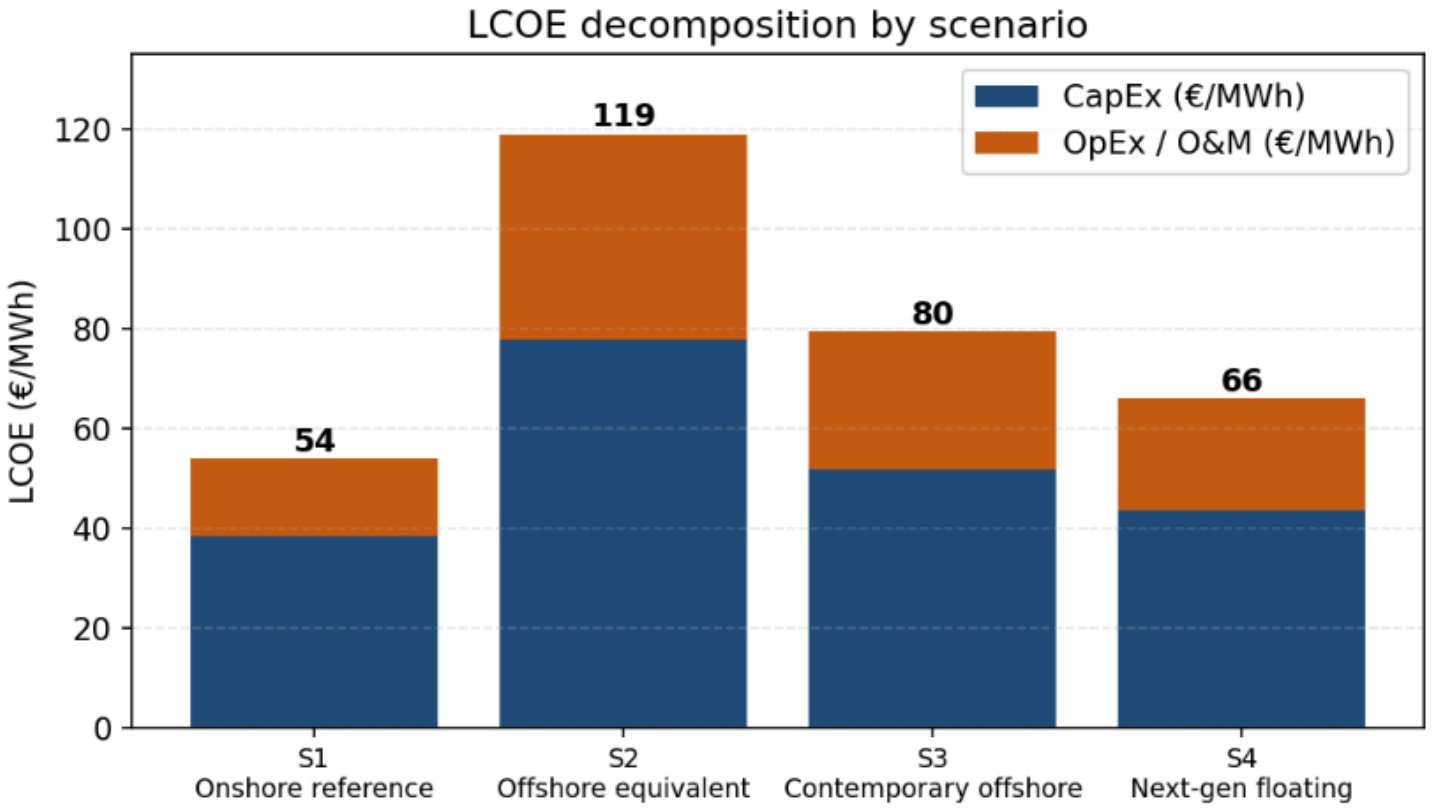


Figure 2

Figure 1. LCOE decomposition by scenario. The reduction from S2 to S4 is driven by lower CapEx per MWh and O&M efficiency gains, both enabled by larger turbines and economies of scale.

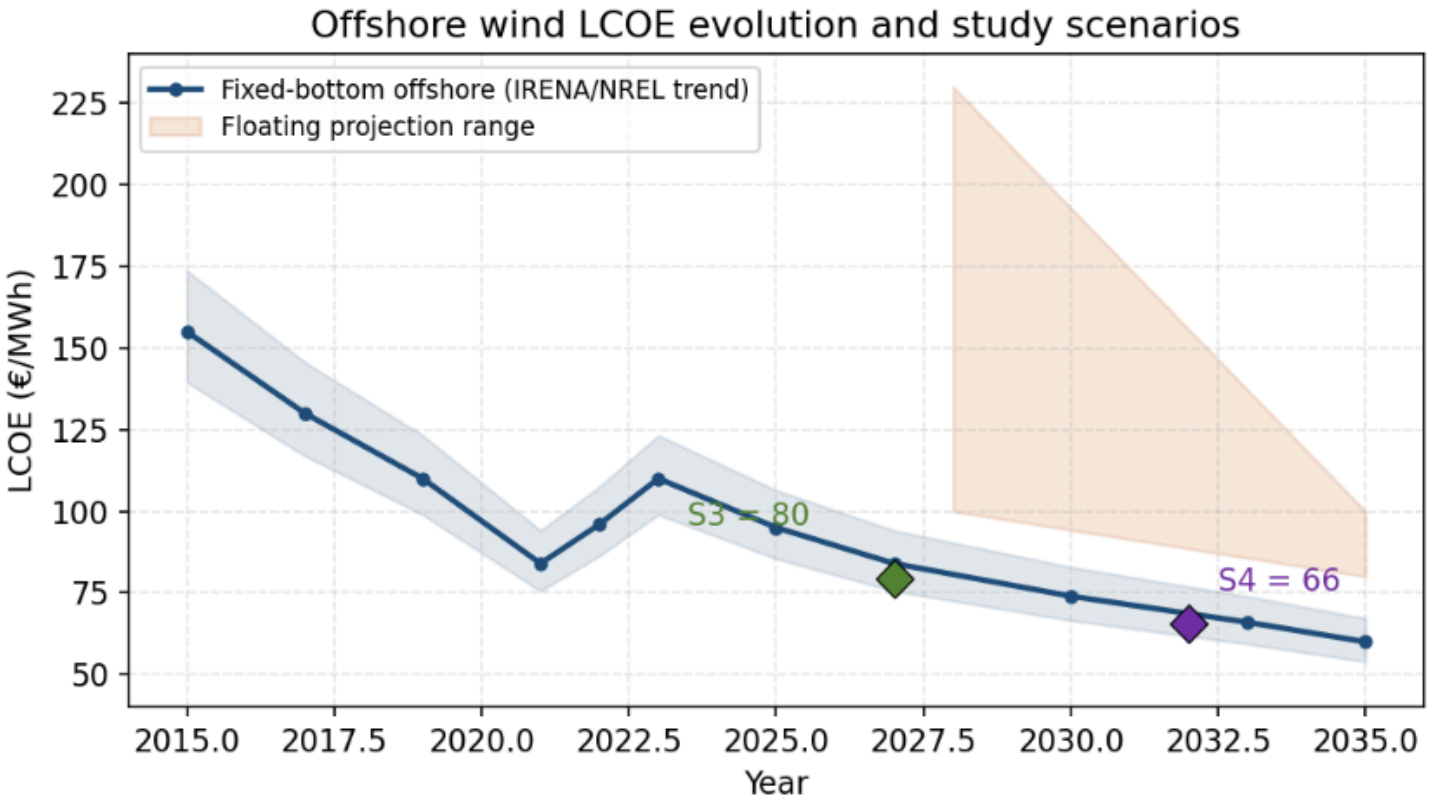


Figure 3

Figure 2. Offshore wind LCOE evolution with the fixed-bottom trend (IRENA/NREL) and the S3 and S4 scenario results positioned on the trajectory. The 2022–2023 inflationary reversal is visible before the projected decline.

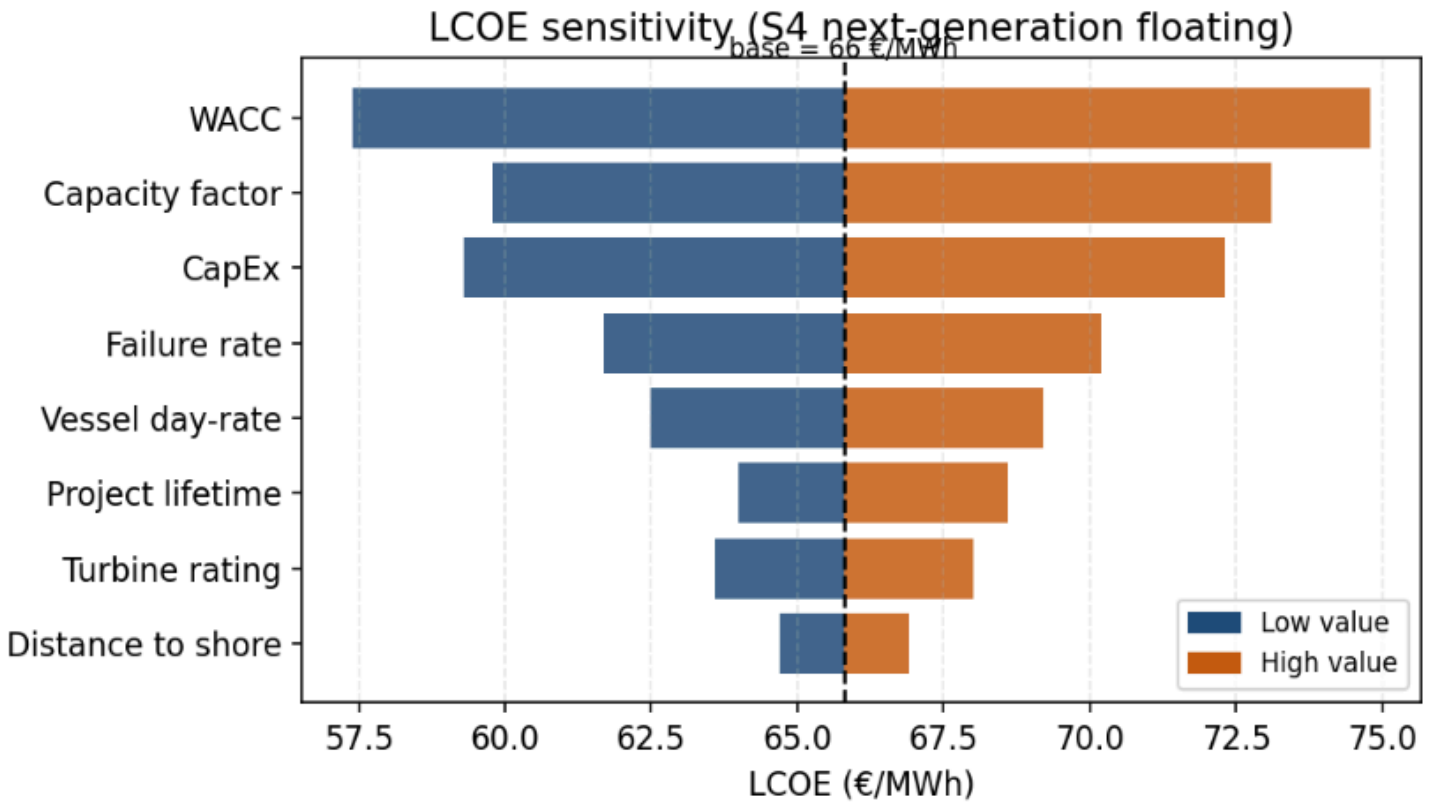


Figure 4

Figure 3. LCOE sensitivity tornado diagram for the S4 next-generation scenario (base = 65.8 €/MWh). WACC dominates, followed by capacity factor and CapEx, with turbine reliability fourth.