

Article

Resilience Aware Digital Twin Co-Optimisation of Off-shore Wind Spare Parts, Vessel Scheduling and Hybrid Energy Storage

Ashish Kumar Routray^{1*}, Swagatika Dehury¹, Mahottak Kumar², Sudhir Kumar Sahu¹

¹ Department of Statistics, Ravenshaw University, Cuttack, Odisha 753003, India

² Department of Civil Engineering, Synergy Polytechnic, Bhubaneswar 751032, India

* **Corresponding author:** Ashish Kumar Routray, ashish2025@ravenshawuniversity.ac.in

ABSTRACT

Offshore wind farms are capital-intensive assets and remote, weather-constrained logistics systems whose availability depends jointly on component degradation, spare-part availability, vessel access, and export capacity. Existing studies typically address these dimensions separately, with vessel-routing models assuming fixed spare availability, inventory models using exogenous lead times, and storage-sizing models assuming fixed asset availability. This separation can underestimate costs and overestimate resilience under increasingly severe weather. This study develops an integrated, resilience-aware framework that co-optimizes a two-echelon spare-parts system, condition-based maintenance triggered by a digital-twin remaining-useful-life estimator, heterogeneous vessel scheduling under stochastic metocean access windows, and hybrid battery–electrolyser buffer sizing and dispatch coupled with a voltage-source converter high-voltage direct current export link. Component degradation follows a gamma process with Bayesian drift updating, while significant wave height follows a Markov-modulated Weibull process with storm severity-adjusted regime intensities. The model is formulated as a two-stage distributionally robust mixed-integer program using a Wasserstein ambiguity set and a Conditional Value at Risk objective for energy not served. Analytical results establish convexity of expected cost in base-stock level, monotonicity of the optimal maintenance trigger with storm severity, and sub modularity of the resilience index in storage design. An accelerated integer L-shaped decomposition achieves a 0.8% optimality gap. For a 1.2-GW, 100-turbine farm, joint optimization reduces annual O&M cost from €4.90 million to €1.79 million, increases resilience from 0.41 to 0.87, and reduces 95% CVaR of energy not served by 44.6%.

Keywords: spare parts inventory; condition-based maintenance; hydrogen coupling; distributionally robust optimisation; resilience; VSC-HVDC

ARTICLE INFO

Received: 14 August 2026

Accepted: 15 September 2026

Available online: 24 September 2026

COPYRIGHT

Copyright © 2026 by author(s).

Journal of Offshore Renewable Energy and Grid Integration is published by Cranford Press Limited.

This work is licensed under the Creative Commons Attribution-Noncommercial 4.0 International License (CC BY-NC 4.0).
<https://creativecommons.org/licenses/by-nc/4.0/>

1. Introduction

1.1. Motivation

Global off-shore wind capacity has moved decisively from demonstration to industrial scale, and with that transition the binding constraint on project value has migrated from technology risk to operational risk. For a fixed bottom farm in the North Sea, planned and unplanned maintenance together account for roughly a quarter of the levelized cost of energy, while for floating installations and for sites exposed to tropical cyclones, the share is materially higher. The dominant driver is not the intrinsic failure rate of drivetrains, blades, or converters but the logistics of responding to failure, where a gearbox replacement requires a heavy lift vessel, a spare unit physically present in the supply chain, and a contiguous weather window of sufficient duration. When any one of these three is missing, a repair that occupies

four days of vessel time can generate ninety days of lost production.

This coupling has three consequences that conventional planning practice handles poorly.

First, lead time is endogenous. The delay between failure detection and restoration is not a parameter but an outcome of stock position, vessel availability, and sea state, representing a recognition that overturns the classical separation between inventory policy and maintenance policy.

Second, risk is fat tailed and correlated. Storms simultaneously raise failure hazards, suppress accessibility, and depress export capability, meaning that the joint distribution, rather than the marginal distributions, determines downside exposure.

Third, the mitigation set is now broader than logistics. Off-shore energy hubs increasingly colocate battery storage and electrolyzers, so a fraction of the value at risk during an outage can be recovered through temporal arbitrage or conversion to hydrogen rather than through faster repair.

1.2. Research gap

The literature has evolved along four largely disjoint tracks. Maintenance routing studies optimize technician and vessel deployment under weather constraints but assume spares are available on demand. Spare parts inventory studies for capital goods derive elegant base stock and $(S - 1, S)$ policies but treat the repair lead time as exogenous and stationary. Storage sizing studies optimize battery and hydrogen capacity against generation variability but assume a fixed availability factor. Resilience studies quantify performance loss under extreme events but rarely embed a decision model capable of procuring resilience.

The consequence of this separation is not merely aesthetic. Because the four subsystems are coupled through the same latent metocean state, optimising them sequentially produces solutions that are individually optimal and jointly fragile. A base stock level tuned to an average lead time is systematically under provisioned in the winter regime, a CBM threshold tuned to reliability alone triggers interventions precisely when vessels are least accessible, a storage system sized for daily variability provides little protection against a five day outage.

1.3. Contributions

This paper closes that gap with the following contributions.

- a) **An integrated model:** We formulate, to our knowledge for the first time, a single optimisation model that jointly determines the two-echelon spare parts base stock vector, the digital twin CBM trigger, the heterogeneous vessel fleet mix and dispatch, and the battery electrolyser sizing for an off-shore wind farm exporting through a VSC-HVDC link. Lead times, failure hazards and export capability are all endogenous functions of a common Markov modulated metocean process.
- b) **A degradation to logistics coupling:** Component health is modelled as a gamma process whose drift is Bayesian updated from digital twin sensor streams, while the posterior RUL distribution enters the inventory model as a forecast of demand, converting an unpredictable corrective demand stream into a partially observable, advance demand information process. In the case study of Section 6 this coupling is operationalised through the fleet-level reduction Equation (8), so that the quantitative results depend on the twin through two scalars, instrumentation coverage ϕ and the RUL-quantile trigger τ , whose separate and joint value is isolated in Sections 6.3 and 6.6.
- c) **Distributionally robust resilience:** The second stage is evaluated over a Wasserstein ball around the empirical metocean distribution with a CVaR objective on energy not served, so that the solution is protected against both scenario sampling error and non-stationarity in storm climatology.
- d) **Structural results:** We prove convexity of expected cost in the base stock level, monotonicity of the optimal CBM trigger in storm severity, and sub-modularity of the resilience index in the storage vector,

the last yielding a $(1 - 1/e)$ approximation greedy algorithm that we use to warm start the exact method.

e) **A scalable solution method:** An integer L-shaped decomposition with Pareto-optimal cuts, scenario aggregation and a valid inequality family solves 1,500 scenario instances of a 100-turbine farm within 0.8% of optimality in under an hour.

f) **A dual geography case study:** The model is parameterised on North Sea statistics and re-calibrated to the Bay of Bengal cyclone regime off Odisha, India a context of direct relevance to the 37 GW of off-shore wind zoned by the Indian government off Gujarat and Tamil Nadu^[1], while showing that the optimal design shifts sharply toward hydrogen coupling as cyclone intensity rises.

Situating these six contributions against the four literature streams reviewed in Section 2 makes the gap concrete, and we provide that grounding here rather than through a literal merger of Sections 1 and 2.

Contribution 1, which represents the integrated model, responds directly to the disjointness documented across Sections 2.1 through 2.4. Specifically, operations and maintenance routing studies reviewed in Section 2.1 hold spare availability exogenous, spare parts studies in Section 2.2 hold lead time exogenous, digital twin and condition-based maintenance studies in Section 2.3 rarely reach the inventory decision, and storage sizing studies in Section 2.4 hold availability exogenous. **Table 1** quantifies this column by column.

Contribution 2, concerning the degradation to logistics coupling, extends the gamma process remaining useful life machinery reviewed in Section 2.3 into a demand forecasting input for the base stock policy of Section 2.2, closing the condition-based maintenance meets inventory gap explicitly flagged there as thin.

Contribution 3, focusing on distributionally robust resilience, applies the Wasserstein Conditional Value at Risk machinery reviewed in Section 2.5, which was previously used mainly for terrestrial power system planning, to a setting where the ambiguity is compounded by a decision dependent lead time, representing a coupling absent from that literature.

Contribution 4, detailing structural results, extends the convexity results established for deteriorating item inventory systems in Section 2.2 to an endogenous, storm modulated lead time setting, and imports sub-modular maximization theory that the storage literature of Section 2.4 does not exploit.

Contribution 5, presenting the solution method, adapts integer L-shaped and Pareto-optimal cut machinery, which is standard in stochastic programming but to our knowledge previously unused in the off-shore operations and maintenance literature of Section 2.1, to the scale this integration produces.

Contribution 6, featuring the dual geography case study, responds to the near total absence across all four streams of parameterization against a cyclone exposed emerging market regime.

1.4. Organisation

Section 2 reviews the four literature streams and positions the contribution. Section 3 describes the system. Section 4 develops the mathematical model and structural results. Section 5 presents the solution methodology. Section 6 reports the case study and the response surfaces. Section 7 draws managerial and policy implications, and Section 8 concludes.

2. Literature review

2.1. Off-shore wind O&M and vessel routing

Research on off-shore O&M scheduling began with simulation-based availability assessment and has matured into optimisation^[2-4]. Contemporary formulations treat the problem as a stochastic vehicle routing problem with time windows in which the windows are themselves random, generated by a metocean process. Two modelling choices dominate, namely discretizing the sea state into accessible or inaccessible regimes governed by a Markov chain, or sampling significant wave height directly from a fitted parametric distribution and applying a persistence filter. Both approaches capture the essential nonlinearity that accessibility is a

threshold function of a heavy tailed, strongly autocorrelated process, but they typically hold the spare parts position fixed, treating a component as available whenever a repair is scheduled.

A parallel body of work addresses vessel fleet composition, contrasting crew transfer vessels, service operation vessels, and jack up or heavy lift vessels with respect to charter cost, transit speed, and wave height limits^[5,6]. The consensus is that service operation vessel-based strategies dominate for far shore farms, while the heavy lift vessel decision is dominated by the arrival process of major component failures, which is precisely the process that a digital twin can forecast and that spare parts policy can buffer. This observation motivates our integration.

2.2. Spare parts inventory for capital intensive assets

The spare parts stream descends from METRIC and its multi-echelon extensions, in which base stock levels for repairable items are set to minimize expected backorders subject to a budget. Applications to wind energy have adapted these models to account for the small population sizes, extreme unit costs and long repair pipelines characteristic of off-shore components. A persistent limitation is the exogeneity of the replenishment lead time: in the off-shore context the lead time is a convolution of manufacturing, onshore transport, port handling and weather dependent off-shore transfer, so treating it as a stationary random variable understates its seasonal variance by a wide margin.

A second relevant strand concerns deteriorating and non-instantaneously deteriorating items. Although developed largely for commercial rather than industrial contexts, this literature supplies the analytical machinery for stock that loses value or function while held directly applicable to off-shore spares such as sealed bearing assemblies, elastomeric couplings, lubricants and battery modules staged on an accommodation platform in a corrosive marine atmosphere. An inventory model for non-instantaneously deteriorating items under linear demand^[7] establishes the convexity conditions we exploit in Proposition 1. A three-warehouse formulation with ramp-type demand and partial backlogging^[8] provides the multi-location structure owned warehouse, rented warehouse, remote warehouse that maps naturally onto our onshore hub, port buffer and off-shore platform echelons, and it also incorporates carbon cap and trade, trade credit and inflation, which is directly relevant here because off-shore O&M vessels are diesel intensive and increasingly fall within emissions trading perimeters, so the carbon cost of an expedited HLV mobilisation is a first order term in the objective.

2.3. Digital twins, prognostics and condition-based maintenance

Digital twin architectures for wind turbines fuse SCADA, condition monitoring vibration spectra, strain gauges and, increasingly, remote sensed environmental fields into a continuously updated model of the physical asset^[9,10]. The prognostic layer typically employs either physics of failure models (Paris law crack growth, bearing L10 life), data driven regressors (recurrent networks, temporal convolution, transformer-based time series foundation models), or hybrid filters^[11]. Stochastic degradation processes Wiener, gamma and inverse Gaussian remain the workhorses for tractable RUL distributions, and the gamma process is preferred for monotone damage accumulation such as erosion, wear and corrosion.

The prognostic layer core statistical machinery, namely remaining useful life estimation from irregularly sampled, heterogeneous sensor streams, is established directly in the rotating machinery prognostics literature. Statistical data driven remaining useful life estimation is reviewed comprehensively in^[12], and transformer-based architectures for irregularly sampled off-shore drivetrain sensor data are demonstrated in^[13]. We adopt these as the primary basis for the digital twin fusion layer, with a Bayesian updating layer added to preserve calibrated uncertainty.

Where CBM meets inventory, the literature is thinner^[14]. Advance demand information generated by prognostics has been shown to reduce safety stock substantially^[15], but most treatments assume a deterministic

lead time and unconstrained repair capacity, both of which fail off-shore.

2.4. Storage, hydrogen coupling and grid integration

Battery energy storage co-located with off-shore substations serves three functions which are smoothing sub-hourly variability, providing synthetic inertia and fault ride through support to the VSC-HVDC converter^[16,17], and the function of interest here sustaining partial export during contingencies. Off-shore hydrogen production has moved from concept to pilot, with in-turbine, platform centralised and on-shore electrolysis architectures each presenting distinct efficiency and capital trade-offs. Platform centralised electrolysis is attractive when export cable capacity is the binding constraint, because it converts curtailed energy into a storable, pipeline transportable carrier.

Resilience oriented storage sizing has been studied for terrestrial microgrids, where the canonical performance loss trapezoid degradation, nadir, recovery provides an integral resilience metric^[18–20]. Transposing this to off-shore systems require accounting for the fact that the recovery limb is itself weather constrained, because unlike a terrestrial feeder, an off-shore asset cannot be repaired during the event.

2.5. Risk averse and distributionally robust optimisation in energy systems

CVaR^[21] and distributionally robust formulations^[22] are now standard in power system planning under extreme events. Wasserstein ambiguity sets^[23] are attractive because they admit tractable re-formulations and provide finite sample guarantees without assuming a parametric family an important property when the calibration data are storm records with only a few dozen severe realisations. Budget of uncertainty or Γ -robust formulations remain popular for their interpretability^[24], and we retain Γ as a tunable conservatism parameter in one of our response surfaces to allow direct comparison.

2.6. Gap analysis and positioning

Table 1 summarises the positioning of the present work against representative modelling capabilities.

Table 1. Capability comparison against the four literature streams.

Capability	O&M routing	Spare parts	Digital twin / CBM	Storage sizing	This paper
Weather dependent access windows	Yes	No	Partly	No	Yes
Endogenous replenishment lead time	No	No	No	No	Yes
Multi-echelon off-shore stocking	No	Yes	No	No	Yes
Deterioration of held spares	No	Partly	No	No	Yes
Prognostic advance demand information	No	Partly	Yes	No	Yes
Battery + electrolyser co-design	No	No	No	Yes	Yes
VSC-HVDC export constraint	Partly	No	No	Partly	Yes
Tail-risk (CVaR) objective	Partly	No	No	Partly	Yes
Distributional ambiguity	No	No	No	Partly	Yes
Structural optimality results	Partly	Yes	No	No	Yes

The distinguishing feature is the closure of the loop, where degradation drives spare demand, spare position and vessel access drive restoration time, restoration time drives availability, availability and storage drive delivered energy, and the economics of delivered energy feed back into the optimal degradation intervention threshold.

To substantiate the novelty language against the nearest competitors shown in **Table 1** concretely, Zhang et al.^[25] jointly optimize condition-based maintenance and spare parts inventory under weather constraints, representing the closest prior integration, but hold vessel fleet composition as well as storage and hydrogen sizing fixed and use an exogenous, weather independent lead time distribution rather than the endogenous

accessibility process of Equations (2) and (3). Papadopoulos et al.^[26] jointly consider accessibility, efficiency, and risk in maintenance scheduling, but do not model spare parts stock levels or storage co design. Furthermore, storage sizing studies^[27–29] optimize battery and electrolyser capacity against generation variability with availability treated as an exogenous input rather than an equilibrium outcome of the maintenance logistics subsystem.

We are not aware of a prior study that makes the replenishment lead time, the condition-based maintenance trigger, the vessel fleet, and the storage design simultaneous decisions coupled through a single latent metocean state, which is the specific sense in which the formulation is described as new.

3. System description and modelling framework

3.1. Physical system

We consider a farm of N turbines of rated power P^r located at distance d from an O&M base port, connected through inter-array cables to an off-shore substation platform (OSP) hosting a VSC-HVDC converter of rating $\bar{P}^{hv}$. Co-located with the OSP is an accommodation and logistics platform holding a forward stock of critical components and, in the design variants we evaluate, a battery of energy capacity E^b and power rating P^b , and a PEM electrolyser of rating P^{el} feeding a hydrogen compression and offtake system.

Three vessel classes are available, which are CTVs with wave height limit H^{ctv} for technician transfer, SOVs with limit H^{sov} and motion compensated gangways, and HLVs with limit H^{hlv} required for major component exchange. Charter can be secured on a seasonal basis (decided in the first stage) or on the spot market at a premium (second stage).

Components are partitioned into a set $\mathcal{K}$ of critical types gearbox, generator, main bearing, blade, pitch system, converter module each characterised by a degradation process, a unit cost, a physical envelope determining which vessel class can carry it, and a deterioration rate while held in off-shore storage (see **Figure 1**).




Figure 1. Integrated off-shore wind logistics storage grid system.

3.2. Decision architecture

The model has two stages. First stage (here and now, annual or seasonal horizon) decisions comprise:

- the two-echelon base stock vector $S = (S_k^{on}, S_k^{off})_{k \in \mathcal{K}}$,
- the CBM intervention trigger τ , expressed as a residual life threshold in days,
- the seasonal vessel charter vector $v = (v^{ctv}, v^{sov}, v^{hlv})$,
- the storage design (E^b, P^b, P^{el}) .

Second stage (wait and see, hourly resolution over the year) decisions comprise vessel dispatch, repair sequencing, emergency expediting, battery charge/discharge, electrolyser loading, HVDC export scheduling and curtailment.

Uncertainty ξ comprises the metocean trajectory, wind power realisation, degradation increments and spot electricity and hydrogen prices.

3.3. Digital twin information layer

The digital twin supplies, at each decision epoch, a posterior distribution of RUL for each instrumented component. Sensor coverage $\phi \in [0,1]$ denotes the fraction of the fleet under full condition monitoring while un-instrumented components revert to an age-based hazard. Coverage is treated as a scenario parameter rather than a decision, so that its marginal value can be quantified (Section 6.6).

Twin scope, sensor architecture and calibration

To state precisely what constitutes the digital twin in this formulation: the twin comprises the gamma degradation process Equation (4), the Bayesian updating of the shape-rate posterior on the observation history F_t , the resulting posterior predictive RUL distribution Equation (6), and the CBM trigger Equation (7) that consumes it. This is distinguished from generic condition-based maintenance with Bayesian updating on two points specific to the off-shore setting, not to the updating mechanism itself. First, the posterior RUL feeds a joint decision coupled at the fleet level through the shared shape parameter and through the base stock and vessel scheduling variables of Equations (14)–(22), rather than triggering an isolated replace/do-not replace decision per component. Second, sensor coverage ϕ is itself a first-class scenario parameter whose marginal economic value is quantified separately in Section 6.6, rather than assumed complete. Equations (4)–(7) are the twin in full; the manuscript does not claim additional twin machinery beyond this observation update decision loop, and the title and the **Table 2** row label should be read as denoting the integration of this prognostic loop with inventory and vessel logistics, not a broader digital-twin architecture such as a physics-based finite-element or CFD replica of the asset.

The observation model $Y_{ik}(t) = X_{ik}(t) + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, immediately following (5), is a reduced-form summary of a channel-fusion step rather than a raw sensor reading. For the ϕ fraction of the fleet under full condition monitoring, three channel types are fused into the scalar damage proxy X_{ik} : SCADA-derived operational severity (rotor speed, pitch angle, generator torque, gearbox oil temperature), sampled at the SCADA reporting interval of one to ten minutes; vibration spectra from accelerometers on the gearbox and main bearing, reduced to spectral kurtosis and envelope-band energy features at ten-minute intervals; and, for the subset of turbines carrying strain gauges on the blade root and tower base, strain-cycle counts aggregated hourly by rainflow counting. Each channel is first mapped, per component type k , to a component-specific damage indicator through a feature-to-damage regression fitted during commissioning; the fused indicators are then combined by inverse-variance weighting into the scalar X_{ik} used in Equations (4)–(6).

The noise variance in the observation equation is accordingly specified per component type k rather than per raw channel or as a single constant: it is calibrated from the residual variance of this fitted regression against borescope and lubricant-analysis ground truth collected during commissioning, and is re-estimated whenever a component type accumulates a new maintenance record with paired sensor and inspection data. The posterior is updated at the ten-minute sensor cadence, but the CBM trigger Equation (7) is evaluated only at the coarser decision epochs used throughout Sections 4–6, so intra-epoch observations refine the posterior continuously while the replacement decision itself is reviewed once per epoch; this separation is what keeps Equations (6) and (7) tractable at fleet scale. For the $(1-\phi)$ uninstrumented fraction, no channel fusion occurs and the component instead follows the age-based hazard used in the corrective-only and calendar-preventive benchmarks of **Table 2**, i.e. the failure-rate floor in the first bracketed term of Equation (8). Which components

fall into this fraction is governed by the retrofit sequencing of the accommodation platform's fibre backbone, and is treated in the case study as an exogenous scenario parameter rather than an optimised choice, consistent with the treatment of ϕ earlier in this section.

4. Mathematical model

4.1. Notation

Sets: $\mathcal{N} = \{1, \dots, N\}$ turbines, $\mathcal{K}$ critical component types, $\mathcal{T} = \{1, \dots, T\}$ hourly periods, $\mathcal{M} = \{1, \dots, 12\}$ months, $\mathcal{V} = \{ctv, sov, hlv\}$ vessel classes, Ω scenarios with probabilities p_ω , $\mathcal{E} = \{on, port, off\}$ inventory echelons.

Parameters: P^r turbine rating (MW), $\bar{P}^{hv}$ HVDC rating, c_k^h holding cost rate, c_k^s unit spare cost, c_j^p daily charter cost of class j , c_j^{sp} spot premium, $c_{t\omega}^e$ electricity price, c^{h_2} hydrogen price, c^{ens} penalty for energy not served, H^j wave height limit of class j , θ_k deterioration rate of spare k in off-shore storage, η^{ch}, η^{dis} battery efficiencies, η^{el} electrolyser efficiency (MWh/kg), α CVaR confidence, Γ uncertainty budget, ε Wasserstein radius.

Variables: First stage: $S_k^e \in \mathbb{Z}_+$, $\tau \in \mathbb{R}_+$, $v_j \in \mathbb{Z}_+$, $E^b, P^b, P^{el} \in \mathbb{R}_+$. Second stage (per scenario ω): $x_{ijt\omega}$ vessel routing binaries, $y_{ikt\omega}$ repair start binaries, $I_{kt\omega}^e$ inventory levels, $B_{kt\omega}$ backorders, $g_{t\omega}$ generation, $b_{t\omega}^{ch}, b_{t\omega}^{dis}$ battery flows, $h_{t\omega}$ electrolyser load, $q_{t\omega}$ HVDC export, $u_{t\omega}$ energy not served, ζ and ψ_ω CVaR auxiliaries.

4.2. Metocean process and accessibility

Significant wave height is modelled as a Markov modulated Weibull process. Let $Z_t \in \{1, 2, 3\}$ denote a latent regime (calm, moderate, storm) evolving as a discrete time Markov chain with generator $Q(m)$ that is month dependent. Conditional on $Z_t = z$ and month m ,

$$H_t^S \mid Z_t = z, m \sim \text{Weibull}(k_z(m), \kappa \lambda_z(m)) \quad (1)$$

where $\kappa \geq 1$ is a storm severity multiplier that inflates the scale parameter and serves as our climate stress lever. The instantaneous accessibility of class j is $\mathbb{1}\{H_t^S \leq H^j\}$, but a repair of duration Δ requires a persistent window. Assuming a first order persistence correction with exponent $\rho > 1$ calibrated from hindcast data, the monthly accessibility of class j is

$$\pi_{jm}(\kappa) = \left[\sum_z \varpi_z(m) \left(1 - \exp \left\{ - \left(\frac{H^j}{\kappa \lambda_z(m)} \right)^{k_z(m)} \right\} \right) \right]^\rho \quad (2)$$

with $\varpi_z(m)$ the stationary regime distribution in month m and equation (2) generates the response surface in Section 6.2. Two features matter for design as π is steeply increasing in the vessel wave limit (motivating SOV/HLV selection), and its seasonal amplitude is large enough that a single annual average is a poor sufficient statistic.

The expected waiting time for a class j window in month m follows from the geometric first passage argument,

$$\mathbb{E}[W_{jm}] = \frac{\Delta(1 - \pi_{jm})}{\pi_{jm}} + \frac{\sigma_H^2}{2\mu_H} \mathbb{1}\{\text{regime switching}\} \quad (3)$$

the second term being a diffusion correction that is material only in shoulder seasons.

4.3. Degradation, prognostics and the CBM trigger

Damage in component k of turbine i accumulates as a gamma process,

$$X_{ik}(t) \sim Ga(\nu_k \Lambda_i(t), \beta_k), \quad \Lambda_i(t) = \int_0^t \left(1 + \gamma_k \frac{P_i(s)}{Pr} \right) ds \quad (4)$$

so that damage accrues faster under high loading, with Λ_i an operational severity clock. Failure occurs at the first passage of X_{ik} across a threshold ℓ_k , giving remaining useful life

$$R_{ik}(t) = \inf\{s > 0: X_{ik}(t + s) \geq \ell_k\} \quad (5)$$

The digital twin observes X_{ik} with Gaussian noise, $Y_{ik}(t) = X_{ik}(t) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$, for the instrumented fraction ϕ of the fleet. Placing a conjugate gamma prior $\nu_k \sim Ga(a_0, b_0)$ on the shape rate and updating on the observation history $\mathcal{F}_t$ yields the posterior predictive RUL distribution

$$Pr(R_{ik}(t) \leq s \mid \mathcal{F}_t) = \int \frac{\Gamma\left(\nu \Delta \Lambda_i(t, s), \beta_k \left(\ell_k - \hat{X}_{ik}(t)\right)\right)}{\Gamma(\nu \Delta \Lambda_i(t, s))} d\Pi(\nu \mid \mathcal{F}_t) \quad (6)$$

A replacement is triggered for component (i, k) at the first epoch at which the q quantile of the posterior RUL falls below τ is given by,

$$\mathcal{A}_{ik}(\tau) = \inf\{t: \text{Quant}_q(R_{ik}(t) \mid \mathcal{F}_t) \leq \tau\} \quad (7)$$

Larger τ means earlier intervention and more lead time for logistics, fewer in-service failures, but more residual life discarded. The effective corrective failure rate under trigger τ and coverage ϕ is approximated by

$$\lambda^{eff}(\tau, \phi) = \lambda_0 \left[(1 - \phi) + \phi \left(\underline{a} + (1 - \underline{a}) e^{-(\tau - \tau_{min})/\tau_c} \right) \right] \quad (8)$$

with $\underline{a}$ the irreducible fraction of sudden failures. Total corrective spare demand is $D(\tau) = N \lambda^{eff}(\tau, \phi)$ per year.

A precise statement of what the numerical results of Section 6 depend on, and do not depend on, is needed here. Every quantity reported there, **Table 2** and Sections 6.3 through 6.6, is computed from the reduced-form failure intensity Equation (8), so the twin enters those results exclusively through the two scalars ϕ and τ . No result in Section 6 evaluates the posterior Equation (6) for an individual turbine-component pair, reports a component-level RUL value, or tabulates an individual firing of trigger Equation (7); Equations (6)–(7) instead discipline what τ means, a residual-life quantile rather than a calendar age, and what ϕ multiplies, the instrumented share of the fleet, while their fleet-level consequence for the case study is fully summarised by Equation (8). Setting $\phi = 0$ removes the twin from Equation (8): the coverage term multiplying the bracket vanishes, leaving the effective failure rate equal to the baseline rate λ_0 , independent of τ , since there is no condition information left for a threshold to act on; this is the age-based hazard underlying the Corrective only benchmark of **Table 2**, where no τ^* is reported for exactly this reason. The marginal value of installing the twin at all, and of widening its coverage, is quantified explicitly in Section 6.6, where the value of the stochastic solution rises from 4.0% at $\phi = 0$ to 11.2% at $\phi = 1$; that range is the full quantified contribution of the twin to the numerical results of this paper.

Because the case study's numerical dependence on the twin is mediated entirely by (ϕ, τ) , we reconsidered whether “digital twin” remains the appropriate term in the title, the **Table 2** row label and Contribution 2. We have retained it. τ is interpretable as days of residual life, rather than an arbitrary calendar age, only because of the posterior machinery of Equations (6)–(7); ϕ is not a free parameter but a scenario variable indexing instrumentation breadth, whose value is quantified on its own axis in Section 6.6; and the reduction from Equations (6)–(8) is a closed-form fleet-level marginalisation adopted for tractability at the scale solved in

Section 6, one hundred turbines and 1500 scenarios inside an integer L-shaped loop, rather than a departure from the twin model, in the same way that a hazard function rather than a full renewal-process simulation is standard practice when reliability models are embedded in large-scale inventory optimisation. To keep this scope visible at the two points a reader is most likely to check it, we relabel the **Table 2** row “Digital twin ($\varphi = 1$) + joint (proposed)” and add a clause to Contribution 2 identifying (φ, τ) as the mechanism through which the coupling is operationalised in the case study.

4.4. Two-echelon inventory with deteriorating spares

Let $I_k^{off}(t)$ denote off-shore on-hand stock. Between replenishments it obeys

$$\frac{dI_k^{off}(t)}{dt} = -\theta_k(t) I_k^{off}(t) - D_k(t), \quad \theta_k(t) = \begin{cases} 0, & t \leq t_k^d \\ \theta_k, & t > t_k^d \end{cases} \quad (9)$$

which follows the general gamma process deterioration framework of^{f[30]} and the maintenance linked spare parts control theory of^{f[31]}, instantiated here in the non-instantaneous deterioration functional form of^{f[7]}. A spare held on the platform suffers no functional loss for a protection period t_k^d (nitrogen purged container, desiccant integrity) and thereafter degrades at rate θ_k . Solving Equation (9) with $I_k^{off}(T_k) = 0$ over a cycle of length T_k gives, for $t > t_k^d$,

$$I_k^{off}(t) = e^{-\theta_k t} \int_t^{T_k} D_k(s) e^{\theta_k s} ds, \quad (10)$$

and the cycle holding cost

$$\mathcal{H}_k = c_k^h \left[\int_0^{t_k^d} I_k^{off}(t) dt + \int_{t_k^d}^{T_k} e^{-\theta_k t} \int_t^{T_k} D_k(s) e^{\theta_k s} ds dt \right] \quad (11)$$

The off-shore echelon is replenished from the onshore hub with a lead time that is itself a random variable determined by the metocean process,

$$L_k(\tau, \kappa) = L_k^{man} + L_k^{tr} + \underbrace{\mathbb{E}[W_{j(k)m}(\kappa)]}_{\text{weather waiting}}, \quad (12)$$

where $j(k)$ is the minimum vessel class able to carry component k . This is the central endogeneity of the paper. L_k depends on τ (through the seasonality of triggered jobs) and on κ (through storm climatology). Under a base stock policy S_k and a compound Poisson demand stream with rate $D_k(\tau)$, the steady state pipeline is $\mu_k = D_k L_k / 365$ with variance $\sigma_k^2 = D_k L_k / 365 + (D_k / 365)^2 \sigma_{L_k}^2$. Using the normal approximation and the standard loss function $\mathcal{L}(z) = \varphi(z) - z(1 - \Phi(z))$, expected backorders and on-hand stock are

$$\mathbb{E}[B_k] = \sigma_k \mathcal{L}(S_k - \mu_k / \sigma_k), \quad \mathbb{E}[I_k] = S_k - \mu_k + \mathbb{E}[B_k] \quad (13)$$

Each backordered unit triggers an emergency expedite whose expected outage duration is Δ^{ex} , so the associated production loss is $\Delta^{ex} P^r CF c^{ens}$ per backorder year, with CF the capacity factor. The reduced form annual cost used to construct the response surface of Section 6.3 is therefore

$$\mathcal{C}(S, \tau) = \underbrace{c^h c^s \mathbb{E}[I]}_{\text{holding}} + \underbrace{\mathbb{E}[B] \Delta^{ex} P^r CF c^{ens}}_{\text{unavailability}} + \underbrace{c^{pm} \tau}_{\text{residual life discarded}} + \underbrace{c^v D(\tau) (\zeta_0 + \zeta_1 e^{-\tau/\tau_v})}_{\text{vessel}}. \quad (14)$$

4.5. Vessel scheduling constraints

For each scenario, vessel dispatch obeys standard routing constraints augmented by weather feasibility:

$$\sum_{i \in \mathcal{N}} \sum_{k \in \mathcal{K}} y_{ikt\omega} \delta_k \leq \sum_{j \in \mathcal{V}} (v_j + v_{jt\omega}^{sp}) a_{jt\omega} \bar{c}_j, \quad \forall t, \omega \quad (15)$$

$$a_{jt\omega} = \prod_{s=t}^{t+\Delta_j-1} \mathbb{1}\{H_{s\omega}^S \leq H^J\}, \quad \forall j, t, \omega \quad (16)$$

$$y_{ikt\omega} \leq \sum_{s \leq t} r_{iks\omega}, \quad \sum_t y_{ikt\omega} \leq 1, \quad \forall i, k, \omega \quad (17)$$

where $r_{iks\omega}$ indicates that the CBM trigger Equation (7) has fired, δ_k is the vessel day requirement, $\bar{c}_j$ the daily job capacity, and v^{sp} spot charters. Constraint Equations (16) encodes the persistence requirement exactly, at the cost of a product of binaries, it is linearised by the standard big- M chain $a_{jt\omega} \leq \mathbb{1}\{H_{s\omega}^S \leq H^J\}$ for each s in the window plus a lower bounding aggregate cut.

Inventory linkage closes the loop:

$$I_{kt\omega}^{off} = I_{k,t-1,\omega}^{off} (1 - \theta_k \mathbb{1}\{t > t_k^d\}) + z_{kt\omega}^{in} - \sum_i y_{ikt\omega} + B_{kt\omega} - B_{k,t-1,\omega} \quad (18)$$

with z^{in} the arriving replenishment, itself gated by accessibility.

4.6. Storage, hydrogen and export

Wind generation $g_{t\omega} = \sum_i P^r \varrho(V_{t\omega}) \chi_{it\omega}$, where ϱ is the power curve and $\chi_{it\omega} \in \{0,1\}$ the availability of turbine i which is determined by the repair state and thus by the inventory and vessel decisions. The energy balance at the OSP is

$$g_{t\omega} + \eta^{dis} b_{t\omega}^{dis} = q_{t\omega} + b_{t\omega}^{ch} + h_{t\omega} + \varphi_{t\omega}, \quad \forall t, \omega \quad (19)$$

with φ curtailment. Storage dynamics and limits:

$$E_{t\omega} = E_{t-1,\omega} + \eta^{ch} b_{t\omega}^{ch} - b_{t\omega}^{dis} / \eta^{dis}, \quad 0 \leq E_{t\omega} \leq E^b \quad (20)$$

$$b_{t\omega}^{ch} + b_{t\omega}^{dis} \leq P^b, \quad 0 \leq h_{t\omega} \leq P^{el}, \quad 0 \leq q_{t\omega} \leq \bar{P}^{hv} \varsigma_{t\omega} \quad (21)$$

where $\varsigma_{t\omega} \in [0,1]$ is the derating of the HVDC link, set to zero during converter or cable contingencies. Hydrogen output is $\mathcal{H}_\omega^2 = \sum_t h_{t\omega} / \eta^{el}$. Energy not served relative to a firm contractual schedule $\bar{q}_t$ is

$$u_{t\omega} = \max\{0, \bar{q}_t - q_{t\omega}\} \quad (22)$$

The critical structural point is that Equations (19)–(22) let the system monetise an outage in two distinct ways deferral via the battery and conversion via the electrolyser whose relative value depends on the duration of the disruption. Batteries dominate for interruptions of hours while electrolysers dominate when the export corridor rather than the generation is impaired, and when the disruption persists for days, as in a cyclone.

4.7. Resilience metric

Following the performance loss convention, define the normalised delivered power trajectory $Q_\omega(t) = q_{t\omega} / \bar{q}_t$ and the resilience index over an event window $[t_0, t_1]$,

$$\mathcal{R}_\omega = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} Q_\omega(t) dt, \quad \mathcal{R} = \sum_\omega p_\omega \mathcal{R}_\omega \quad (23)$$

$\mathcal{R}$ integrates the depth of the nadir, the speed of recovery and the residual capability during the event, and

is bounded in $[0,1]$. Section 6.4 maps $\mathcal{R}$ over the storage design space and shows the underlying trajectories.

4.8. Risk measure and distributional ambiguity

Let Φ_ω denote the second stage cost, including ENS penalties net of hydrogen revenue. We use a mean CVaR objective,

$$CVaR_\alpha = \zeta + \frac{1}{1-\alpha} \sum_{\omega} p_{\omega} \psi_{\omega}, \quad \psi_{\omega} \geq \Phi_{\omega} - \zeta, \psi_{\omega} \geq 0 \quad (24)$$

and protect against misspecification of $\{p_{\omega}\}$ with a type-1 Wasserstein ball $\mathbb{B}_\varepsilon(\hat{\mathbb{P}}_N)$. The distributionally robust counterpart of the risk term admits the finite convex reformulation

$$\sup_{\mathbb{P} \in \mathbb{B}_\varepsilon(\hat{\mathbb{P}}_N)} \mathbb{E}_{\mathbb{P}}[\Phi(\xi)] = \inf_{\varrho \geq 0} \left\{ \varrho \varepsilon + \frac{1}{N} \sum_{n=1}^N \sup_{\xi} [\Phi(\xi) - \varrho \|\xi - \hat{\xi}_n\|] \right\} \quad (25)$$

which, because Φ is piecewise linear in ξ for fixed integer decisions, reduces to a finite set of linear constraints on ϱ and the dual multipliers of the second stage.

4.9. Complete formulation

$$\begin{aligned} \min_{S, \tau, v, E^b, P^b, P^{el}} \quad & \sum_k \left(c_k^h c_k^s \mathbb{E}[I_k] + c_k^s \text{crf} \sum_e S_k^e \right) + \sum_j c_j^v v_j + \text{crf} (c_E^b E^b + c_P^b P^b + c^{el} P^{el}) \\ & + (1 - \varrho_r) \mathbb{E}_{\mathbb{P}}[\Phi(\xi)] + \varrho_r \sup_{\mathbb{P} \in \mathbb{B}_\varepsilon} CVaR_\alpha[\Phi(\xi)] \\ \text{s.t.} \quad & (13), (15)-(22), \mathcal{R} \geq \underline{\mathcal{R}}, \\ & \sum_k \sum_e c_k^s S_k^e \leq \bar{B}, \quad E^b \leq Y P^b, \quad P^{el} \leq \bar{P}^{hv}, \\ & S \in \mathbb{Z}_+^{|\mathcal{K}| \times |\mathcal{E}|}, \tau \in [\tau_{min}, \tau_{max}], v \in \mathbb{Z}_+^3, E^b, P^b, P^{el} \geq 0, \end{aligned} \quad (26)$$

with crf the capital recovery factor, $\varrho_r \in [0,1]$ the risk aversion weight, $\bar{B}$ the spares budget, Y the maximum energy to power ratio and $\underline{\mathcal{R}}$ a regulatory resilience floor.

4.10. Structural results

Theorem 1 (Convexity in the base stock level): For fixed τ , κ and ϕ , and for any lead time distribution with a log-concave density, the reduced form cost $\mathcal{C}(S, \tau)$ in Equation (14) is convex in S on $\mathbb{R}_+$, and the optimal integer base stock level satisfies the newsvendor type condition

$$Pr(\mathcal{D}_L \leq S^* - 1) < \frac{c^u}{c^u + c^h c^s} \leq Pr(\mathcal{D}_L \leq S^*), \quad c^u = \Delta^{ex} P^r C F c^{ens} \quad (27)$$

Proof of Theorem 1. $\mathbb{E}[I]$ is affine in S plus $\mathbb{E}[B]$, $\mathbb{E}[B] = \mathbb{E}[(\mathcal{D}_L - S)^+]$ is the expectation of a convex function of S and therefore convex, and its first difference is $-\text{Pr}(\mathcal{D}_L > S)$. Summing a convex and an affine term preserves convexity, log-concavity of the lead time demand density guarantees unimodality of the first difference so the discrete condition is necessary and sufficient. $\square$

This convexity structure follows the classical multi-echelon base stock theory of^[32] and the gamma process deterioration framework of^[30], it is also the off-shore analogue of the pseudo-convexity established for non-instantaneously deteriorating item systems with linear demand^[7]. The deterioration term θ_k enters only through μ_k and σ_k and therefore does not disturb convexity, though it shifts S^* upward.

Theorem 2 (Monotone CBM trigger in storm severity). If λ^{eff} is convex and decreasing in τ as in Equation (8), and the expected weather waiting time $\mathbb{E}[W]$ in Equation (3) is increasing in κ and decreasing in τ with

$\partial^2 \mathbb{E}[W] / \partial \tau \partial \kappa \leq 0$, then the optimizer $\tau^*(\kappa)$ is non-decreasing in κ .

Proof of Theorem 2. The cross partial of the objective satisfies $\partial^2 \mathcal{C} / \partial \tau \partial \kappa \leq 0$, so $\mathcal{C}$ is sub-modular in (τ, κ) on a lattice, Topkis' theorem^[33] then yields monotone comparative statics. $\square$

Managerially, harsher weather warrants earlier intervention, not merely more stock. The intuition is that in a storm dominated climate the option value of an early, weather flexible repair slot exceeds the cost of discarding residual life.

Theorem 3 (Sub-modularity of resilience in the storage vector). Let $\mathcal{R}(\mathbf{c})$ be the resilience index as a function of a vector of discrete storage increments $\mathbf{c}$. Under Equations (19)–(23) with non-increasing marginal displacement of unserved energy, $\mathcal{R}$ is monotone and sub-modular, and greedy selection attains at least $(1 - 1/e)$ of the optimum.

Proof of Theorem 3. Each increment reduces the ENS integral by the energy it can displace, which is bounded by the residual deficit. The residual deficit is non-increasing in the set of already installed increments, giving diminishing returns. Monotonicity is immediate as increments never increase ENS. The Nemhauser Wolsey Fisher bound applies^[34].

Corollary 1. If the storm severity multiplier κ exceeds the threshold $\kappa^\dagger$ at which the expected uninterrupted export deficit exceeds E^b , the marginal resilience value of electrolyser capacity strictly exceeds that of battery capacity. This is the analytical statement behind the crossover observed at $\kappa \approx 1.3$ in Section 6.4. $\square$

5. Solution methodology

Below is the methodology solution is shown in which Methodological work flow from hindcast data to design outputs are shown (see **Figure 2**).




Figure 2. Methodological work flow from hindcast data to design outputs.

5.1. Scenario generation

Metocean and wind scenarios are generated by a two-layer procedure. The regime chain Z_t is simulated from month specific transition matrices estimated by maximum likelihood from 40 years of ERA5 class hindcast^[35], conditional wave heights are drawn from the fitted Weibull families and wind speeds from a Gaussian copula linking H^s and V with an empirical rank correlation of 0.71. Extreme events are super-sampled using importance sampling with a likelihood ratio correction so that the 1in 50-year storm is represented with adequate frequency without biasing the mean. Component degradation increments are drawn from Equation (4) conditional on the simulated loading. Prices are drawn from a vector autoregressive model with regime dependent innovations. Scenario reduction by fast forward selection on a Wasserstein metric compresses 20,000 raw paths to $|\Omega| \in [50, 3000]$.

Formally, let $\{\hat{\xi}_n\}$ for $n = 1, \dots, 20,000$ denote the raw simulated paths under the empirical distribution

$\hat{\mathbb{P}}$ with mass $1/20,000$ each. Wasserstein scenario reduction selects a retained index set J and a reweighting so as to minimise the type 1 Wasserstein distance, defined as $D_W(\hat{\mathbb{P}}, \mathbb{P}_J) = \min_{\pi \in \Pi(\hat{\mathbb{P}}, \mathbb{P}_J)} \sum_i \sum_{j \in J} \pi(i, j) \|\hat{\xi}_i - \hat{\xi}_j\|$, where the minimization is taken over all couplings π with the correct marginal distributions. Because optimal subset selection is combinatorial, we use the fast forward selection heuristic of^[36]. Starting from $J = \emptyset$, at each step we add $j^* = \arg \min_{j \notin J} \sum_{i \in J} \min_{l \in J \cup \{j\}} \|\hat{\xi}_i - \hat{\xi}_l\|$, which represents the scenario whose inclusion most reduces the residual nearest neighbour distance of the excluded paths, and redistribute the probability mass of each deleted scenario to its Wasserstein nearest survivor. The procedure is run to $|\Omega| = 3,000$ and truncated at checkpoints (1,500,500,250,50) to produce the scenario count sweep of Section 6.6. The sample size $|\Omega| = 1,500$ is retained as the operating point because Section 6.6 shows the value of the stochastic solution has plateaued by that count.

5.2. Decomposition

Problem (26) is a two-stage stochastic mixed integer program with integer recourse in the routing variables. We apply an integer L-shaped method:

- a) **Master problem:** Relax the second stage to a set of optimality cuts $\vartheta \geq \pi_\ell^\top (h_\ell - T_\ell x)$ and solve for $(S, \tau, v, E^b, P^b, P^{el})$.
- b) **Subproblems:** For each ω , fix first stage variables and solve the hourly dispatch routing LP/MIP. Where the routing binaries are integral in the optimal LP relaxation which held for 94% of subproblems in our experiments owing to the interval structure of (16) the standard Benders cut is valid, otherwise, a Laporte Louveaux integer cut is added.
- c) **Pareto-optimal cuts:** Among the alternative dual solutions, we select the one maximising the cut value at a core point x^0 , which reduced iteration counts by roughly 40%.
- d) **Valid inequalities:** Two families are added: a base stock lower bound from Proposition 1 evaluated at the seasonal maximum lead time, and a storage lower bound from the greedy sub-modular solution of Proposition 3.
- e) **Warm start.** The greedy solution of Proposition 3 combined with the closed form S^* of Proposition 1 provides an incumbent within 6.4% of the final solution on average, materially shortening the branch and bound tree.

Algorithm 1: Pareto-optimal cut generation.

Input: incumbent $\hat{x}$, core point x^0 , scenario ω

- 1 Solve SP $(\hat{x}, \omega)$; obtain $\Phi(\omega; \hat{x})$ and check integrality of the LP relaxation
- 2 if LP relaxation integral then
- 3 Re-solve SP (x^0, ω) restricted to the dual-optimal face of SP $(\hat{x}, \omega)$; obtain π_{ω^0}
- 4 Cut $\leftarrow \vartheta \geq \pi_{\omega^0}^\top (h_{\omega} - T_{\omega} x)$
- 5 else
- 6 Cut $\leftarrow$ Laporte–Louveaux integer L-shaped cut at $\hat{x}$
- 7 return Cut

On the Pareto-optimal cut specifically, this construction is not proposed for the first time in this paper: it adapts the core point method of^[37] to the present two stage stochastic integer setting. Among the generally nonunique optimal duals of SP $(\hat{x}, \omega)$, the Pareto optimal cut uses the duals obtained by resolving the subproblem SP at an interior core point $x^0 \in \text{ri}(\mathcal{X})$, representing the analytic center of the box constraints on x computed once at initialisation, on the dual optimal face of SP $(\hat{x}, \omega)$, given by $\pi_{\omega^0} \in \arg \max \{\pi_{\omega}^\top (h_{\omega} -$

$T_\omega x^0 \mid \pi_\omega$ is dual feasible for $\text{SP}(\hat{x}, \omega)\}$.

Proposition (Pareto-optimality of the core point cut, adapted from^[37]). If x^0 is interior to $\mathcal{X}$ and π_ω^0 solves the above program, the cut $\vartheta \geq (\pi_\omega^0)^\top (h_\omega - T_\omega x)$ is not dominated by any other valid optimality cut on $\mathcal{X}$.

Proof sketch. Linear programming duality gives $\pi_\omega^\top (h_\omega - T_\omega x) \leq \Phi(\omega; x)$ for every dual feasible π_ω and every x , with equality holding where π_ω is optimal for $\text{SP}(x, \omega)$. The cut generated at x^0 is therefore tight at x^0 , and because x^0 is interior, any cut that is both valid on $\mathcal{X}$ and at least as tight at x^0 must coincide with the one obtained by maximising the cut value at x^0 , which is exactly how π_ω^0 is chosen, thereby ruling out a strictly dominating alternative. This completes the proof sketch; see^[37] for the general result. $\square$

The pseudocode for cut generation is stated as Algorithm 2. The master problem and subproblem are stated explicitly for completeness. Write $\mathbf{x} = (\mathbf{S}, \boldsymbol{\tau}, \mathbf{v}, \mathbf{E}^b, \mathbf{P}^b, \mathbf{P}^{el})$ for the first stage vector, $\mathcal{X}$ for the first stage feasible region defined by the budget, energy to power, and integrality constraints in the last two lines of Equation (26), and $\mathbf{c}(\mathbf{x})$ for the deterministic first stage cost given by the first line of Equation (26). The master problem at iteration ℓ is $\text{MP}(\ell)$, formulated as $\min_{\mathbf{x} \in \mathcal{X}} \{\mathbf{c}(\mathbf{x}) + \boldsymbol{\vartheta}\}$ subject to $\boldsymbol{\vartheta} \geq \boldsymbol{\pi}_l^\top (\mathbf{h}_l - \mathbf{T}_l \mathbf{x})$ for $l = 1, \dots, L$, where each optimality cut is generated from a previously solved subproblem. For scenario ω , fixing $\mathbf{x}$ at the current incumbent $\hat{\mathbf{x}}$, the subproblem $\text{SP}(\hat{\mathbf{x}}, \omega)$ is the hourly dispatch routing program $\min \Phi(\omega; \hat{\mathbf{x}}) = \sum_t \left[c_{t\omega}^e q_{t\omega} - \frac{c^{h2} H_\omega^2}{T} + c^{ens} u_{t\omega} \right]$ subject to Equation (15) through Equation (22) evaluated at $\hat{\mathbf{x}}$, with dual variables π_ω attached to the linking constraints Equations (15) and (18).

5.3. Algorithm

Algorithm 2. Accelerated integer L-shaped algorithm with distributional robustness.

Input: scenario set Ω , radius ε , confidence α , tolerance δ

- 1 Generate scenarios; reduce to Ω by Wasserstein fast-forward selection
- 2 Compute greedy storage vector $\mathbf{c}^0$ (Prop. 3) and $\mathbf{S}^0$ (Prop. 1); set incumbent $\mathbf{x}^0$
- 3 Initialise master with valid inequalities VI-1, VI-2; $\text{UB} \leftarrow C(\mathbf{x}^0)$, $\text{LB} \leftarrow -\infty$
- 4 repeats
- 5 Solve master $\rightarrow (\mathbf{x}, \vartheta)$; $\text{LB} \leftarrow \text{objective}(\text{master})$
- 6 for each $\omega \in \Omega$ (parallel) do
- 7 Solve dispatch–routing subproblem $\text{SP}(\mathbf{x}, \omega) \rightarrow \Phi_\omega$, duals π_ω
- 8 if LP relaxation fractional then generate integer L shaped cut
- 9 Solve inner sup over Wasserstein ball via (25) $\rightarrow$ robust risk term
- 10 Aggregate Pareto-optimal optimality cut at core point $\mathbf{x}^0$; add to master
- 11 $\text{UB} \leftarrow \min(\text{UB}, \text{first-stage cost} + \text{robust risk term})$
- 12 until $(\text{UB} - \text{LB})/\text{UB} \leq \delta$, $\mathbf{x}^* = (\mathbf{S}^*, \boldsymbol{\tau}^*, \mathbf{v}^*, \mathbf{E}^{b*}, \mathbf{P}^{b*}, \mathbf{P}^{el*})$, R, CVaR_α

To make the stopping rule precise, LB is the optimal objective value of the current relaxed master problem $\text{MP}(\ell)$, referenced in line 5 of Algorithm 2, which serves as a valid lower bound on Equation (26) because the optimality cuts underapproximate the true recourse function $\boldsymbol{\vartheta}$ from below by linear programming and Bender’s duality.

UB is the best feasible objective found so far, given by $\mathbf{c}(\mathbf{x})$ plus the exact rather than relaxed robust risk term, computed by resolving all subproblems at the incumbent $\mathbf{x}$, referenced in line 11, which serves as a valid upper bound because $\mathbf{x}$ is feasible for Equation (26).

LB is nondecreasing and UB is nonincreasing over iterations, and both converge to the optimal value of Equation (26) as $L \rightarrow \infty$ by standard L shaped convergence established by^[38].

The algorithm therefore terminates finitely at any tolerance $\delta > 0$, since the first stage integer variables take finitely many values on the bounded set $\mathcal{X}$.

5.4. Computational performance

Instances were solved with a commercial MIP solver on a 32-core workstation. For $|\Omega| = 1500$, $N = 100$, $T = 8760$ with 3-hour aggregation in the master, the algorithm reached a 0.8% gap in 47 minutes, the monolithic formulation failed to close below 14% in six hours. Parallel subproblem evaluation scaled near linearly to 24 threads.

6. Case study and numerical results

6.1. Data and parameterization

The reference case is a 1.2 GW farm of $N = 100$ turbines rated at $P^r = 12$ MW, located 78 km from the O&M base, exporting through a 1.2 GW symmetrical-monopole VSC-HVDC link. Metocean parameters are fitted to North Sea hindcast statistics, the Weibull scale in month m follows $\lambda(m) = 1.55 + 0.95\cos(2\pi(m-1)/12)$ m with shape $k(m) = 2.05 + 0.12\cos(2\pi(m-1)/12)$, and the persistence exponent is $\rho = 1.65$. Wave height limits are $H^{ctv} = 1.5$ m, $H^{sov} = 2.5$ m and $H^{hlv} = 1.8$ m for lifting operations. Baseline major component failure intensity is $\lambda_0 = 0.42$ per turbine-year across the six critical types, the irreducible sudden failure fraction is $\underline{a} = 0.38$ and the prognostic decay constant $\tau_c = 28$ days.

Economic parameters are these followings, spare unit cost $c^s = \text{€}420,000$ (weighted across component types), holding rate $c^h = 11\%$ per annum inclusive of insurance, preservation and platform space, HLV charter $c^v = \text{€}42,000$ per day, and a combined lost revenue plus imbalance penalty of $c^{ens} = \text{€}145/\text{MWh}$. The emergency expedite outage is $\Delta^{ex} = 60$ days and the capacity factor $\text{CF} = 0.46$. Storage capital costs are $\text{€}195/\text{kWh}$ for the battery, $\text{€}118/\text{kW}$ for its power stage and $\text{€}1,290/\text{kW}$ for the platform marinised PEM electrolyser, all annualised at a 7% discount rate over 20 years while the hydrogen offtake price is $\text{€}4.60/\text{kg}$. Wasserstein radius ε is selected by cross validation at $\varepsilon = 0.06$, and $\alpha = 0.95$.

For the Indian variant, the seasonal structure is replaced by the Bay of Bengal bimodal cyclone climatology (pre-monsoon April-June and post-monsoon October-December peaks), the storm severity multiplier is calibrated to $\kappa = 1.62$ against IMD landfall records for the Odisha coast^[39,40], and vessel day rates are reduced by 22% to reflect regional charter markets while spare-part transport lead times increase by 19 days owing to import dependence.

6.2. Accessibility and the weather constraint

Figure 3 shows the accessibility surface. Three design relevant features emerge. First, the surface is strongly convex in the wave height limit in the range 1.5-2.2 m, so a modest capability upgrade from a conventional CTV to a motion compensated SOV gangway buys a disproportionate gain in accessible hours, at $H^c = 1.5$ m the December accessibility is 0.21, rising to 0.47 at 2.2 m. Second, the seasonal amplitude is severe. July accessibility at the HLV limit is 3.1 times the January value. Third, the surface is markedly flatter in summer, meaning that summer scheduling is robust to wave limit assumptions while winter scheduling is not an argument for concentrating major component campaigns in the shoulder months even at the cost of discarded residual life. This is the mechanism underlying Proposition 2.




Figure 3. Weather window accessibility surface π_m as a function of the vessel wave height limit and month of year, computed from Equation (2) with $\kappa = 1$.

6.3. Joint inventory maintenance optimisation




Figure 4. Expected annual major component logistics cost as a joint function of the off-shore base stock level S and the condition-based maintenance trigger τ , from Equation (14).

Figure 4 displays the cost surface of Equation (14). The surface is convex in S for every τ , confirming Proposition 1 numerically, and the minimum lies in the interior at $S^* = 3$ off-shore units with a trigger of $\tau^* = 50$ days of residual life, yielding €1.79 M per annum. The geometry is informative. Along the S axis the cost rises steeply to the left of the optimum under stocking is punished by the €1.15 M expected cost of each backorder year through the expedite channel but only gently to the right, where the penalty is the €46.2 k annual carrying cost per unit. This asymmetry, a direct consequence of the newsvendor ratio $c^u / (c^u + c^h c^s) = 0.96$ in Proposition 1, argues for deliberate over provisioning relative to a symmetric loss intuition.

Along the τ axis the trade-off is between discarded residual life and logistics flexibility. Very short triggers ($\tau < 15$ days) leave insufficient time to secure an HLV window in the winter regime. The resulting mix of emergency spot charters and extended outages drives cost to €3.4 M. Very long triggers ($\tau > 65$ days)

waste component life and inflate the preventive term. Note that $\tau^* = 50$ days is substantially longer than the 20–30 day triggers commonly reported in reliability centred studies that ignore vessel access, the difference is precisely the option value of a weather window.

Table 2. Optimal design and performance by policy.

Policy	S^*	τ^* (d)	Holding (M€)	Vessel (M€)	Downtime (M€)	Preventive (M€)	Total (M€)	Availability	$\mathcal{R}$
Corrective only	1	–	0.14	2.90	1.86	0.00	4.90	0.882	0.41
Calendar preventive	2	–	0.31	2.10	0.62	0.35	3.38	0.928	0.56
Condition-based, decoupled	2	32	0.22	1.35	0.34	0.50	2.41	0.951	0.68
Digital twin ($\phi = 1$) + joint (proposed)	3	50	0.19	0.92	0.13	0.55	1.79	0.968	0.87




Figure 5. Decomposition of annual O&M cost by policy.

Table 2 and **Figure 5** show that the 63.5% cost reduction relative to the corrective benchmark is achieved by shifting expenditure rather than merely cutting it: holding cost rises by 36% and preventive cost is created from nothing, while vessel charter falls by 68% and downtime cost by 93%. The comparison between rows 3 and 4 isolates the value of joint versus sequential optimisation at €0.62 M per annum, or 25.7% a figure that is invisible to any model that optimizes inventory and maintenance separately.

On the attribution of this gain: the €0.62 M, 25.7% gap between rows 3 and 4 of Table 2 isolates joint versus sequential optimization of the base stock level and the CBM trigger at fixed sensor coverage, and so does not by itself identify the marginal contribution of the digital twin. That contribution is isolated separately in Section 6.6, where sensor coverage ϕ is swept from 0 to 1 inside the joint (row 4) formulation and the value of the stochastic solution rises from 4.0% at $\phi = 0$ to 11.2% at full coverage. The row label “Digital twin + joint” should accordingly be read as the performance of the full model at $\phi = 1$, with its improvement over row 3 (decoupled, decision-level jointness only, no explicit twin credit) reflecting two distinct and multiplicative rather than additive effects: jointness, quantified here at €0.62 M, and instrumentation coverage, quantified in Section 6.6 as the 4.0% to 11.2% rise in VSS. Because the two interact multiplicatively, their separate figures do not sum to the row 3 to row 4 gap, each isolate one axis of the joint benefit while holding the other fixed, and Section 6.6 is the source of the twin’s marginal value on its own.

A clarification on the internal consistency of **Table 2** is needed. At 1.2 GW and $CF = 0.46$, annual production at full availability is approximately 4,836 GWh, so a one percentage point change in a production-based availability measure is worth approximately €7.0 M at $c^{\text{ens}} = \text{€}145/\text{MWh}$. Applying this conversion to the 8.6 point rise in the Availability column, from 0.882 to 0.968, suggests approximately €60 M of

recovered value, far exceeding the €1.73 M fall in the Downtime column. The two columns are not measuring the same quantity, and the apparent factor of 35 discrepancy is definitional rather than an arithmetic error.

The Availability column reports turbine time availability, representing the fraction of turbine hours the machine is mechanically able to generate, including many short interruptions completed within a scheduled weather window from on hand off-shore stock that never trigger the emergency expedite channel and therefore never register as a production shortfall against the firm export schedule. The Downtime column, by contrast, is only the second or unavailability term of the reduced form cost Equation (14), given by $\mathbb{E}[B] \cdot \Delta^{\text{ex}} \cdot P^r \cdot \text{CF} \cdot c^{\text{ens}}$, which prices only backorder driven emergency expedite events, evaluated at the expedite outage $\Delta^{\text{ex}} = 60$ days used throughout Section 6. Its 93% fall across the four policies is a statement about the near elimination of backorder events specifically, not a statement about total lost production across every downtime hour, most of which is absorbed without triggering the emergency channel once base stock and condition-based maintenance triggers are set jointly. The system level energy not served accounting that would support the calculation on the scale of €60 M noted above is reported separately as a distributional quantity, namely the Conditional Value at Risk of energy not served, in Section 6.5. The two are deliberately not both charged in **Table 2**, since summing the Section 6.5 Conditional Value at Risk term into the reduced form logistics cost of **Table 2** would double count the risk premium already captured by the Conditional Value at Risk term in objective Equation (26).

6.4. Storage design and the resilience frontier




Figure 6. Resilience index $\mathcal{R}$ over the battery electrolyser design space.

Figure 6 maps the resilience index over the storage design space. The surface is monotone and concave in each coordinate, consistent with the sub-modularity of Proposition 3, and the contour projection reveals a pronounced ridge as for a given resilience target there is a locus of battery electrolyser combinations, and the cheapest point on that locus is generally not a corner. At $\mathcal{R} = 0.85$ the frontier runs from $(E^b, P^{\text{el}}) = (680, 90)$ to $(310, 380)$, the cost-minimising point on this segment is $(450, 300)$, the configuration adopted in the reference design.

The concavity is economically important. The first 250 MWh of battery capacity raises $\mathcal{R}$ from 0.37 to 0.66; the next 250 MWh adds only 0.11. Beyond roughly 700 MWh the marginal resilience per euro falls below that of the electrolyser branch for every storm severity multiplier we tested. The mild curvature reversal at high P^{el} reflects a real effect. An oversized electrolyser diverts energy that would otherwise be exported under the firm schedule, so hydrogen coupling is a complement to, not a substitute for, export capacity.




Figure 7. Performance trajectories $Q(t)$ during a 1 in 50-year storm for four storage configurations.

Figure 7 makes the mechanism visible at event resolution. Without storage, delivered performance collapses to 0.28 at the storm nadir and does not recover until repair vessels regain access some 40 hours after the wind subsides, giving $\mathcal{R} = 0.41$. A battery only configuration lifts the nadir but exhausts within 9 hours. Hydrogen only sustains a shallower but longer plateau because the electrolyser continues to absorb generation when the HVDC link is derated, converting energy that would otherwise be curtailed. The hybrid configuration dominates, nadir 0.86, recovery within 14 hours of the wind event, $\mathcal{R} = 0.87$. Corollary 1 is confirmed numerically the crossover at which electrolyser capacity becomes the higher value marginal investment occurs at $\kappa^\dagger = 1.31$, which is below the $\kappa = 1.62$ calibrated for the Odisha coast. In cyclone exposed geographies, hydrogen coupling is a resilience technology, not merely a decarbonisation one. To make the mechanism underlying **Figure 7** auditable against the energy balance Equation (19), **Table 3** instantiates that balance at four representative hours of the 1 in 50-year storm event for the hybrid configuration, specifically six hours before the nadir, the nadir hour itself, positive **14 h**, which is the recovery point reported in the text, and positive **48 h**, representing the post event state. The reported values are an illustrative instantiation consistent with the aggregate trajectory of **Figure 7** and satisfy the energy balance Equation (19) exactly at each row. The underlying hourly simulation record is available from the corresponding author per the Data Availability statement.

Table 3. Illustrative hourly energy balance Equation (19) during the 1 in 50-year storm, hybrid configuration.

Hour	$g_{t\omega}$	$\eta^{\text{dis}} b_{t\omega}^{\text{dis}}$	$b_{t\omega}^{\text{ch}}$	$h_{t\omega}$	$q_{t\omega}$	$\phi_{t\omega}$	$Q(t) = q_{t\omega}/\bar{q}$
t0–6 h (pre)	522	0	18	24	480	0	0.98 ($\bar{q} = 490$)
Nadir	40	128	0	0	168	0	0.86 ($\bar{q} = 195$)
+14 h (recov.)	310	0	0	15	295	0	0.98 ($\bar{q} = 300$)
+48 h (post)	515	0	30	20	465	0	0.97 ($\bar{q} = 480$)

Note: $g_{t\omega}$ is generation, $\eta^{\text{dis}} b_{t\omega}^{\text{dis}}$ is discharge output, $b_{t\omega}^{\text{ch}}$ is battery charge (corrected from an erroneous repeated “dis” superscript in the submitted version), $h_{t\omega}$ is hydrogen conversion load, $q_{t\omega}$ is firm export, and $\phi_{t\omega}$ is curtailment/spill, all in MWh; the last column reports the ratio of realised to reference export $q_{t\omega}/\bar{q}$, with the reference level $\bar{q}$ given per row.

The energy balance check per row satisfies $g_{t\omega} + \eta^{\text{dis}} b_{t\omega}^{\text{dis}} = q_{t\omega} + b_{t\omega}^{\text{ch}} + h_{t\omega} + \phi_{t\omega}$ with all quantities expressed in MWh.

6.5. Tail risk and robustness




Figure 8. 95% conditional value at risk of energy not served as a function of the uncertainty budget Γ and the storm severity multiplier κ .

Figure 8 characterises downside exposure. At the nominal severity $\kappa = 1.0$ and $\Gamma = 0$, $\text{CVaR}_{0.95}$ of energy not served is 58.6 GWh per annum, raising the uncertainty budget to $\Gamma = 6$ reduces it to 32.5 GWh, a 44.6% improvement, at a first stage cost increase of 8.9%. The marginal return to robustness declines sharply beyond $\Gamma \approx 7$, where the additional protection is purchased mainly through vessel over chartering with little further reduction in tail exposure the classic price of robustness saturation.

The severity axis is more alarming. Holding the design fixed at its $\kappa = 1$ optimum and stressing the climatology to $\kappa = 1.8$ raises CVaR by a factor of 2.9, whereas re-optimising the design at $\kappa = 1.8$ recovers 61% of that deterioration. The implication for asset owners and insurers is that resilience is substantially repurchasable through logistics and storage decisions, but only if the design is revisited when climatology is re-estimated. Static designs age badly under non-stationary storm regimes.

The distributionally robust layer contributes an out of sample improvement of its own. Testing designs obtained under the sample average approximation against a held-out set of 5000 unseen scenarios, the SAA design overspends its predicted cost by 7.8% on average, whereas the Wasserstein design with $\varepsilon = 0.06$ overspends by 1.9%, at an in-sample cost penalty of 2.4%.

6.6. Tri-objective trade-offs and the value of information

A clarification on the construction of **Figure 9** is needed, since Equation (26) is, as its statement makes explicit, a single objective program in which resilience enters only through the constraint $\mathcal{R} \geq \underline{\mathcal{R}}$ and storage enters the cost objective directly, meaning it is not solved as a vector optimization problem.

Figure 9 is instead generated by an outer ε constraint sweep wherein model Equation (26) is resolved 420 times, each time fixing the resilience floor $\underline{\mathcal{R}}$ and, orthogonally, a hydrogen offtake requirement at distinct grid values, and recording the resulting minimised levelised cost. The 420 triples comprising levelised cost of energy, resilience, and hydrogen yield are then filtered for nondominance in the usual sense.

This is the standard ε constraint device for tracing an approximate Pareto-frontier from a family of single objective programs indexed by right hand side parameters, and the resulting object is legitimately called a Pareto set, representing a set of realised nondominated outcomes rather than an artefact of solving Equation (26) as a three objective problem.




Figure 9. ε constraint trade-off (Pareto) set in the (LCOE, resilience, hydrogen yield) space, obtained by re-solving the single objective program (26) at 420 grid points of the resilience floor and hydrogen offtake requirement, coloured points are non-dominated.

Figure 9 presents the tri-objective trade-off among levelised cost, resilience and hydrogen yield across 420 sampled designs, of which 200 are non-dominated. Two structural insights follow. First, the Pareto set is highly non-convex in the resilience direction. Between $\mathcal{R} = 0.55$ and $\mathcal{R} = 0.80$ the LCOE penalty is modest at roughly €0.9/MWh per 0.1 of resilience, but beyond $\mathcal{R} = 0.90$ it steepens to €4.7/MWh per 0.1. A regulatory resilience floor set at 0.85-0.88 therefore extracts most of the available protection at modest cost, while a floor at 0.95 would be disproportionately expensive. Second, hydrogen yield and resilience are complementary over most of the frontier designs that produce more hydrogen are typically also more resilient because the electrolyser absorbs energy during export constraints but they compete at the extreme high-yield end, where oversized electrolysers cannibalise firm export.




Figure 10. Value of the stochastic solution as a function of scenario count and digital twin sensor coverage ϕ .

Figure 10 quantifies the value of stochastic modelling and of instrumentation jointly. The value of the stochastic solution rises steeply with scenario count up to roughly 1,200 scenarios and then plateaus, justifying our operating point of $|\Omega| = 1500$. More interesting is the sensor coverage axis as at $\phi = 0$, no condition monitoring, age based hazards only VSS is 4.0%, whereas at full coverage it reaches 11.2%. The interaction is multiplicative rather than additive, prognostic information is most valuable precisely when the decision model is capable of exploiting it through joint stock and vessel positioning. An operator who installs condition monitoring but retains a deterministic, decoupled planning process captures roughly a third of the available

benefit.

6.7. Sensitivity analysis

A one at a time sensitivity over $\pm 30\%$ ranges rank the drivers of total cost as: expedite outage duration Δ^{ex} (elasticity 0.61), HLV wave limit H^{hlv} (-0.48), storm severity multiplier κ (0.44), spare unit cost (0.19), electricity price (0.17) and holding rate (0.06). The dominance of Δ^{ex} and H^{hlv} over purely financial parameters reinforce the paper’s central claim. In off-shore O&M, physical accessibility is the binding economic constraint. A global Sobol analysis on the same range’s attributes 71% of output variance to first order effects and 29% to interactions, of which the $\kappa \times H^{hlv}$ pair alone accounts for 12% confirming that severity and vessel capability must be considered jointly.

6.8. Indian coastal variant

Under the Bay of Bengal calibration, the optimal design shifts materially. S^* rises from 3 to 5 off-shore units (driven by the 19 day import lead time extension), τ^* lengthens from 50 to 63 days (Proposition 2 in action, as κ rises from 1.00 to 1.62), and the optimal electrolyser rating grows from 300 MW to 415 MW while battery capacity falls from 450 MWh to 335 MWh. Total annualised cost per MW installed is 14% higher than the North Sea reference despite lower charter rates, and $\mathcal{R}$ at the optimum is 0.81 rather than 0.87 the residual gap reflecting the irreducible severity of cyclone season inaccessibility. The policy reading is that Indian off-shore wind zones will require either deeper strategic spare part reserves held domestically or a materially larger hydrogen offtake infrastructure to reach European resilience norms.

6.9. Comparison against alternative solution methods

To assess performance against alternative solution methods, beyond sensitivity analysis alone, **Table 4** compares the proposed accelerated integer L-shaped method described in Section 5 against three alternatives on the reference 1.2 GW instance, namely a plain integer L-shaped method of Laporte and Louveaux^[38] without Pareto optimal cuts or the Proposition 1 and Proposition 3 warm start, the monolithic mixed integer programming formulation already reported in Section 5.4, and a population based metaheuristic of the Non-dominated Sorting Genetic Algorithm II style of the kind used in the operations and maintenance routing literature reviewed in Section 2.1, given the same 47 minute computational budget as the proposed method.

Table 4. Solution method comparison on the reference 1.2 GW, 100 turbine instance ($|\Omega| = 1,500$).

Method	Gap at stop	Time to 0.8%gap	Notes
Proposed (Pareto cuts and warm start)	0.8%	47 min	reported in Section 5.4
Plain integer L shaped	0.8%	118 min	no Pareto cuts (approximately +40%iterations), no warm start
Monolithic MIP	14.0%	did not close in 6 h	reported in Section 5.4
NSGA II metaheuristic (47 minute budget)	4.1%*	not applicable	*best found gap versus proven optimum, with $\pm 2.3\%$ run to run variance over 10 seeds and no optimality certificate

The comparison shows that both the Pareto optimal cuts and the structural warm start are individually load bearing for tractability, where removing them roughly increases the solve time by a factor of 2.5 to reach the same gap. Furthermore, the results show that the monolithic formulation is not competitive at this scale, and that a metaheuristic baseline, while feasible within the same wall clock budget, neither certifies optimality nor matches the exact method’s solution quality. It also exhibits materially higher run to run variance, which represents a property with direct consequences for reproducibility of the resilience and cost figures reported in Section 6.

7. Managerial, regulatory and policy implications

For operators: Three actionable rules emerge. (i) Stock asymmetrically. The newsvendor ratio in off-shore major components exceeds 0.95, so the cost of a stockout dwarfs the cost of carrying, base stock levels should be set well above mean pipeline demand. (ii) Trigger earlier in harsher climates. Proposition 2 and Section 6.8 show that the optimal CBM threshold lengthens with storm severity, reliability centred thresholds transplanted from onshore practice are systematically too short. (iii) Buy accessibility, not just capacity. The elasticity of cost to the HLV wave limit (-0.48) exceeds that to spare cost (0.19) by a factor of 2.5, so investment in motion compensation, walk to work systems and higher limit lifting technology outperforms equivalent spending on inventory.

For system planners and regulators: The Pareto geometry of **Figure 9** suggests that a resilience floor in the range 0.85–0.88 is the efficient regulatory instrument, floors above 0.92 impose costs that rise faster than the protection delivered. Where off-shore hubs are being planned as multi-energy nodes, hydrogen coupling should be evaluated on a joint decarbonisation and resilience basis, valuing electrolyzers only on hydrogen revenue understates their system value by the CVaR reduction they deliver, which in our reference case is worth €0.41 M per annum, equivalent to roughly 9% of the electrolyser’s annualised capital cost.

For insurers and financiers: Because the design is re-optimisable against updated climatology, business interruption exposure should be underwritten against the re-optimized rather than the as-built configuration, with premium structures that reward periodic re-optimisation. The 61% recovery of storm induced deterioration through re-design (Section 6.5) is a quantifiable, contractible quantity.

For carbon accounting: Vessel activity is diesel intensive, and the 68% reduction in charter days delivered by joint optimisation carries an emissions dividend of approximately 3100 tCO_{2e} per annum for the reference farm. Where off-shore support vessels fall within emissions trading perimeters, this dividend is monetizable, and the carbon cap and trade and trade credit structures analysed in the deteriorating inventory literature^[7,41] provide a ready modelling template for internalising it in the objective.

8. Conclusions and future research

This paper has developed and solved an integrated model that co-optimizes spare parts inventory, digital twin triggered condition-based maintenance, vessel scheduling and hybrid battery electrolyser storage for off-shore wind farms exposed to extreme weather. The core methodological contribution is the endogenisation of the replenishment lead time in a metocean process that simultaneously governs failure exposure, repair accessibility and export capability, embedded in a distributionally robust two stage stochastic program with a CVaR objective on energy not served.

Three structural results convexity in the base stock level, monotonicity of the optimal maintenance trigger in storm severity, and sub-modularity of resilience in the storage vector provide both managerial intuition and computational leverage, the last supporting a $(1 - 1/e)$ greedy warm start within an accelerated integer L-shaped algorithm that closes realistic instances to 0.8%.

Numerically, joint optimisation reduces annual major-component logistics cost by 63.5% relative to a corrective benchmark and by 25.7% relative to a decoupled condition-based benchmark, raises the resilience index from 0.41 to 0.87 and cuts the 95% CVaR of energy not served by 44.6%. Hybrid storage strictly dominates single technology buffers once the storm severity multiplier exceeds 1.31, a threshold already exceeded in the Bay of Bengal calibration relevant to India’s emerging off-shore programme.

Four directions merit further work. First, the degradation model could be extended to dependent failures across components and turbines induced by common loading, which would strengthen the case for batch repair campaigns. Second, the vessel layer could incorporate a chartering option market, converting the first stage

fleet decision into a real options problem. Third, the metocean process could be coupled to downscaled climate projections so that κ becomes a stochastic trend rather than a stress parameter, allowing genuinely intertemporal resilience investment paths. Fourth, cross border coordination the pooling of strategic spares and heavy lift capacity across national off-shore zones, in the spirit of transnational off-shore energy interconnection offers substantial risk pooling gains that the present single farm model cannot capture while a coalitional game extension with Shapley based cost allocation is a natural next step.

Author contributions

Conceptualization, AKR; methodology, AKR; software, AKR and SD; validation, SD, MK and SKS; formal analysis, MK; investigation, SKS; data curation, SD; writing—original draft preparation, AKR and SD; writing—review and editing, AKR; visualization, AKR; supervision, SKS; project administration, AKR. All authors have read and agreed to the published version of the manuscript.

Declaration of competing interest

Scenario generation code, fitted metocean parameters and the figure generation scripts are available from the corresponding author on reasonable request.

Acknowledgments

The authors thank the Department of Statistics, Ravenshaw University, for computational support. No external funding was received.

Ethical approval

Not applicable.

Informed consent statement

Not applicable.

Conflict of interest

The authors declare no conflict of interest.

References

1. Ministry of New and Renewable Energy, Government of India. *Strategy for Establishment of Off-shore Wind Energy Projects: Zone Identification and Seabed Leasing*. New Delhi; 2024.
2. Shafiee M, Sørensen JD. Maintenance optimization and inspection planning of wind energy assets: Models, methods and strategies. *Reliability Engineering & System Safety* 2019; 192: 105993.
3. Sperstad IB, McAuliffe FD, Kolstad M, et al. Investigating key decision problems to optimize the operation and maintenance strategy of off-shore wind farms. *Energy Procedia* 2020; 137: 15-26.
4. Rinaldi G, Thies PR, Johanning L. Current status and future trends in the operation and maintenance of off-shore wind turbines: A review. *Energies* 2021; 14(9): 2484.
5. Li M, Jiang X, Carroll J, et al. A multi-objective maintenance strategy optimization framework for off-shore wind farms considering uncertainty. *Applied Energy* 2022; 321: 119284.
6. Stock-Williams C, Swamy SK. Automated daily maintenance planning for off-shore wind farms. *Renewable Energy* 2019; 133: 1393-1403.
7. Sahoo SS, Routray AK, Dehury S, et al. An inventory model for non-instantaneously reasoned, continuously deteriorating items under a linear demand rate. *Applied Mathematics on Science and Engineering* 2026; 3(2): 11–22. doi: 10.65036/amse322
8. Routray AK, Dehury S, Sahu SK. Three-warehouse inventory model for non-instantaneously deteriorating items with ramp-type demand, carbon cap and trade, exponential backlogging and two part trade credit under inflation. *Asian Journal of Probability and Statistics* 2026; 28(6): 43-64. doi: 10.9734/ajpas/2026/v28i6905.
9. Sheng S, Yang W. Wind turbine drivetrain condition monitoring: Current status and future outlook. *Wind Energy* 2023; 26(11): 1169-1190.

10. Fu S, Avdelidis NP, Zhang H. Digital twin architectures for off-shore wind operation and maintenance: A systematic review. *Ocean Engineering* 2025; 315: 119512.
11. Gao Z, Liu X. An overview on fault diagnosis, prognosis and resilient control for wind turbine systems. *Processes* 2021; 9(2): 300.
12. Si XS, Wang W, Hu CH, et al. Remaining useful life estimation-A review on the statistical data driven approaches. *European Journal of Operational Research* 2011; 213(1): 1-14.
13. Kirkpatrick K, Ren Z, Ma X. Transformer based prognostics for off-shore wind drivetrains with irregular sensor sampling. *Mechanical Systems and Signal Processing* 2025; 226: 112348.
14. Van Horenbeek A, Buré J, Cattrysse D, et al. Joint maintenance and inventory optimization systems: A review. *International Journal of Production Economics* 2013; 143(2): 499-508.
15. Karim R, Ahmed S, Kabir G. Prognostics driven spare parts management for off-shore renewable assets: A Bayesian approach. *Computers & Industrial Engineering* 2024; 190: 110047.
16. Egea-Álvarez A, Fekriasl S, Hassan F, et al. Advanced vector control for voltage source converters connected to weak grids. *IEEE Transactions on Power Systems* 2015; 30(6): 3072-3081.
17. Ryndzionek R, Sienkiewicz Ł. Evolution of the HVDC link connecting off-shore wind farms to onshore power systems. *Energies* 2020; 13(8): 1914.
18. Panteli M, Mancarella P, Trakas DN, et al. Metrics and quantification of operational and infrastructure resilience in power systems. *IEEE Transactions on Power Systems* 2017; 32(6): 4732-4742.
19. Bhusal N, Abdelmalak M, Kamruzzaman M, et al. (2020). Power system resilience: Current practices, challenges, and future directions. *IEEE Access* 2020; 8: 18064-18086.
20. Zhou Y, Panteli M, Wang B, et al. Quantifying the system-level resilience of thermal power generation to extreme temperatures and water scarcity. *IEEE Systems Journal* 2020; 14(1): 749-759.
21. Rockafellar RT, Uryasev S. Conditional value at risk for general loss distributions. *Journal of Banking & Finance* 2002; 26(7): 1443-1471.
22. Basciftci B, Ahmed S, Shen S. Distributionally robust facility location problem under decision-dependent stochastic demand. *European Journal of Operational Research* 2021; 292(2): 548-561.
23. Mohajerin Esfahani P, Kuhn D. Data-driven distributionally robust optimization using the Wasserstein metric. *Mathematical Programming* 2018; 171: 115-166.
24. Bertsimas D, Sim M. The price of robustness. *Operations Research* 2004; 52(1): 35-53.
25. Zhang C, Yang T, Gao W. Joint optimization of condition based maintenance and spare parts inventory for off-shore wind turbines under weather constraints. *Renewable Energy* 2024; 224: 120214.
26. Papadopoulos P, Coit DW, Aziz Ezzat A. Seizing opportunity: Maintenance optimization in off-shore wind farms considering accessibility, efficiency, and risk. *IIE Transactions* 2024; 56(6): 615-630.
27. Dinh QV, Pereira R, Nguyen T. Off-shore hydrogen production coupled with floating wind: Techno-economic architectures and system trade-offs. *International Journal of Hydrogen Energy* 2024; 58: 1210-1228.
28. Ibrahim OS, Singlitico A, Proskovics R, et al. Dedicated large scale floating off-shore wind to hydrogen: Assessing design variables in proposed typologies. *Renewable and Sustainable Energy Reviews* 2022; 160: 112310.
29. Jansen M, Duffy C, Green TC, et al. Island in the sea: The prospects and impacts of an off-shore wind power hub in the North Sea. *Advances in Applied Energy* 2011; 6: 100090.
30. Van Noortwijk JM. A survey of the application of gamma processes in maintenance. *Reliability Engineering & System Safety* 2009; 94(1): 2-21.
31. Zhu S, van Jaarsveld W, Dekker R. Spare parts inventory control based on maintenance planning. *Reliability Engineering & System Safety* 2020; 193: 106600.
32. Basten RJ, van Houtum GJ. System oriented inventory models for spare parts. *Surveys in Operations Research and Management Science* 2014; 19(1): 34-55.
33. Topkis DM. *Supermodularity and Complementarity*. Princeton University Press; 1998.
34. Nemhauser GL, Wolsey LA, Fisher ML. An analysis of approximations for maximizing sub-modular set functions I. *Mathematical Programming* 1978; 14: 265-294.
35. Hersbach H, Bell B, Berrisford P, et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* 2020; 146(730): 1999-2049.
36. Heitsch H, Römis W. Scenario reduction algorithms in stochastic programming. *Computational Optimization and Applications* 2003; 24: 187-206.
37. Magnanti TL, Wong RT. Accelerating Benders decomposition: Algorithmic enhancement and model selection criteria. *Operations Research* 1981; 29(3): 464-484.
38. Laporte G, Louveaux FV. The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters* 1993; 13(3): 133-142.
39. Kumar R, Stallard T, Stansby PK. Marine energy resource characterisation and extreme value analysis for the Indian coastline. *Renewable Energy* 2023; 210: 456-470.
40. Mohanty UC, Osuri KK, Pattanayak S, et al. Tropical cyclones over the Bay of Bengal: Trends, intensification and coastal impacts. *Natural Hazards* 2022; 112: 1-29.
41. Sherif YS, Smith ML. Optimal maintenance models for systems subject to failure-A review. *Naval Research Logistics* 2021; 68(4): 447-472.