Probabilistic Modelling of O&M Activities and Strategies in Wind Energy

by

Álvaro Cano

to obtain the degree of Master of Science in Aerospace Engineering
at Delft University of Technology and to obtain the degree of
Master of Science in Wind Energy at Technical University of Denmark.
To be defended publicly on Thursday 20th August, 2026 at 15:00.

Research carried out at:

Faculty of Aerospace Engineering, Delft University of Technology

Department of Wind and Energy Systems, Technical University of Denmark

Thesis committee:	
Chair:	Prof.dr. Simon Jonathan Watson
Supervisors:	Dr. Donatella Zappalá, TU Delft
	Prof. Nikolay Dimitrov, DTU
	Dr. Moritz Johann Gräfe, DTU
	Dr. Vasilis Pettas, DTU
Place:	Faculty of Aerospace Engineering, Delft
Project Duration:	November 2025 - July 2026
Student numbers:	6292763 (TU Delft), s244068 (DTU)

Abstract

Offshore wind operation and maintenance is modelled by two families of tools that do not meet. Discrete-event simulation carries the mechanisms that drive offshore downtime but costs many runs per converged answer, whereas the published analytical models answer for a fraction of the computational cost yet are exponential, steady-state, and window-based, and the effects of these assumptions have not been benchmarked against a modern simulator.

This thesis develops a hierarchy of analytical availability models that occupies that gap, establishing what each member can be trusted to estimate, at what accuracy, and at what cost. The family pairs steady-state and time-marching solvers under Continuous Time Markov Chain and a Semi-Markov Process frameworks, plus two simulator twins, all consuming one shared input pipeline. The reference case is a bottom-fixed 12 MW turbine with 22 corrective failure modes over twenty years of hourly ERA5 weather, and the discrete-event reference, NREL's WOMBAT (0.13.3), is audited, minimally patched, and characterised through a pre-registered mechanism-accretion ladder.

Under aligned conditions the steady-state semi-Markov model reproduces the reference raw, at $z = +0.20$ over 200 seeds and $z = -1.02$ against the thousand-seed population, certifying 97.9903% time-based availability, 97.8243% production-based availability, and an annual direct cost of 576,393.60 USD for the selected case study turbine. The memoryless simplification is evaluated and proves adequate for mean-determined steady-state quantities, although it cannot express production-based availability accurately, overstates the cold-start boundary credit by 46%, and is blind to the replacement waves that make a twenty-year life at $\beta = 2$ one long transient, missing 1.47 percentage points of year-one availability. Process interruptibility treatment is also evaluated, with results indicating a second-order convention worth 0.167 percentage points for the short repairs, and the enabling assumption for heavy replacement modelling, whose contiguous execution diverges at the reference site.

Finally, the 13 ms solve enables rapid scenario screening and sensitivity analysis, being used to propagate forty thousand joint uncertainty samples, price vessel capability, lifting envelopes, mobilisation lead time, and component reliability with break-evens attached, therefore turning the validated hierarchy into a decision instrument at a cost measured in seconds.

Acknowledgments

En primer lugar, quiero agradecer a mis padres y a mi familia todo el apoyo que me habéis dado a lo largo no solo del desarrollo de esta tesis, sino de toda mi carrera. En los momentos más difíciles me habéis dado ánimos y fuerza para seguir adelante, y a encontrar la confianza y determinación necesarias para superar cualquier obstáculo. Todos mis éxitos, pasados y futuros, son fruto de las oportunidades que me habéis brindado a costa de vuestro esfuerzo y sacrificio. Gracias por tanto.

I must also extend a profound thank you to my supervisors, Donatella, Moritz, and Vasilis, for their continued guidance, availability, and specially for their patience throughout the development of this thesis. Their advice and feedback were invaluable in helping me refine ideas and communicate them more effectively, as well as to overcome the scoping challenges that arose along the way, and overall bring this work to completion. I am truly grateful for all the time, support, and expertise they generously shared with me throughout this process.

Finally, I would like to thank all my friends and people who became part of this journey through Delft and Copenhagen and made it so special. During these last two years I've learned as much outside the university as within it, and I owe much of that growth to the friends I was fortunate enough to meet, with all the experiences we shared and memories we created along the way. Each of you has left a lasting mark on me, and wherever life takes us next, I will carry a part of you with me forever.

Contents

List of Figures

List of Tables

Nomenclature

The abbreviations, designators and symbols of this thesis are listed below. Every entry keeps the meaning fixed at its point of first use. The index i runs over the failure modes of the reference case.

Abbreviations

ABM	Agent-based model
AEP	Annual energy production
CapEx	Capital expenditure
CDF	Cumulative distribution function
CTMC	Continuous-time Markov chain
CTV	Crew transfer vessel
DES	Discrete-event simulation
ERA5	Fifth-generation ECMWF atmospheric reanalysis
FCR	Fixed charge rate
FOW	Floating offshore wind
HLV	Heavy-lift vessel
IEA	International Energy Agency
IEC	International Electrotechnical Commission
KPI	Key performance indicator
LCN	Large-crane capability class of the reference simulator, carried by the heavy-lift vessel
LCoE	Levelised cost of energy
MC	Monte Carlo
MDP	Markov decision process
MTBF	Mean time between failures
MTTF	Mean time to failure
MTTR	Mean time to repair
MUSD	Million United States dollars
NREL	National Renewable Energy Laboratory
O&M	Operation and maintenance
OpEx	Operating expenditure
P5, P50, P90, P95	5th, 50th, 90th and 95th percentiles
POMDP	Partially observable Markov decision process
pp	Percentage point(s)

SCADA	Supervisory control and data acquisition
SD	Standard deviation
SE	Standard error
SES	Surface-effect ship
SMDP	Semi-Markov decision process
SMP	Semi-Markov process
SOV	Service operation vessel
SPN	Stochastic Petri net
WMEP	Wissenschaftliches Mess- und Evaluierungsprogramm, the German onshore field measurement programme
WOMBAT	Windfarm Operations and Maintenance cost-Benefit Analysis Tool, the discrete-event reference simulator

Model and experiment designators

SS-C	Markovian description solved in steady state, the closed-form CTMC solver
TM-C	Markovian description marched in time, the Kolmogorov forward solver
SS-S	Semi-Markov description solved in steady state
TM-S	Semi-Markov description marched in time, the Markov renewal solver
MC-C	Monte Carlo sampler twin of the Markovian description
MC-S	Monte Carlo sampler twin of the semi-Markov description
B0–B3	Build ladder of the reference simulator, with B3 the certified build
ML-0 to ML-7	Rungs of the mechanism-accretion validation ladder
A-1 to A-11	Entries of the assumption register
RQ1 to RQ4	Research questions
O1 to O5	Thesis objectives
x10 to x19	Registered batch labels of the validation campaign, with x14 to x18 the five 200-seed reference batches and x19 the $\beta = 2$ anchor

Latin symbols

A	Time-based availability	[–]
A_E	Production-based availability	[–]
a_i	Operational age of mode i since its last renewal, age vector $\boldsymbol{a}$	[h]
C	Annual direct cost, cumulative $C(t)$, split into C^{eq} and C^{mat}	[USD/yr]

c_i^{day}, c_i^{mob}, c_i^{mat}	Vessel day rate, mobilisation fee and materials cost of mode i	[USD]
D	Mean regeneration cycle length	[h]
$\bar{D}(L)$	Start-weighted mean completion span at work content L	[h]
E_i^{down}	Potential-energy content of a mode-i down window, mean $\bar{E}_i^{\text{down}}$	[MWh]
E_{gross}, E_{net}	Gross and net annual energy production	[MWh/yr]
$e(h)$	Potential power of record hour h, long-run mean $\bar{e}$	[MW]
$F_L^{\text{span}}(t)$	Marginal completion-span law at work content L	[–]
$F_i^{\text{UP}}(u)$	Conditional up-time law given failing mode i	[–]
$g(s)$	Binary accessibility mask at record hour s	[–]
H_i	Active repair time of mode i	[h]
H_s	Significant wave height, record series $h_s(s)$, vessel limit h_s^{lim}	[m]
$h_i(a)$	Weibull hazard of mode i at operational age a	[h^{-1}]
J	Identity of the failing mode	[–]
L_i	Banked work content of mode i, $L_i = H_i + 2T_i$	[h]
ℓ_i	Transit leg of mode i	[h]
n	Number of failure modes	[–]
P_j	Limiting occupancy of state j, transition functions $P_{jk}(t)$	[–]
$\boldsymbol{p}(t)$	State-probability vector at time t	[–]
p	Electricity price	[USD/MWh]
$\mathbf{Q}$	CTMC generator matrix, off-diagonal entries q_{jk}	[h^{-1}]
$Q_{jk}(t)$	Semi-Markov kernel from state j to state k	[–]
q_i	Stationary mode share, $q_i = \Pr(J = i)$	[–]
$R_i(a)$	Weibull survival function of mode i	[–]
S_{rec}	Length of the weather record	[h]
$\mathcal{S}$	State space, with $\lvert\mathcal{S}\rvert = 1 + 2n$	[–]
$\text{span}(s; L)$	Completion span from start hour s at work content L	[h]
s	Record hour, start hour of a restoration	[h]
T_i	One-way transfer time of mode i	[h]
$\text{TC}(p)$	Total annual cost at electricity price p	[USD/yr]
t, u	Calendar time and elapsed time	[h]
U	Operational up-time to the next failure	[h]
$v(s)$	Wind speed at record hour s, vessel limit v^{lim}	[m/s]

$w^{\mathrm{cw}}(s; L)$	Contiguous-window wait from start hour s, mean $\bar{w}^{\mathrm{cw}}$	[h]
z	Standardised score, $z = (m - \bar{x})/\mathrm{SE}$	[–]

Greek symbols

α_i	Exit rate of wait state W_i, with $\alpha_i = 1/\tau_i^{\mathrm{W}}$	[h^{-1}]
β_i	Weibull shape parameter of mode i	[–]
Δ	Paired difference, as in ΔA	[–]
δ_{jk}	Kronecker delta	[–]
ε	Half-width tolerance in the seed-count relation	[–]
$\varepsilon_X(\theta)$	Elasticity of output X with respect to parameter θ	[–]
η_i	Weibull scale parameter of mode i	[h]
λ_i^h	Hourly failure rate of mode i	[h^{-1}]
μ_i	Restoration rate of mode i, with $\mu_i = 1/\tau_i^{\mathrm{R}}$	[h^{-1}]
ν_j	Stationary jump-chain probability of state j	[–]
π_j	Stationary CTMC probability of state j	[–]
$\rho(s)$	Start-hour weighting	[–]
σ	Standard deviation, log-scale width $\sigma_{\ln}$	[–]
τ_{UP}	Mean operational up-time	[h]
τ_i^{W}, τ_i^{R}	Mean wait and restoration sojourns of mode i	[h]
τ_i^{mob}	Mobilisation duration of mode i	[h]
ω_i	Task frequency of mode i	[h^{-1}]

Operators

$\mathbb{E}[\,\cdot\,]$	Expectation
$\Pr(\,\cdot\,)$	Probability
$\mathbf{1}\{\,\cdot\,\}$	Indicator function
$\Gamma(\,\cdot\,)$	Gamma function

Subscripts, superscripts and state labels

UP	Operating state
W_i	Wait state of mode i, mobilisation and transit leg
R_i	Restoration state of mode i
cw	Contiguous-window variant

day	Vessel day-rate basis
down	Down window
E	Energy basis, as in A_E
eq, mat	Equipment and materials cost components
gross, net	Before and after availability losses
h	Per-hour basis, as in λ_i^h
i, j, k	Failure-mode and state indices
lim, max	Vessel operating limit
mob	Mobilisation
rec	Weather record

1 Introduction

This introductory chapter provides the motivation behind the thesis and defines the problem it addresses. It first places offshore wind within the energy transition and locates its central weakness, a cost of energy that remains roughly twice that of its onshore counterpart. It then shows where operation and maintenance (O&M) sits within that cost, and describes how offshore maintenance is executed in practice, since those operational mechanisms are precisely what a credible O&M model must reproduce. Currently, the simulation models trusted for offshore O&M assessment are computationally too demanding for design-scale applications, whereas the analytical models suitable for such studies are constrained by several modelling assumptions and have not yet been validated against these higher-fidelity tools. Readers already familiar with the offshore O&M context may proceed directly to Section 1.2, from where the remainder of the chapter states the research problem formally before closing with the outline of the thesis.

1.1 Background and motivation

1.1.1 Offshore wind in the energy transition

Climate change has moved from an emerging international concern to a defining constraint on energy system development, with successive IPCC assessments stating that it is unequivocal that human influence has warmed the atmosphere, ocean and land [9]. This consensus underpins the objective of rapid decarbonisation formalised in the Paris Agreement, which aims to hold the increase in global average temperature "well below" 2°C and to pursue efforts to limit it to 1.5°C above pre-industrial levels [10]. Following this objective, the Renewable Energy Directive of the European Council has set a target of at least 42.5% renewable energy by 2030, as energy use and production account for 75% of total EU greenhouse gas emissions [11].

However, the transition is not driven by climate policy and decarbonisation alone. Increasing wind and solar capacity has been empirically shown to measurably reduce reliance on imported fossil fuels [12], and renewable deployment is therefore also framed as a matter of energy security and strategic autonomy. This security framing has become particularly prominent in Europe since 2022, with the REPowerEU plan of the European Commission explicitly aiming to end dependence on Russian fossil fuels by combining supply diversification with an accelerated transition to renewable energy [13, 14].

A third driver, arguably as decisive as policy support for deployment, is cost competitiveness. In their 2025 LCOE+ report, analysts from the asset management firm Lazard conclude that on an unsubsidised basis renewables remain a cost-competitive form of new generation [1]. The comparison in Figure 1 places onshore wind among the cheapest energy sources available, yet it also identifies the cost disadvantage motivating this thesis, as offshore wind remains roughly twice as expensive as its onshore counterpart and utility-scale solar.

Its continued expansion therefore depends on reducing this cost gap rather than maintaining an existing competitive advantage.

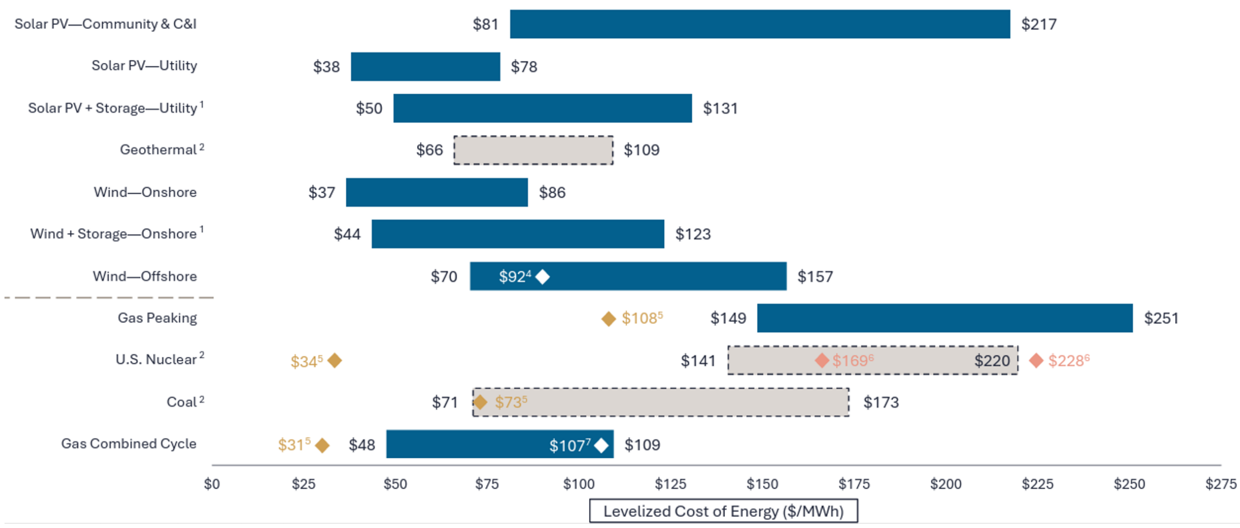

Figure 1: Illustrative levelised cost of energy (LCoE) ranges for selected generation technologies. Adapted from [1].

The industry attempting to close that gap has grown steadily. Global installed wind capacity increased from roughly 417 GW in 2015 to more than 1.1 TW by the end of 2024, while offshore capacity rose from less than 12 GW to almost 80 GW over the same period [15]. Industry outlooks anticipate continued strong growth through 2030, with offshore wind accounting for an increasing share of new capacity. However, the IEA notes that offshore deployment remains particularly vulnerable to permitting delays, vessel shortages, and grid constraints, leading to downward revisions of projections in some major markets [16, 17].

This expansion has been enabled by rapid technological scaling. In the United States, no turbine installed in 2013 had a rotor diameter of 115 m or more, while 98% of those installed in 2023 met or exceeded this threshold [18]. Offshore turbine ratings have consequently increased in parallel, with the average capacity of newly ordered offshore turbines reaching 14.8 MW in 2024 [19]. This scaling (illustrated in Figure 2) is central to offshore economics, as larger machines spread fixed installation and infrastructure costs over a greater energy output. However, it is also worth noting that this growth also amplifies the consequences of failures. Major components in turbines of these sizes can usually only be exchanged by specialised heavy-lift vessels, and a single stopped turbine now removes more than a dozen megawatts of generating capacity. Therefore, every failure event carries larger production losses and logistical demands than its onshore and smaller sized counterparts [20, 21]. This relationship between turbine scale and maintenance burden is examined further in Subsection 1.1.3.

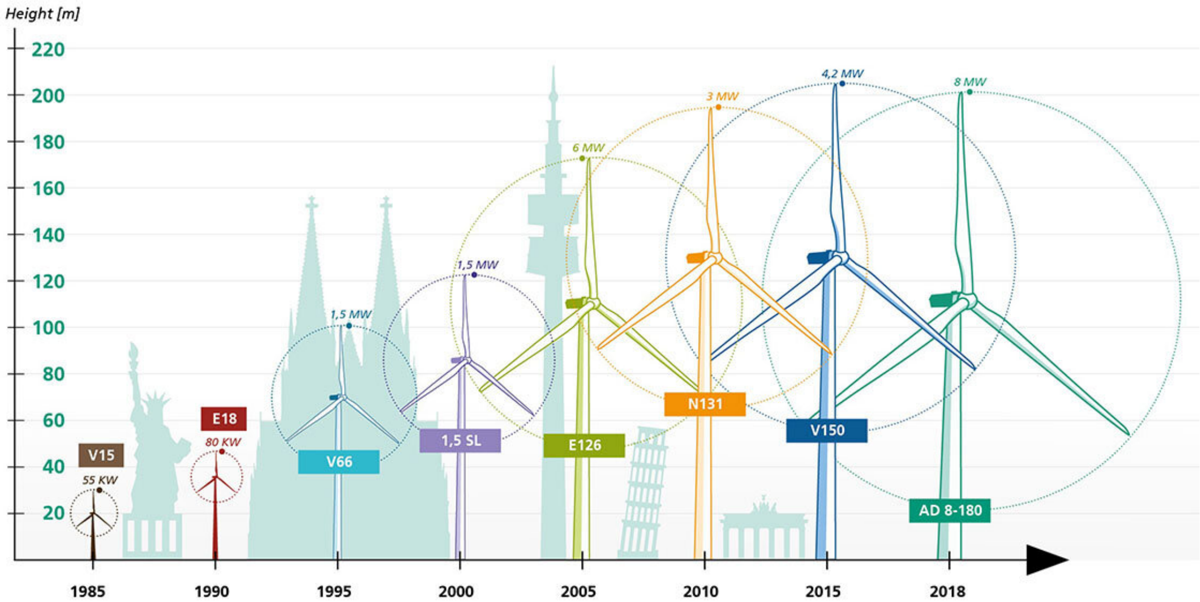

Figure 2: Evolution of turbine size with selected buildings and landmarks as reference. Adapted from [2].

Moving wind energy generation offshore is motivated by resource quality and siting feasibility, as winds at sea are stronger and steadier, with new projects expected to reach capacity factors of roughly 40–50% [22], and there is more space for larger turbines and farms. Furthermore, offshore generation allows production close to the coastal regions that host much of global demand, including 12 of the 15 largest cities in the world [23]. The marine environment, on the other hand, introduces new challenges like combined wind and wave loading, aggressive corrosion, and a complex electrical system of subsea cables and offshore substations with failure modes of their own [20]. Additionally, projects built farther from shore face longer transits and harsher met-ocean regimes [21].

Within offshore wind, two main categories can be distinguished based on the support structure used by the turbine. Bottom-fixed substructures dominate deployment and are expected to account for roughly 95% of global offshore additions between 2024 and 2033 [21, 24], while floating turbines accounted for only approximately 278 MW by the end of 2024 [17], and still carry reference cost estimates of about 181 $/MWh against the 117 $/MWh of their bottom-fixed counterparts [3]. Given this maturity gap, the specific engineering and logistics of floating offshore wind, such as mooring systems, dynamic cables, or tow-to-port heavy maintenance [25], lie beyond the scope of this work.

1.1.2 The cost of offshore wind and the role of O&M

From a whole-life perspective, offshore project expenditures are commonly grouped into development and consenting costs, capital expenditures (CapEx), operational expenditures (OpEx), and end-of-life and decommissioning liabilities (DecEx), with CapEx as the dominant block. A representative breakdown of the CapEx is shown in Figure 3a, obtained from the 2024 NREL Cost of Wind Energy Review, which prices a reference 600 MW plant built from 12 MW turbines. Although specific values vary by region and project, the breakdown highlights a key structural feature of offshore wind: a large fraction of project cost is tied to marine construction scope and electrical infrastructure rather than to the turbine alone [3]. These same site factors, namely water depth, distance to shore and grid, met-ocean severity, and supply-chain maturity, propagate not only into installation upfront costs, but also into the long-run operational and maintenance burden carried across the project life [21].

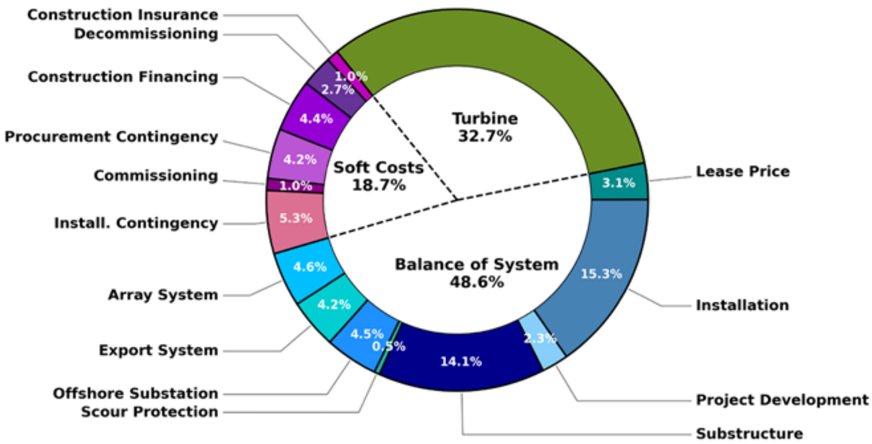

(a) Capital expenditure distribution.

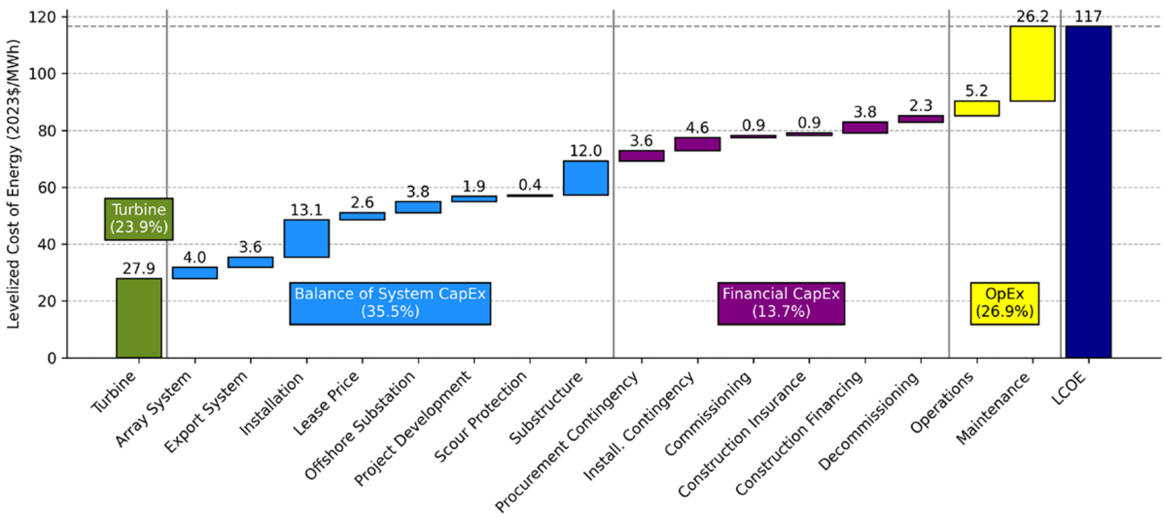

(b) Levelised cost breakdown.

Figure 3: Cost structure of the NREL 2024 reference fixed-bottom offshore wind farm [3].

The levelised cost of energy (LCoE) is a standard metric for comparing electricity generation options on a consistent \$/MWh basis. In its commonly used form, it can be written as [3]

$$\mathrm{LCoE} = \frac{\mathrm{CapEx} \cdot \mathrm{FCR} + \mathrm{OpEx}}{\mathrm{AEP}_{\mathrm{net}}}, \tag{1}$$

where FCR is a fixed charge rate reflecting financing and tax assumptions, OpEx represents annual operating expenditures, and $\mathrm{AEP}_{\mathrm{net}}$ is the net annual energy production.

Within this formulation, O&M affects the LCoE through two coupled channels. Firstly, it contributes directly to OpEx through labour, vessels, spares, logistics, and inspection or repair campaigns. Secondly, O&M decisions shape the $\mathrm{AEP}_{\mathrm{net}}$ via turbine availability, since failures, access delays and repair durations remove production hours. This joint cost-production coupling is one of the reasons offshore O&M is repeatedly identified as a central determinant of economic performance and a priority area for innovation and modelling [20]. Focusing on the direct channel alone, OpEx represents a 26.9% of the levelised cost for the reference plant, as displayed in Figure 3b.

This weight is likely to grow rather than shrink. As the technology matures and capital costs decline, the relative importance of the operating phase will increase, and industry analyses argue that O&M will rise to represent a large fraction of life-time costs [26, 27]. The cost reduction pathway in Figure 4 makes this dependence explicit, as part of the projected decline to 2035 is attributed to improved O&M strategies. Evaluating such strategies is, however, a modelling problem, because costs and availability must be assessed jointly and across many candidate configurations, which motivates the need for O&M models that are both faithful and computationally affordable [20, 25]. The next subsection discusses the processes these models intend to capture and represent.

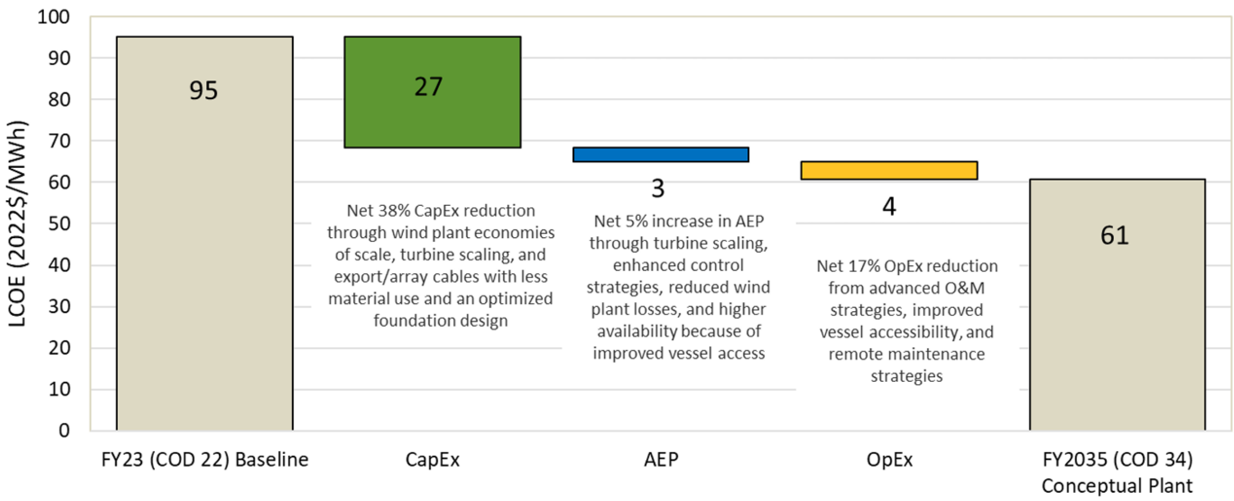

Figure 4: Cost reduction pathway from 2023 to 2035 for fixed-bottom offshore wind [3].

1.1.3 Offshore maintenance in practice

Offshore O&M can be viewed as a closed operational loop that starts with a failure or degradation signal and ends with return to service. Failures are detected through SCADA alarms, condition monitoring systems, or inspections, with operators aiming to maintain continuous monitoring capability [28]. Although certain events can be resolved remotely through manual resets, the remainder trigger physical mobilisation [20]. Operations teams then diagnose and classify each task by criticality, resource needs, and logistics, and a common planning framework classifies repair work into categories, from manual resets through minor, medium, and major repairs to major replacements. Each of these is associated with a typical time to repair and a required vessel and crew configuration [29]. Execution is finally conditioned by the weather, as limits on tolerable wind speed and significant wave height are converted into go or no-go decisions for specific vessels and tasks, therefore driving accessibility, which drives downtime and cost through waiting time [30].

The vessel fleet supporting this loop is organised in a hierarchy of capability and cost, summarised in Table 7. Routine interventions at short to medium distances are executed by crew transfer vessels (CTVs) operating from an onshore base, which suit inspections, resets, and minor repairs but are limited by safe transfer at the turbine interface and by daily shift patterns that reduce effective working time offshore [28, 29]. As distance to shore increases, service operation vessels (SOVs) provide accommodation, spares storage, and motion-compensated access at the price of higher day rates [30, 31]. Heavy maintenance, such as gearbox, generator, main bearing, or blade exchanges, requires heavy-lift jack-up vessels (HLVs), whose combined limits on wave height and wind speed make them the most weather-constrained assets in the fleet, and which are scarce, expensive, and chartered in discrete campaigns [29]. Mobilisation should consequently be understood as far more than transit, since it includes charter procurement, crew availability, port readiness and loading, sailing time, standby due to weather and, for heavy operations, jacking and lifting preparation, and it can extend from weeks to months depending on regional vessel supply and competing demand [29, 30]. It is also worth noting the spread of values visible in Table 7. Published charter rates differ by factors of two or more between the 2015 modelling reference cases and current industry guidance [8, 32]. Recent 15 MW baselines place heavy-lift day rates several times above the ranges reported a decade earlier [33, 34], and rates swing further with charter strategy, market cycle, and season [35]. For that reason O&M models treat vessel characteristics as explicit inputs rather than constants, and the values adopted for the reference case of this thesis are given in Chapter 3.

Table 7: Vessel classes supporting offshore corrective maintenance, with representative characteristics and the spread reported across sources. Rates are quoted in the currency of each source, and values are indicative. The exact parameters adopted for the reference case of this thesis are given in Chapter 4.

	CTV (crew transfer vessel)	**SOV** (service operation vessel)	**HLV** (heavy-lift jack-up vessel)
Typical tasks	Inspections, manual resets, minor repairs [28]	Minor-to-medium repairs at far-from-shore sites, onboard workshop and spares [30, 32]	Major component exchanges [29]
Access limits	$H_s \approx$ 1.5–2 m at crew transfer [28, 29], with 1.5 m the common modelling baseline [8, 34]	$H_s \approx$ 2.5 m gangway baseline [34], design envelopes extending towards 3 m [31]	$H_s \approx$ 2 m and wind $\approx$ 10 m/s for lifting [29], with H_s = 2.8 m used in a recent 15 MW baseline [34]
Shift and endurance	Day trips from port, shifts of ~12 h [28, 29]	On station for up to ~4 weeks, 24 h accommodation [32]	Continuous operation while on site during campaigns [29]
Mobilisation	None, port-based on continuous charter [8, 34]	Usually none, long-term charter based at the wind farm [32, 34]	Weeks to months of lead time [29, 30], 60 days with an EUR 2.1M fee in a recent 15 MW baseline [34], fees scaling with turbine rating [36]
Indicative charter rate	~ £1,750 per day in the 2015 reference cases [8], EUR 3,000 in a recent 15 MW baseline [34]	~ £9,500 per day as a 2015 field support vessel [8], ~30,000 (GBP or EUR) per day in current guidance and baselines [32, 34]	US$60,000–300,000 per day [33], EUR 420,500 in a recent 15 MW baseline [34], strongly dependent on strategy and season [35]

The consequence of these constraints is that offshore downtime is dominated by waiting rather than by active repair, as turbines stand still while weather windows, vessels, technicians, and spares are awaited [30, 25]. Long downtimes arise when suitable windows are sparse, when specialised vessels have long lead times, and when the failed component is critical, and downtime risk shifts towards the scarcer HLV-dependent events as failure burdens increase [29]. Subsea cable faults stand on the extreme, with average repair downtimes of roughly 40 days for array cables and 60 days for export cables [37, 38]. A second consequence is that offshore repairs are rarely completed in one visit, as work begins in a calm spell, pauses when the weather worsens, and resumes in the next viable window. Additionally, shift distributions or, for instance, labour laws that limit work to daytime, have a similar effect of stretching the calendar span of a task well beyond the working hours it contains.

An O&M strategy defines the trigger and the decision timing under which such interventions are planned, and five archetypes dominate the offshore literature: corrective maintenance (CM) acting only after failure, time-based preventive maintenance on fixed intervals, condition-based maintenance triggered by monitoring indicators, opportunistic maintenance bundling tasks into mobilisations that are already planned or taking place, and risk-based maintenance weighting failure probability by consequence [30, 39, 40, 41]. Reviews of the field agree that comparisons between them are only meaningful when weather-window feasibility and resource constraints are represented explicitly rather than simplified away [30, 42, 25, 43]. All strategies, however, operate on top of the same execution layer, since weather-driven access, mobilisation, and interruptible work constrain corrective and preventive interventions alike. For that reason this thesis is scoped to corrective maintenance, the setting in which that execution layer appears in its most direct form, and whose conclusions are expected to be most transferable to any potential future strategy studies built upon it.

1.2 Research problem

The operational picture above fixes what a credible model must represent. Most offshore downtime accrues while waiting for a weather window, for a vessel to mobilise, or for crew and spares, rather than during repair itself [30, 25], and repairs fragment across weather windows with progress retained between visits. Any model intended to support the techno-economic assessments motivated in Subsection 1.1.2 must therefore represent weather-driven access, mobilisation, and interruptible repair execution.

Offshore O&M models serve several purposes. Some optimise decisions, selecting vessel fleets or maintenance schedules, and others emulate slower models to accelerate such searches. All of them, however, rest on the same underlying capability: evaluating the availability, downtime, and cost of a given configuration [30, 25]. For a probabilistic model of the kind this thesis develops, that evaluation can be computed in only two ways, either by simulating the system and averaging over many sampled histories, or by solving the equations that govern its probabilities [44, 45]. These two routes have grown into two distinct tool groups, and Chapter 2 reviews both.

The first route is discrete-event simulation (DES), in which a tool recreates the whole operational history event by event, covering failures, weather, vessels, and repairs [8]. Although simulation is faithful to the real O&M activities, each run realises only one draw of the underlying random processes, so single runs are inherently subjected to sampling variability and are therefore not representative of the overall system. The averages of many repetitions are needed before the numbers settle, and the rare (and often expensive) failures converge slowest. Wrapping such a simulator inside an optimisation loop or a large sensitivity study therefore quickly becomes impractical [30, 25].

The second route is analytical modelling, in which those governing equations are solved directly, either in closed form or numerically, rather than sampled. Continuous-time Markov chains (CTMCs) and related stochastic processes compute availability, downtime, event counts, and cost this way, and consequently can run in fractions of a second. Recent work has started applying these methods to offshore wind [46]. However, as the review in Chapter 2 shows, the analytical models published so far rely on strong simplifications: they assume exponential distributions for all durations, they usually give only the long-run answers rather than time histories, and they often treat weather access as a single uninterrupted window, which Subsection 1.1.3 showed is not necessarily how offshore repairs are performed in reality. Furthermore, no analytical model in the literature has yet been tested and compared, mechanism by mechanism, against a modern simulator under carefully aligned inputs and conditions.

The field is consequently left with two tool groups that do not meet. Simulation is trusted and complete, but too expensive for studies that need thousands of model evaluations, while analytical models are fast and transparent, but rest on assumptions whose accuracy for offshore O&M has been asserted or doubted more often than it has been measured. This directly motivates the following open question: how much of the fidelity of a modern O&M simulator can an analytical model reproduce, at what fraction of the cost, and where exactly do the simplifications taken by the analytical model break down?

1.3 Approach and thesis outline

This thesis tackles the open question by developing four analytical availability models, two following a Markovian and another two a semi-Markov description of the system, with both approaches being solved in steady state and marched in time, thus obtaining four models. These are complemented with Monte Carlo sampler twins for cross-verification. The semi-Markov side exists because the two mechanisms at the centre of the question, ageing and weather-interrupted repair, produce durations that no exponential description can carry. The family of developed models is validated, mechanism by mechanism, against an aligned open-source simulator, and the study runs on a single 12 MW bottom-fixed reference turbine serviced by a CTV and an HLV under corrective maintenance.

The thesis is organised in six chapters:

- Chapter 2 (*Literature Review*) reviews the modelling landscape in depth, condenses the identified limitations into a capability matrix, formalises the research gap, and derives the research questions and the objectives that answer them.
- Chapter 3 (*Methodology and Model Development*) builds the four models and their samplers from a shared system definition, so that they differ by one structural modelling assumption at a time.

- Chapter 4 *(Model Validation)* presents the DES reference (the open-source WOMBAT simulator [5]), explains what was aligned and under which conditions comparison is fair, and executes the mechanism-led validation programme.
- Chapter 5 (*Results and Discussion*) uses the developed models and runs experiments to answer the research questions, from reference behaviour and the interruptibility comparison to sensitivity, uncertainty propagation, and application studies.
- Chapter 6 (*Conclusions*) answers the research questions explicitly, states the contributions of the thesis and their limits, and lays out suggestions for future work.

Figure 5 maps this argumental pipeline onto the chapters. The research gap is substantiated by the review before it is formalised, then the research questions derive from the gap and the objectives from the questions, and the bridge closing Chapter 2 maps the gap onto the model family that the remaining chapters build, validate, and employ. Table 9 in Chapter 2 later refines this pipeline to element level, tracing each research question to the gap elements it derives from and the objectives that answer it.

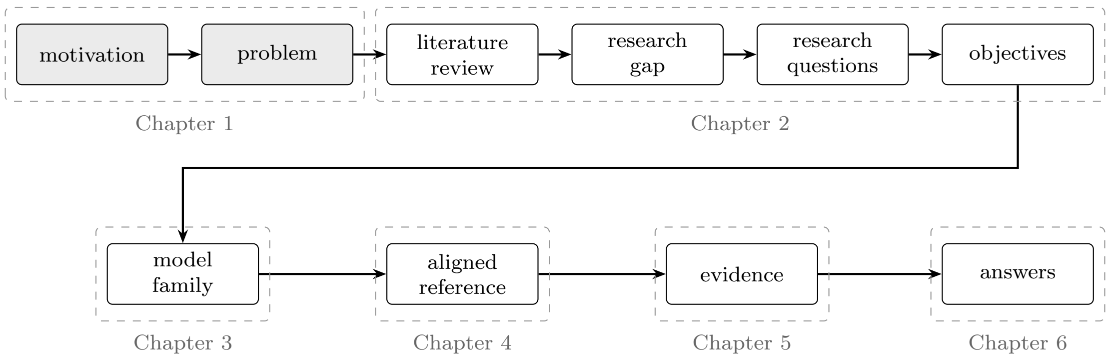

Figure 5: The argument of the thesis from motivation to conclusions, mapped onto its chapters, with shaded stages belonging to this chapter.

2 Literature review

2.1 Introduction and review approach

This chapter reviews how offshore wind O&M has been modelled and identifies the gap this thesis aims to fill. The review follows the two solution routes introduced in Section 1.2. The first route simulates the maintenance process event by event and averages over many sampled histories. The second solves the equations that govern the probabilities of the same process, mainly through Markov chains and their extensions. Figure 6 maps the modelling landscape covered in this chapter and shows where the thesis sits within it.

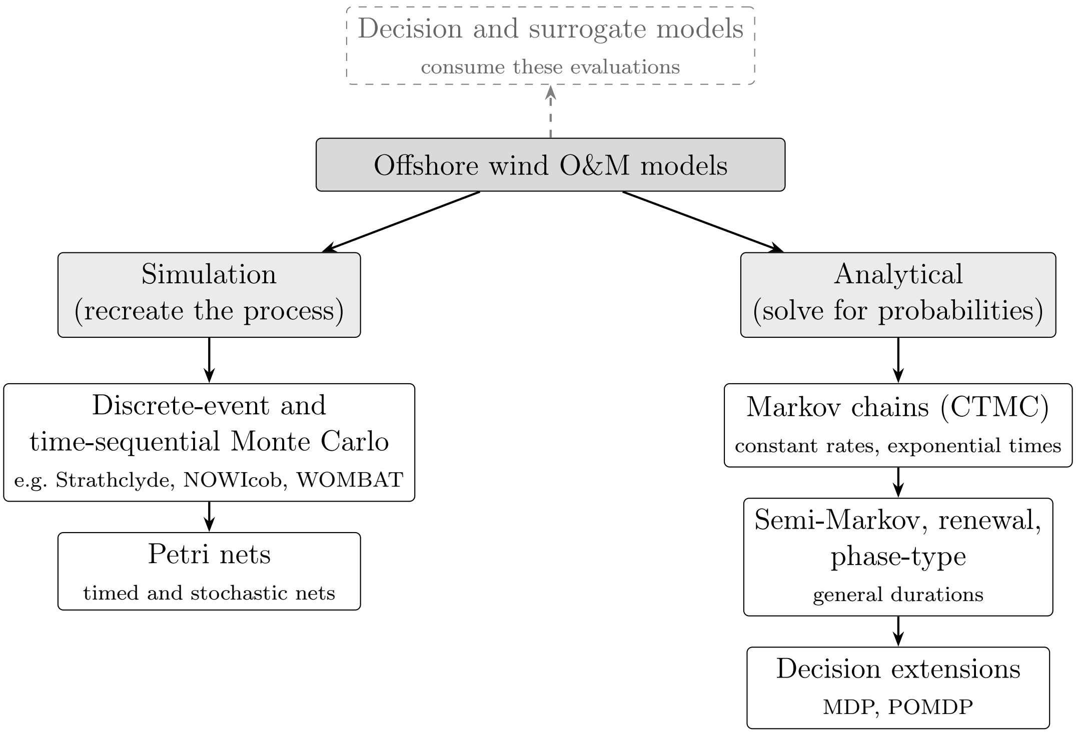

Figure 6: Map of the modelling landscape reviewed in this chapter. Simulation recreates the maintenance process event by event and averages sampled histories, whereas analytical models solve directly for state probabilities. Decision and surrogate models consume such evaluations and lie outside the scope of the review. This thesis develops the CTMC and semi-Markov branches and validates them against a simulation reference.

The review is bounded from two sides. Firstly, models that optimise decisions, such as vessel fleet selection, scheduling, or routing [47, 48], consume availability and cost evaluations rather than produce them. The same holds for surrogate models that emulate slower tools to accelerate such searches. Therefore, these models appear here only where they evidence the behaviour of the evaluators underneath them. The field itself separates these categories when it benchmarks, as one comparison exercise ran simulation tools, an optimisation tool, and an analytic spreadsheet tool on a single case [44]. Secondly, expected-value calculators of the spreadsheet kind belong to the solve route with the distributions discarded, and they are reviewed only in passing for the same reason.

The review proceeds as follows. The operational context comes first in Section 2.2, which explains what a corrective repair actually involves offshore, how the field defines availability, and why waiting dominates downtime. Simulation tools and their main weakness, namely computational cost, are covered in Section 2.3. Analytical Markov models follow in Section 2.4, first in general reliability theory and then in offshore wind. The models that relax the Markov assumptions, namely semi-Markov processes, renewal theory, and phase-type distributions, are treated in Section 2.5. The representation of weather, access, and interrupted repairs is reviewed in Section 2.6, the reliability data that parametrise every model in the field are examined in Section 2.7, and the verification and validation practice of the field is reviewed in Section 2.8. Finally, Section 2.9 condenses everything into a capability matrix and a five-element gap statement. From this gap, the research questions (Section 2.10) and the objectives that answer them (Section 2.11) are then derived, and Section 2.12 closes the chapter by mapping the gap onto the model family built in Chapter 3.

Regarding source selection, sources were gathered from Google Scholar and general search engines, using keyword families around offshore wind O&M, availability, Markov, semi-Markov, weather window, accessibility, and related terms. These searches were complemented by backward and forward citation chasing from the review papers analysed. Applied offshore work prioritised the use of recent sources, while the theoretical thread was followed to its classical sources regardless of age, since the mathematical groundwork this thesis uses dates from the 1950s and 1960s.

2.2 Offshore wind O&M and the role of availability

Chapter 1 established that offshore O&M is expensive and that most of the expense traces back to access. This section fixes the definitions and the field evidence that any model must answer to. Reference cost studies attribute a large share of the levelised cost of offshore energy to operations [49], and reviews of the field agree on the mechanism, as turbines can only be reached in suitable weather and with suitable vessels and both are scarce [50, 30]. Field data confirms where the time goes. Failure statistics from operating turbines show that a small number of major failures cause a large share of total downtime, because these failures require specialised equipment and servicing with long mobilisation lead times [4, 7, 51]. Figure 7 makes this relation visible.

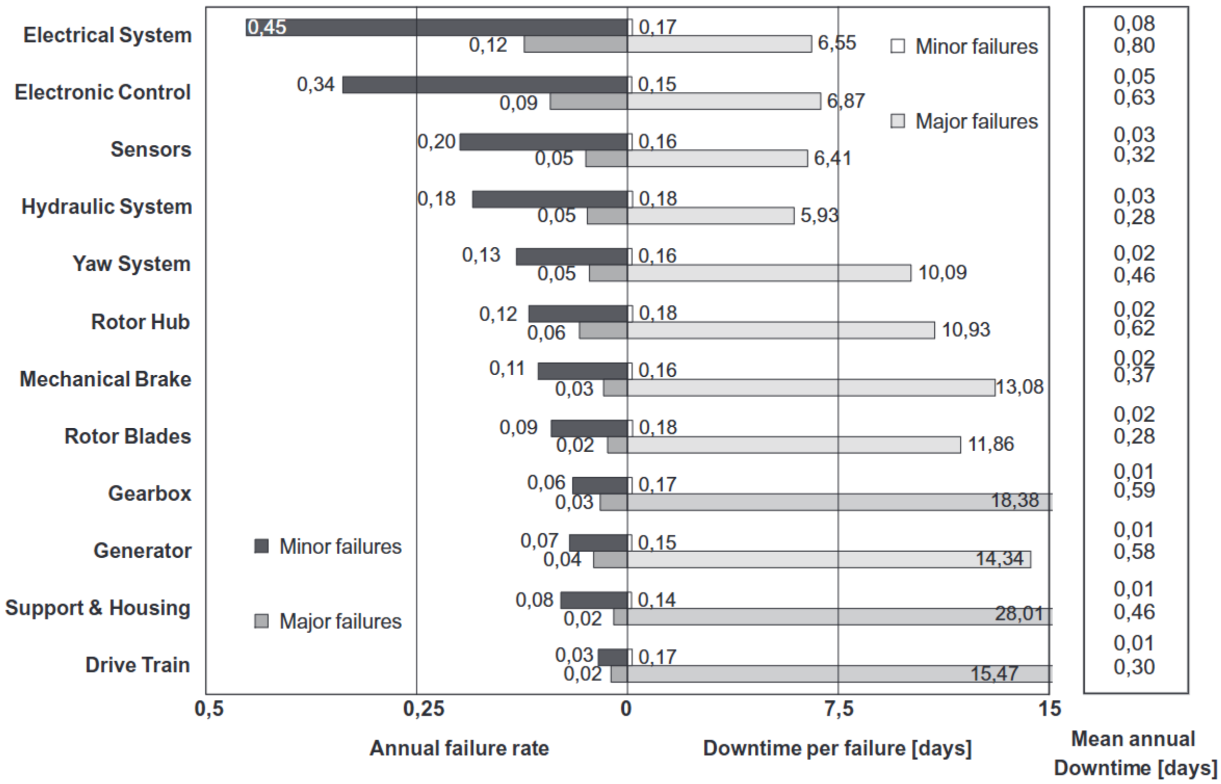

Figure 7: Annual failure rate (left), downtime per failure (right), and mean annual downtime (outer column) per subassembly component, split into minor and major failures, from the onshore WMEP field programme. Frequently failing subassemblies are repaired in fractions of a day, while the rarer major failures take several days each, so total downtime concentrates in the rare heavy events. Although this data is from onshore turbines, the authors claim that the pattern is expected to strengthen offshore. Extracted from [4].

Figure 8 shows the anatomy of a single corrective repair episode. A component fails and, if a chartered vessel is needed, mobilisation takes days to weeks. The vessel then sails to the turbine and must transfer crew, which requires a calm interval, after which work proceeds only in accessible weather. In this example, when conditions worsen mid-repair, the technicians do not remain at the turbine. Instead, they transfer back to the vessel, work pauses with the progress kept, and when conditions improve then a new transfer opens and the next work fragment is completed. Finally, once the work is complete, the crew transfers off and the vessel returns to port. Although the figure is for representational purposes only and block durations are not based on real data, the figure aims to highlight how the hands-on repair (the dark segments) can be a small part of the total downtime, and the rest is waiting and logistics.

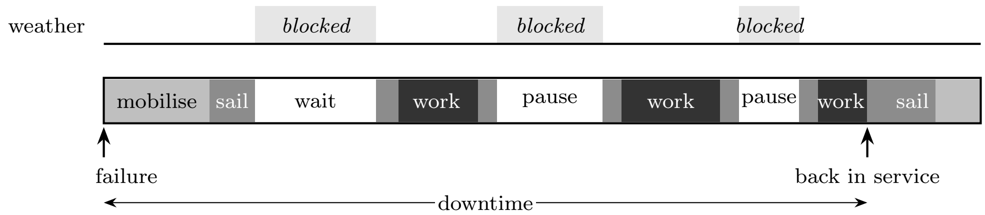

Figure 8: Anatomy of an example corrective repair episode. Work (dark) proceeds only in accessible weather and pauses when the weather worsens, with progress kept between fragments. Here, crews do not remain at the turbine through blocked weather, so every work fragment begins and ends with a crew transfer (medium grey).

Two consequences follow for modelling. Firstly, the metric that matters is availability with its full delay structure, not component reliability alone. Secondly, the delays are strongly non-exponential, since a duration is only memoryless when the distribution of its remaining time does not depend on how long it has already lasted. This property defines the exponential distribution and enables constant-rate Markov models. Mobilisation violates it in one direction, as a lead time quoted at fourteen days has one day remaining after thirteen days have elapsed, not a fresh fourteen-day draw. Weather-driven waiting violates it in the other direction, since sea states persist and follow seasons, so the time already spent waiting carries real information about the time still to come [52].

The outcome of all these delays is expressed through availability, for which two distinct definitions exist, both standardised in IEC TS 61400-26 [53]. Time-based availability A is the fraction of calendar time in which the turbine is technically capable of operating. Production-based availability A_E weights that capability by the energy that could be produced, so it is the fraction of potential energy actually captured. Offshore, the two diverge for a physical reason that unfolds in three steps. Storms halt maintenance, so downtime accumulates preferentially in stormy periods. However, stormy hours also carry the strongest winds, and therefore the largest share of producible energy. By this relation, downtime hours are consequently worth more energy than average hours, which makes production-based availability A_E systematically lower than time-based availability A offshore. Any model that reports only A or computes energy production from it therefore hides part of the economic picture, and this thesis aims to compute both parameters and treat their gap as a quantifiable output.

The literature attaches to availability-focused decision-making questions like which vessel fleet to charter and under which strategy [54], which access thresholds and shift rules to apply, or which O&M cost and downtime consequences follow from siting and technology choices. Sperstad et al. identify the decision variables to which simulated O&M results are most sensitive, and weather access limits ranked at the top [55]. A model intended for these questions must therefore resolve weather-driven waiting well, and it should ideally be affordable enough to evaluate across many candidate strategies. This pair of requirements frames everything that follows.

2.3 Simulation-based O&M modelling

2.3.1 The discrete-event and Monte Carlo approach

The dominant response to this complexity has been to simulate it directly. A discrete-event simulation (DES) keeps a clock and a list of scheduled future events. It repeatedly takes the earliest event, advances the clock to it, updates the system state, and schedules the events that follow, as shown in Figure 9. Event times can come from any distribution, and activities can be paused and resumed because progress is trackable. Therefore, a DES model can reproduce an episode like the one in Figure 8 directly.

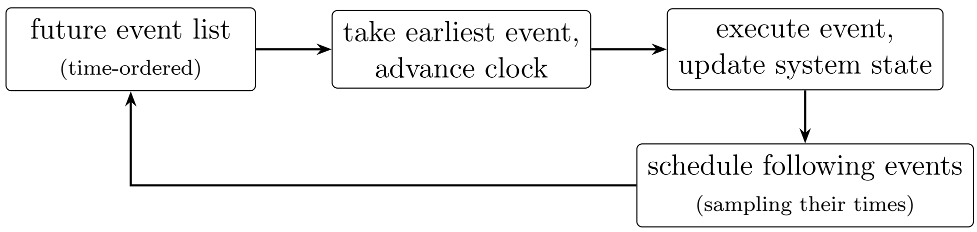

Figure 9: The discrete-event simulation loop. The clock jumps from event to event, and event times are sampled from arbitrary distributions, so no memoryless assumption is needed.

The offshore wind DES approach began with academic tools, whose first generation was catalogued in the review of Hofmann [56]. The Strathclyde model from Dinwoodie set the time-sequential template, combining sampled failures, a metocean time series, repair classes, and vessel logistics [57], while the model NOWIcob added rich vessel-fleet and logistics options [58]. Petri nets offered an alternative event-driven representation for the same processes [59, 60]. Around these cores, a large applied literature used simulation to study specific decisions, such as maintenance-strategy comparison [61], vessel and transport selection [54], spare-parts logistics [62], technician routing [48], and fleet sizing, often by coupling the simulator to an optimiser [47]. Furthermore, Sperstad et al. and most recently Huang et al. showed how strongly the answers to such decision problems depend on modelling choices inside the tools [55, 63].

The current DES reference point in this thesis is WOMBAT, an open-source tool from NREL [5], whose handling of a single event is illustrated in Figure 10. Turbines, their subassemblies, and the servicing vessels run as concurrent processes on a shared event calendar. Failures raise repair requests, and vessels mobilise, sail, transfer crew in weather windows, and execute repairs with pause-and-resume mechanics. The main output is an event log from which availability, downtime, and cost can be computed. WOMBAT matters for this thesis for two reasons. Firstly, it is open-source and can be locally edited, which the audit of Section 4.2 takes advantage of. Secondly, it implements the full mechanism set that the analytical models approximate. For these reasons it serves as the code-to-code validation reference throughout the thesis.

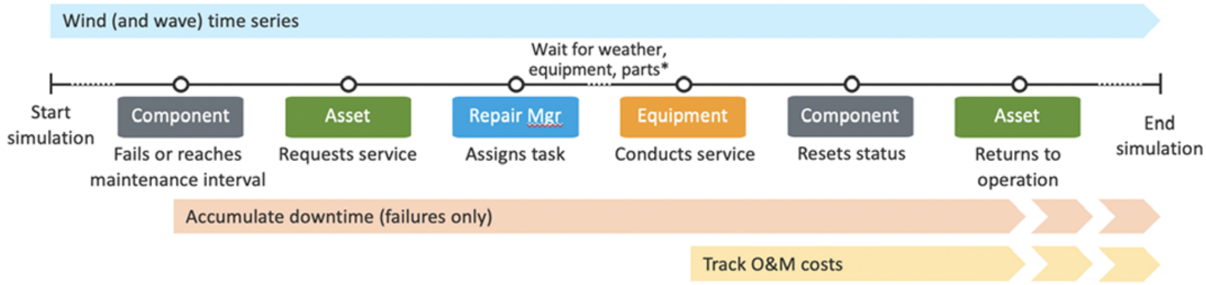

Figure 10: Conceptual overview of a single event within the simulation architecture of WOMBAT. Extracted from [5].

2.3.2 The computational-cost critique

The weakness of the DES approach is not fidelity but cost, and the reviews of the field say so consistently [30, 64]. The cost has a simple statistical root. One simulation run is one random sample path, useful numbers are obtained from averages of many runs, and their statistical error shrinks only with the square root of the number of runs. Therefore, halving the error may cost four times the computation. Furthermore, the rare and expensive events, namely the major replacements that dominate cost, converge slowest and consequently set the budget.

The cost becomes prohibitive exactly where fast models are most wanted, that is, in the many-query settings that Section 1.2 anticipated. Optimising a maintenance strategy, sweeping a design parameter, or propagating input uncertainty each require anywhere between dozens to thousands of scenario evaluations, and each evaluation may already require hundreds or thousands of individual simulations. Studies that couple a simulator to an optimiser feel this cost directly [47, 65]. Furthermore, this acts directly against the stated primary purpose of WOMBAT, which is "to help users compare O&M scenarios" [5]. A second, quieter weakness is the parameter burden. Detailed simulators need many inputs while offshore wind data is scarce, so a heavily parametrised model can look realistic while relying on guesses [8], similarly to the simpler analytical models. Both weaknesses point in the same direction, towards fast and transparent models that keep the decision-relevant behaviour, provided such models can be shown to be accurate. That condition is the subject of the rest of this chapter.

The conclusion this review draws from the DES modelling approach is therefore deliberately double-edged. On the one hand, DES is the right reference, because it carries the relevant mechanisms directly and its remaining error is quantifiable sampling noise. On the other hand, DES is the wrong screening tool, because the studies that create the most engineering value need evaluation counts it cannot afford [55]. The role as a reference is further developed in Chapter 4, while the role as a fast screening tool is what the analytical family developed in this thesis must earn.

2.4 Analytical reliability and availability modelling

2.4.1 Markov models in reliability theory

The analytical alternative is rooted in reliability engineering. A Markov model describes a system by a finite set of states (for example: operating, waiting, and under repair), together with constant rates for jumping between them. Its defining assumption is that the future depends only on the current state, and is therefore independent from how long the system has been in it. From this assumption everything of interest follows by linear algebra. State probabilities over time come from a matrix exponential, and long-run probabilities come from a small linear system [66, 67], from which availability, expected downtime, event frequencies, and costs then follow directly [68, 69]. In this formulation, the system is described by a stochastic process whose probability distribution evolves deterministically. Analytical Markov models solve for that distribution directly, whereas Monte Carlo models approximate it through repeated sampling. Reliable numerical methods for these computations are long established [70]. Nevertheless, the price of this convenience is fixed by the mathematics, as the Markov assumption forces the memoryless property and therefore every duration in the model to be exponential.

2.4.2 Analytical models of offshore wind availability

Wind energy adopted Markov models early, mostly for maintenance planning and cost-rate analysis [71, 65]. Huang et al. applied a CTMC to wind-farm reliability at the system level [72], analytical availability expressions have also been built for the electrical collection and transmission system [73, 74], and time-window formulations brought maintenance-access constraints into analytical planning models [75].

However, the closest antecedent to this thesis is the model of Centeno-Telleria et al. [46], which has later been used for geospatial screening studies [76] and for tow-to-port modelling for floating wind. Their model represents each failure mode by a small chain of the type shown in Figure 11. The turbine is up, a failure moves it into a waiting state covering mobilisation and access delay, a repair state follows, and the turbine then returns to service. Solving the chain gives availability and O&M cost far faster than simulation, which is precisely what their geospatial siting studies needed. Their own validation illustrates both the promise of the route and the open problem. The authors compared the model against two simulation-based case studies reproduced from the literature, one on floating wind and one on wave energy, and reported matching their endpoint availability and cost figures at a computational burden at least five orders of magnitude lower [46]. A comparison of that kind corroborates the speed claim. However, reproducing published cases cannot align all the specific input configuration (particularly the metocean data), cannot isolate which assumption causes which residual, and cannot probe the regimes the assumptions exclude.

Therefore, it validates the model inside its assumption envelope without measuring where the envelope breaks. Furthermore, while the authors claimed that the models showed "good agreement", a 10% OpEx difference and a 3% availability gap were measured, the latter hiding a doubling of the downtime in the FOW case and a sevenfold increase in the wave converter case [77]. These gaps highlight modelling discrepancies that cannot be ignored, and that must be attributed to a modelling choice or assumption. Moreover, since all durations are exponential, only the steady-state or long-run solutions are obtained and there are no time histories. Additionally, weather enters as a single access delay before the repair, a single-window treatment in which, once work starts, it is expected to run uninterrupted. The pause-and-resume behaviour of Figure 8 therefore lies outside the model class. Finally, the AEP is computed by multiplying the potential power production by the time-based availability A, rather than the production-based availability, which is not obtained [46].

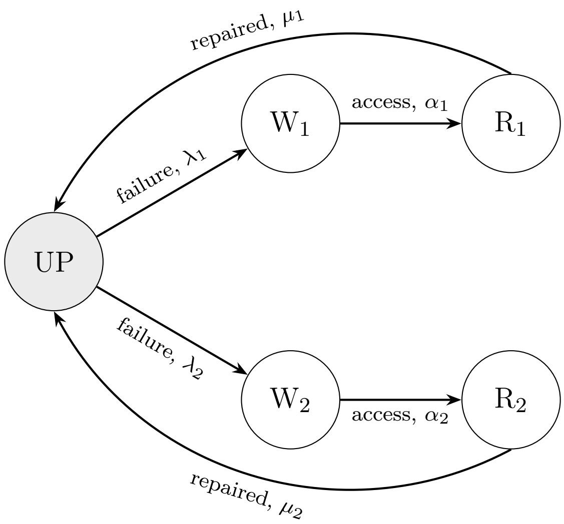

Figure 11: Typical structure of the Markov chain models in offshore wind literature, here with two failure modes. Each arrow carries a constant transition rate, so every duration is exponential.

The overall picture of the analytical wind literature is therefore consistent. The models are fast and transparent, and they have proven useful for screening and planning. However, they share three restrictions, anticipated in Section 1.2. Firstly, every stochastic duration is exponential, covering times to failure, mobilisation and access delays, and repair times alike. Secondly, only long-run steady-state answers are produced, with no time histories. Thirdly, weather receives the single-window treatment just described, in which access is one delay and execution is never interrupted. Whether these restrictions matter, and by how much, has not been measured against a mechanism-complete reference, and that measurement is what this thesis provides.

2.5 Semi-Markov, renewal, and general-distribution approaches

2.5.1 Semi-Markov processes and their reliability applications

Applied probability offers a well-understood way to keep the state-space structure of a Markov model while dropping the exponential restriction. In a semi-Markov process (SMP), the sequence of visited states is still a Markov chain, but the time spent in each state can follow any distribution. The theory goes back to Pyke, to the regenerative processes of Smith, and to the Markov renewal theory of Çinlar [78, 79, 80]. It is consolidated in the monographs of Limnios and Oprişan and of Grabski, and in the standard reliability curriculum of Rausand and Høyland [81, 82, 83]. Two results matter most in practice. Firstly, the long-run state probabilities have a simple closed form, in which the share of time spent in each state is proportional to how often it is visited, multiplied by how long a visit lasts on average. This implies that steady-state availability depends on the holding-time distributions only through their means, which entails that steady-state solutions can never certify distributional structure. Secondly, the time-dependent probabilities solve a system of integral equations for which standard numerical schemes exist [84, 85]. The classic treatment from Howard connects both results to rewards and costs [69].

2.5.2 Semi-Markov and general-distribution models in wind energy

Wind applications of these ideas exist, but they are scattered and partial. D'Amico et al. used semi-Markov chains to model the wind-speed process itself [86], Dawid et al. applied a semi-Markov description to offshore access states [87], and Welte et al. advocated phase-type stage models for deterioration and maintenance durations in wind assets [88]. What is missing from all of them is the combination this thesis targets: a semi-Markov model of the full availability cycle, covering failure, logistics, weather, interrupted repair, and restoration, solved both in time and in steady state, and validated against a simulator. The published wind SMP work treats isolated pieces of the problem, none of it carrying all the identified mechanisms or having been benchmarked against a DES reference.

2.6 Modelling weather, accessibility, and interruptibility

2.6.1 Weather and accessibility representations

Weather enters O&M models through accessibility, since an hour is determined as workable only when wind and wave conditions are below the limits imposed by the vessel and task. The classical analytical treatment asks for a full weather window, that is, an uninterrupted run of workable hours at least as long as the task. Persistence analyses characterise how long calm spells last [52], and Feuchtwang and Infield derived closed-form delay distributions for waiting on a window [6], a formulation whose consequences are illustrated in Figure 12. Later work refined access modelling with multivariate sea-state criteria and richer metocean statistics [89, 90, 91], uncertainty-aware window forecasting [92], and stochastic weather generators for simulation input [93]. Data-collection practice for these inputs is documented in the industry literature [94].

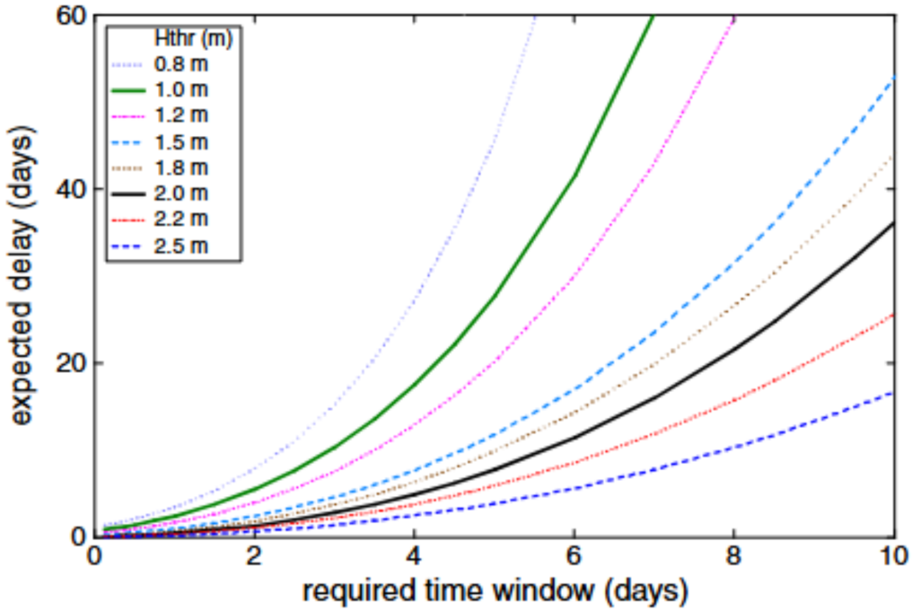

Figure 12: Expected access delay against the required contiguous window length (repair time) at the Barrow offshore wind farm, for a family of significant wave height thresholds, computed with the closed-form window formulation. Extracted from [6].

2.6.2 The interruptibility phenomenon and its treatment

The window formulation contains a strong hidden assumption, namely that the task needs one contiguous span of calm weather. Real offshore practice, and the reference DES, does something different. Work starts as soon as any workable interval opens, pauses when the weather closes (or the shift ends), and resumes later with the progress kept. Figure 13 contrasts the two readings, and the gap between them widens with task length. For long tasks in harsh climates, a single window of the required length may be rare, while the same work content is easily accumulated across several shorter spells. Therefore, the contiguous assumption overstates waiting and downtime, while fragmentation stretches the calendar duration of the repair itself.

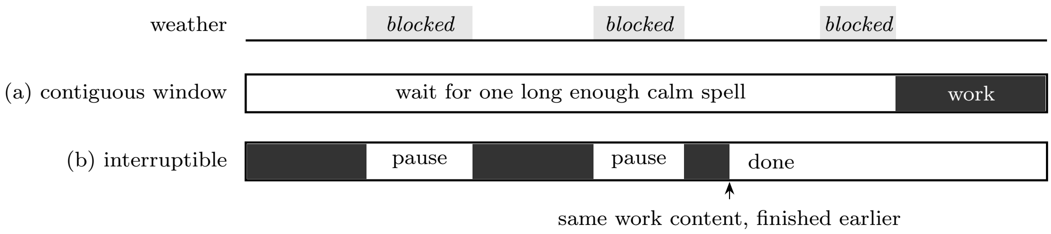

Figure 13: Two readings of weather-constrained work, on the same weather sequence. (a) The classical window assumption waits for one uninterrupted calm spell of the full task length, after which the work runs without interruption. (b) Interruptible execution accumulates the same work across shorter spells, pausing and resuming with progress kept. Usually, simulators implement (b), and the analytical literature assumes (a).

The treatment of this mechanism splits the literature quite cleanly. On the simulation side, pause-and-resume is standard, as the Strathclyde model, NOWIcob, and WOMBAT all accumulate repair progress through accessible hours [57, 58, 5]. WOMBAT makes the distinction explicit in its source code, since crew transfer and vessel travel demand an uninterrupted window while the repair itself accumulates good hours across interruptions (Section 4.2). On the analytical side, the window-based delay models stop at access [6], and the availability chains of Subsection 2.4.2 treat repair as an uninterrupted block [46]. Interruptibility also reshapes the duration distributions themselves. The calendar duration of an interrupted repair is the sum of its net work content and the closed-weather gaps that intersect it, producing a seasonal, weather-dependent relation that no exponential can express and that cannot be specified independently of the weather record. A model that wants these durations must therefore construct them from the metocean series itself, an observation that motivates the weather-walk kernel of Chapter 3.

2.7 Reliability data and failure statistics

Every O&M model, simulated or analytical, stands on a small set of published reliability datasets, and the review must establish both what they provide and what they can legitimately parametrise. The anchor offshore dataset is that of Carroll, McDonald, and McMillan [7]. It reports failure rates (Figure 14), repair times, and repair costs per subassembly in three material-cost severity bins, namely minor repair, major repair, and major replacement, with events lacking cost information kept as a separate category, all from an offshore sample of roughly 350 turbine years.

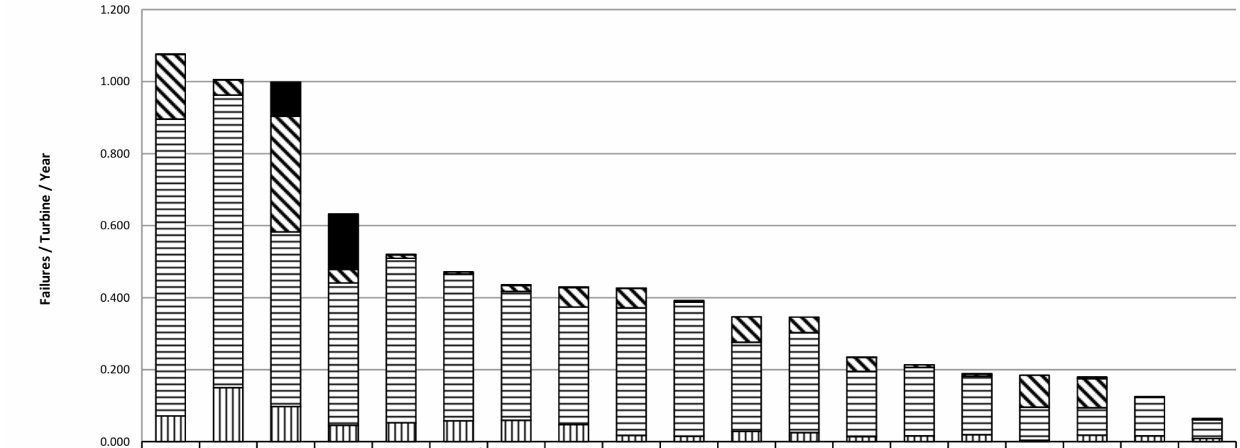

	Pitch / Hyd	Other Components	Generator	Gearbox	Blades	Grease / Oil / Cooling Liq.	Electrical Components	Contactor / Circuit Breaker / Relay	Controls	Safety	Sensors	Pumps/ Motors	Hub	Heaters / Coolers	Yaw System	Tower / Foundation	Power Supply / Converter	Service Items	Transformer
Major Replacement	0.001	0.001	0.095	0.154	0.001	0.000	0.002	0.002	0.001	0.000	0.000	0.000	0.001	0.000	0.001	0.000	0.005	0.000	0.001
Major Repair	0.179	0.042	0.321	0.038	0.010	0.006	0.016	0.054	0.054	0.004	0.070	0.043	0.038	0.007	0.006	0.089	0.081	0.001	0.003
Minor Repair	0.824	0.812	0.485	0.395	0.456	0.407	0.358	0.326	0.355	0.373	0.247	0.278	0.182	0.190	0.162	0.092	0.076	0.108	0.052
No Cost Data	0.072	0.150	0.098	0.046	0.053	0.058	0.059	0.048	0.018	0.015	0.029	0.025	0.014	0.016	0.020	0.004	0.018	0.016	0.009

Figure 14: Failure rates per turbine per year by subassembly and severity category. Extracted from [7].

Two properties of such datasets matter for correct use. Firstly, the rates are reported per turbine per calendar year, whereas an availability model may consume rates per operating hour, and the conversion through 8,760 hours per year introduces a calendar-versus-operational gap that must be handled explicitly (Chapter 3). Secondly, the datasets report rates and mean durations rather than duration distributions. Therefore, any distributional assumption placed on top of them is exactly that, an assumption, and the remainder of this section asks which distributional claims the field evidence can actually support.

The onshore literature supplies the long-horizon statistics. Tavner et al. and Spinato et al. analyse large European datasets and report subassembly failure intensities, including the familiar early-life and ageing patterns [95, 96]. Faulstich et al. fit an additive multi-mechanism hazard and find that it improves on a constant rate only marginally on held-out data, an underappreciated result that supports the exponential base case [97]. Additionally, Harzendorf et al. connect drivetrain concepts to failure behaviour and O&M cost consequences [98].

The most recent large-sample evidence spans both environments. Walgern et al. analyse maintenance reports from more than 1,000 onshore and offshore turbines covering over 4,200 operating years, using Nelson–Aalen plots for failure trends over operating age and Poisson-process regression for design and technology effects [99]. They find that reliability generally improves with later commissioning years, that higher rated powers carry higher failure intensities, and that offshore turbines are generally more reliable than onshore ones, with the yaw system as the exception. They also find that design choices such as the drive-train concept measurably shift subsystem failure intensities. The rated-power finding sharpens the scale argument of Chapter 1, since the largest machines not only lose more per failure but also fail more often. However, like most of the field statistics above, their fitted quantity is the failure intensity of a repairable fleet, and whether such intensities can supply renewal kernels is a question of estimands, addressed next.

Using these sources to justify non-exponential failure clocks requires an estimand check that the wind literature routinely skips. A shape fitted to the rate of occurrence of failures against calendar or operating age in a fleet describes a repairable population whose parts remain, at each failure, essentially as old as before, the classical setting of Ascher and Feingold and of Lindqvist [100, 101]. As it will be shown later in Chapter 3, the semi-Markov kernel of this thesis needs a different object. It needs the time-to-failure distribution of one failure mode on its own operational clock, renewed to as-good-as-new at each repair, a restoration assumption positioned within the virtual-age framework of Kijima [102]. The two estimands coincide only in the exponential case. Furthermore, fleet heterogeneity alone can manufacture apparently decreasing intensities, as Slimacek and Lindqvist demonstrate with explicit frailty modelling, so published shapes below one are largely artefacts of pooling [103].

The only field-fitted shape located by the review that passes this check is the drivetrain renewal value $\beta = 1.43$ of Shrivastava and Baghel, from a small population and a modest venue [104], while Asgarpour and Sørensen offer directional support for wear-type ageing in mechanical components without fitted values [105].

The consequence for this thesis is that the baseline build sets $\beta = 1$ for all failure modes, which matches what the data can actually defend and keeps the validated reference intact. Shape effects are then studied separately as controlled arms, with a wear arm carrying the one defensible field shape plus conventional wear values on the replacement bins, and an adversarial arm with mild infant-mortality shapes. The invariance property of the steady-state solution guarantees that all shape effects appear only in the transient, and Chapter 5 executes both arms.

2.8 Verification, validation, and benchmarking practice

How does this field decide whether a model is right? High-quality field data at the event level is scarce, so the community has relied on code comparison, that is, running different tools on the same defined cases and comparing outputs. The most prominent exercise is the case study of Dinwoodie et al., in which four simulation tools ran identical scenarios [8], followed by the IEA Wind expert comparison [106]. One related exercise even reached across categories. It ran four simulation tools, one optimisation tool, and one analytic spreadsheet tool on a common reference farm, and found only partial agreement on the resulting vessel-fleet rankings [44]. Two findings from these exercises shape the present work. Firstly, even well-built tools disagree when their internal conventions differ, so input and accounting alignment is a precondition for any meaningful comparison, and the spread in Figure 15 shows the size of the effect. Secondly, the largest disagreements trace to exactly the servicing-protocol details discussed above, namely shift handling, window logic, and repair execution [44]. Sperstad et al. also showed that such modelling choices can change the selected decisions and strategies, not just the numbers [55]. The tradition remains active and remains internal to the simulation route, as its most recent instalment applies a structured verification framework between two DES models on a harmonised floating reference case [63].

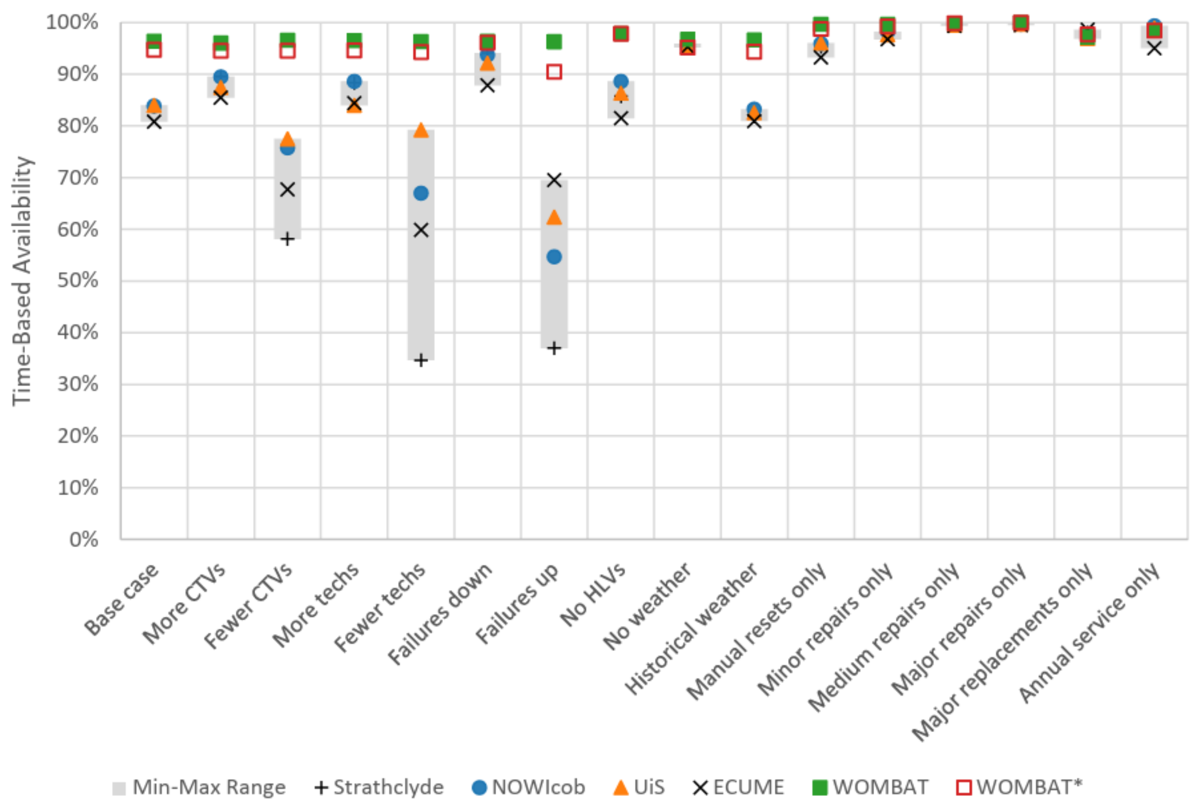

Figure 15: Mean availability estimates from the 2015 Dinwoodie et al. case comparison [8], with simulation results from WOMBAT. Extracted from [5].

The general simulation literature adds the vocabulary this thesis uses throughout. Verification asks whether the implementation solves its own equations correctly, validation asks whether the model answers for the system, and the two must be treated separately [107].

The validation programme of Chapter 4 is built on structuring comparisons as ladders, in which mechanisms are added one at a time so that each test isolates one cause, similarly to the process followed by Huang et al. [63]. This section therefore records a gap, as the benchmarking practice has been applied between simulators but never across the simulation–analytical divide at mechanism level under aligned conventions.

2.9 Synthesis and research gap

2.9.1 Synthesis of the reviewed literature

The review can be reduced to a trade-off, sketched in Figure 16. Simulation sits at high fidelity and high cost. It captures every mechanism that drives offshore downtime, namely non-exponential durations, logistics, weather, and interrupted repairs, although each converged answer costs many runs, which makes it inefficient for many-query studies. The analytical models sit at low cost and reduced fidelity. They answer much quicker, but the published offshore wind models are exponential, steady-state, and contiguous window-based. Between the two sits unoccupied territory, namely a Markov-based model that keeps analytical costs while recovering the mechanisms that its peers lack. The semi-Markov theory of Subsection 2.5.1 indicates towards this territory being reachable in principle, yet nobody has built and validated a model that occupies it.

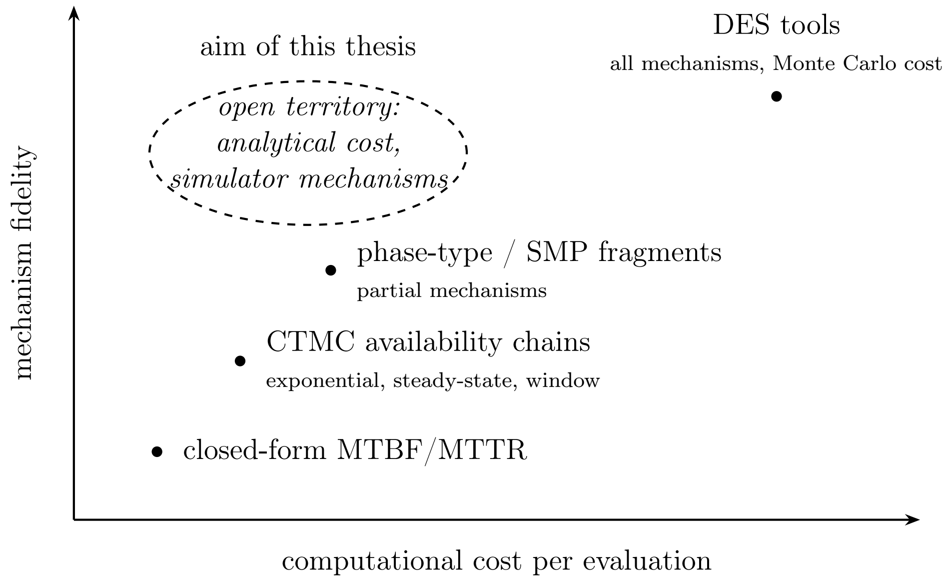

Figure 16: Positioning of the reviewed model families on the cost–fidelity plane.

2.9.2 The capability matrix

The four thesis rows of the matrix require a brief introduction before the comparison, since Chapter 3 builds them in full but the matrix uses them now. They arise from the two-by-two design previewed in Section 1.3: a Markovian description of the availability cycle solved in steady state (SS-C) and marched in time (TM-C), and a semi-Markov description solved in steady state (SS-S) and marched in time (TM-S). Two Monte Carlo twins (MC-C, MC-S) sample the same model content for cross-verification. Moving between rows changes one named assumption at a time, which is what makes the filled and empty cells attributable to specific modelling choices.

Table 8: Capability matrix comparing offshore wind O&M availability models in the literature with the models developed in this thesis. ✓ = capability present, ✗ = absent, (✓) = partial (see notes). "Transient" and "Steady state" refer to analytical solutions, simulation tools produce time-resolved sample paths by construction but not analytical solutions. Computational cost is per model evaluation: *high* = minutes–days (Monte Carlo over a lifetime), *low* = sub-second–seconds (linear algebra / quadrature).

Model / reference	Paradigm	Transient (analyt.)	Steady state	Non-exp. durations	Interrupt. repairs	Weather / access.	Validated vs. DES	Comp. cost
Simulation-based models								
WOMBAT [5]	DES	✗	✗	✓	✓	✓	✓	high
StrathOW-OM [57]	MC time-domain	✗	✗	✓	✓	✓	✓	high
NOWIcob [58]	DES/MC	✗	✗	✓	✓	✓	✓	high
UiS model [108]	ABM/DES	✗	✗	✓	✓	✓	✓	high
Dalgic et al. [54]	MC time-domain	✗	✗	✓	✓	✓	✗	high
Petri nets [60, 59]	SPN (MC-solved)	✗	✗	✓	✓	✓	✗	high
Analytical models								
Feuchtwang & Infield [6]	Closed-form prob.	✗	✓	(✓)[a]	✗	✓	✗	low
Besnard et al. [71]	Markov/queueing	✗	✓	✗	✗	(✓)[b]	✗	low
Huang et al. [72]	CTMC	✗	✓	✗	✗	✓	✗	low
Collection-system [94, 73, 74]	Markov/UGF	✗	✓	✗	✗	(✓)[b]	✗	low
Centeno-Telleria et al. [46]	CTMC	✗	✓	✗	✗	✓	(✓)[c]	low
Adjacent semi-Markov applications in wind energy								
Wind-speed SMC [86]	Semi-Markov	✓	✓	✓	n/a[d]	n/a[d]	✗	low
Weather-aware MDP [65]	MDP	✗	✓	✗	✗	✓	✗	low
SMDP planning [87]	Semi-MDP	✗	✓	✓	✗	✓	✗	low
This thesis								
Time-marching CTMC (Subsection 3.4.1)	CTMC	✓	–	✗	✗	✓	✓	low
Steady-state CTMC (Section 3.4)	CTMC	–	✓	✗	✗	✓	✓	low
Time-marching SMP (Section 3.8)	Semi-Markov	✓	–	✓	(✓)[e]	✓	✓	low
Steady-state SMP (Section 3.7)	Semi-Markov	–	✓	✓	(✓)[e]	✓	✓	low

[a] Weibull persistence of weather windows, but exponential failure and repair completion. [b] Weather appears through aggregate accessibility factors rather than an explicit weather process. [c] Compared against two simulation-based case studies reproduced from the literature, discrepancies in availability and cost without convention alignment or mechanism isolation (Subsection 2.4.2). [d] Models the wind resource, not the maintenance process. [e] Marginal (duration-distribution) sense of interruptibility: the SMP kernel carries the interruption-inclusive completion-time distributions derived from the reference model. The measured cost of this marginal approximation is quantified in Chapter 5.

Table 8 condenses the review into a single comparison. Each row is a model or model group from this chapter, plus the four analytical models developed in this thesis. The columns record what each model can do: whether it gives time-resolved and long-run answers analytically, whether durations may be non-exponential, whether repairs may be interrupted, whether weather accessibility is built in, whether the model has been validated against a DES reference under controlled conditions, and the rough cost of one evaluation. Marks in parentheses mean partial capability, and the table notes explain each.

The matrix makes the gap visible at a glance. The simulation rows are full in every capability column but expensive in the cost column, while the analytical rows invert the pattern. Three columns, namely transient analytical solution, non-exponential durations, and interruptible repairs, are empty for every analytical offshore wind model in the literature. Furthermore, no existing row combines analytical cost with any of those three capabilities while also being validated against a DES reference.

The matrix also dictates the methodological choices of the thesis rather than merely recording the gap. The three empty analytical columns map one-to-one onto the design just introduced. The semi-Markov axis of the two-by-two exists to fill the non-exponential and interruption columns, the time-marching axis exists to fill the transient column, and the Markovian half is retained so that the price of memorylessness can be measured rather than assumed. The validation column dictates the rest of the thesis, since filling it honestly requires the aligned, mechanism-accreting campaign of Chapter 4, and Section 2.12 records which model closes which element of the gap.

2.9.3 Statement of the research gap

The research gap this thesis tackles has five elements:

1. No analytical offshore wind availability model handles non-exponential durations. The published chains are exponential throughout, although mobilisation is near-deterministic, failures age, and interrupted repairs are multi-modal (Subsection 2.4.2).

2. No analytical offshore wind availability model represents interruptible repairs. Pause-and-resume execution exists in every simulator but is absent from analytical formulations (Subsection 2.6.2).

3. Time-resolved analytical solutions are unused. The analytical literature takes only steady-state answers, although the semi-Markov machinery for transients exists and matters for horizon-limited questions and degradation (Subsection 2.5.1).

4. Ageing under realistic repair is not solved analytically. When only the failed part is renewed and the others keep their operating age, the standard closed forms do not apply, and the wind literature avoids the regime rather than solving it (Subsection 2.4.2 and Subsection 2.5.2).

5. No analytical model has been validated against a modern DES reference at mechanism level under aligned inputs and conventions. The benchmarking practice of the field (Section 2.8) has never been applied across the simulation–analytical divide in that form, and the closest analytical antecedent was compared only against endpoint results reproduced from published cases (Subsection 2.4.2).

2.10 Research questions

Four research questions follow directly from the gap statement above. Chapter 6 answers each verbatim, and Table 9 at the close of Section 2.11 maps each question to the gap elements it derives from, the objectives that answer it, and the chapters that produce the evidence.

RQ1 (fidelity) Can an analytical availability model, built on strictly aligned inputs, reproduce the availability and cost estimates of a certified DES reference within the statistical resolution of the reference itself?

RQ2 (exponential adequacy boundary) For which output parameters and operating regimes does the memoryless simplification remain adequate, and where does the semi-Markov description become necessary?

RQ3 (interruptibility) How large is the pause-and-resume effect of weather-interrupted repairs, and does a semi-Markov model carrying interruption-shaped durations capture it?

RQ4 (computational reach and model applicability) What can the analytical route resolve and afford that the simulation route cannot, and what engineering studies does that reach enable?

2.11 Research aim and objectives

The aim of this thesis is to close this gap: to develop a hierarchy of analytical availability models for offshore wind O&M, and to establish, by validating against a certified discrete-event simulator as reference, which output quantities each model in the hierarchy can be trusted to estimate, at what accuracy, and at what computational cost. Five objectives structure the work required to answer the research questions:

1. **O1:** Develop the analytical model family: steady-state and transient solvers for both the Markovian and the semi-Markov description (SS-C, SS-S, TM-C, TM-S), plus simulator twins (MC-C, MC-S) that sample the exact processes and state transition mechanics the solvers integrate. All six consume one shared input pipeline. It is the foundation on which all four research questions rest.

2. **O2:** Align and certify the DES reference: a patched WOMBAT build with a decoded event protocol, verified input parity, and a pooled 1,000-seed benchmark with quantified convergence. It produces the reference against which RQ1 is judged.

3. **O3:** Execute a pre-registered validation campaign across the simulation–analytical divide, structured as a reference-build ladder and a mechanism-accretion ladder with attributable outputs and deltas. It produces the comparison that answers RQ1.

4. **O4:** Quantify the mechanisms that the simplifications of the analytical literature discard, namely interruptibility, non-exponential durations, and ageing, in steady state and in transient. It answers RQ2 and RQ3.

5. **O5:** Quantify what the analytical route can resolve and afford: runtimes, resolution relative to the DES noise floor, and the study types the speed unlocks. It answers RQ4.

Table 9: Derivation map from the gap statement to the research questions and the objectives that answer them. Each row traces one research question to the gap elements it derives from, the objectives that produce its answer, and the chapters where the evidence lands, with Chapter 6 stating the verbatim answers.

Gap element(s)	Research question	Objectives	Primary evidence
5	RQ1 (fidelity)	O1, O2, O3	Chapter 4, Chapter 5
1, 3, 4	RQ2 (adequacy boundary)	O1, O4	Chapter 5
2	RQ3 (interruptibility)	O1, O4	Chapter 5
cost column of Table 8	RQ4 (computational reach)	O1, O5	Chapter 5

2.12 From gap to models

The chapter closes by connecting its findings to the constructive work ahead. Table 10 states, for each model of the family, its purpose in the design, the key assumption that defines it, the gap elements it exists to close, and where it is built and validated. This table is the contract the next chapters execute. Chapter 3 builds the four models and their twins from one shared system definition, and Chapter 4 and Chapter 5 then test them against the reference in the order that the validation column of Table 8 demands.

Table 10: The four analytical models as answers to the research gap. Each row states the purpose of the model within the two-by-two design, the assumption that defines it, the gap elements of Subsection 2.9.3 it addresses, and where it is built and validated.

Model	Purpose	Key assumption	Gap elements	Built in	Validated in
SS-C	Baseline of the literature class, prices the simplifications	Memoryless durations, long-run answer	5	Section 3.4	Chapter 4, Chapter 5
TM-C	Adds time resolution to the Markovian description	Memoryless durations, transient solved	3, 5	Subsection 3.4.1	Chapter 5
SS-S	Drops memorylessness, carries interruption-shaped duration means	General durations through the empirical kernel	1, 2, 5	Section 3.7	Chapter 4, Chapter 5
TM-S	Full description: general durations, transients, ageing arms	General durations, Markov-renewal transient	All 5	Section 3.8	Chapter 5

3 Methodology and Model Development

3.1 Introduction and framework overview

Table 10 closed the literature review with a proposal, four models answering five gap elements, and this chapter executes its constructive half. The chapter builds the four analytical models the thesis runs on: a continuous-time Markov chain (CTMC) and a semi-Markov process (SMP) description of the same system, each solved in a steady-state and a time-marching version, plus a Monte Carlo sampler twin for each description. The four are a nested family. All consume the same inputs through one parameter pipeline, all share one state space and one metric set, and each differs from its neighbours at exactly one named place. The CTMC pair replaces every duration law by the exponential of equal mean, and each time-marching model adds only a transient solution operator to its steady-state partner. Therefore, when two models disagree, the cause is traceable to a named assumption. Figure 17 draws these relationships. The two columns are the two descriptions, the two rows are the two solution operators, every edge carries the single assumption that separates its endpoints, and the dashed twins sample the same model content that the solvers integrate. Table 11 states the same structure cell by cell, and the reader is encouraged to return to both throughout the chapter.

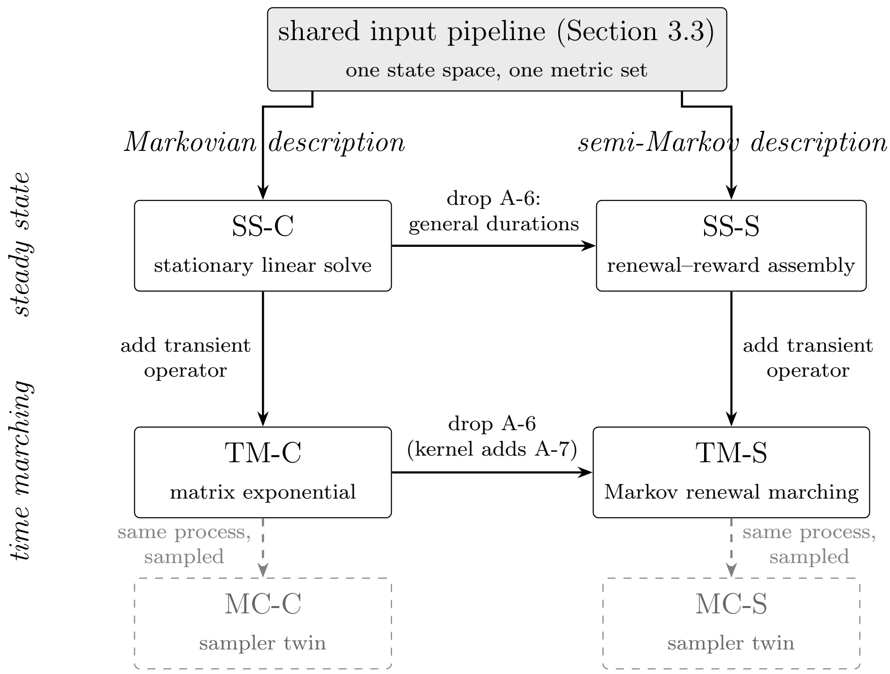

Figure 17: The model family and its single-assumption relationships. Columns are the two descriptions of the system, rows are the two solution operators, and every edge names the one assumption that separates its endpoints. The dashed samplers draw realisations of exactly the process that the solvers of their column integrate. One pipeline feeds all six, so no model ever sees different data.

Table 11: The model family at a glance. The state space Equation 2, the input pipeline, and the metric and cost conventions are shared by all models.

	SS-CTMC	TM-CTMC	SS-SMP	TM-SMP
Failure side	exponential superposition, rates λ_i	as SS-CTMC	exact operational-age kernel $(q_i, \tau_{\mathrm{UP}})$	stationary-share kernel Equation 17
Restoration side	exponential at mean Equation 8	as SS-CTMC	span mean Equation 6	full span law Equation 7
Solution operator	stationary linear solve	matrix exponential	renewal–reward Equation 16	Markov renewal marching
Outputs	long-run A, ω_i, costs	trajectories	long-run, ageing- and interruption-exact means	trajectories, span-faithful
Key assumptions	A-6	A-6	A-3; exact means	A-3, A-7

The chapter is structured to present the models in an increasing order of complexity. Section 3.2 defines the overall system, its state space, and the assumptions followed. Section 3.3 specifies the input pipeline and its two weather operators, which create every duration statistic the models consume. Section 3.4 presents the Markovian pair and its sampler twin, Section 3.5 argues why a semi-Markov pair is needed, and Section 3.6 to Section 3.8 build it along its own sampler twin. Section 3.9 discusses the expected computational cost of each model, and Section 3.10 provides a summary and closes the chapter. Throughout, each section opens with the main idea presented briefly, and states only the equations the solver actually evaluates. Additionally, the supporting mathematics, the CTMC machinery, the renewal and Markov renewal theory, the invariance theorem, the competing-risks derivations, and the exactness proof of the weather operator, lives in Appendix A, behind pointer sentences where each result is used.

3.2 System definition and modelling framework

3.2.1 System boundary and servicing concept

The system is a single offshore wind turbine, decomposed into subassemblies carrying corrective failure modes. Each failure mode is treated as the basic modelling unit and is defined by a failure-time distribution, a repair task specifying the active repair duration, vessel class and material requirements, and the corresponding logistics. Two vessel classes are considered: a crew-transfer vessel (CTV) for minor repairs, resident at port and thus with no mobilisation period, and a heavy-lift vessel (HLV) used for major replacements, chartered on demand with a long mobilisation period. Each maintenance episode includes the failure event, mobilisation where required, transit to the turbine, a weather-dependent crew transfer, repair work completed over one or more accessible weather windows, the final crew transfer and the return journey. All activities are evaluated against an hourly metocean time series and the operational limits of the assigned vessel. Scheduled maintenance, cable failures and non-failure curtailments lie outside the scope of the thesis, as it matches the corrective single-asset setting fixed in Chapter 1, and it is the setting in which the mechanisms under study, ageing and weather-interrupted repair, appear in their most direct form. Conclusions therefore transfer most directly to corrective maintenance strategies on comparable assets, and the composition rules of A-9 state how far beyond a single asset the framework reaches.

3.2.2 Hierarchy and state space

Concept. The state space records what the asset is doing (either up, waiting for mode i, or under restoration for mode i) and nothing else.

One pipeline feeds four solvers and a sampler into one metric set, so no model ever sees different data. For n failure modes, the shared state space per asset follows:

$$\mathcal{S} \;=\; \{\,\mathrm{UP}\,\} \,\cup\, \big\{\,\mathrm{W}_i,\ \mathrm{R}_i \;:\; i = 1,\dots,n\,\big\}, \qquad |\mathcal{S}| = 1 + 2n, \tag{2}$$

with the cycle UP $\to$ W_i $\to$ R_i $\to$ UP: a mode-i failure sends the asset into the wait state W_i (mobilisation plus transit leg), then the restoration state R_i (the weather-gated span banking transfers and repair), then back up.

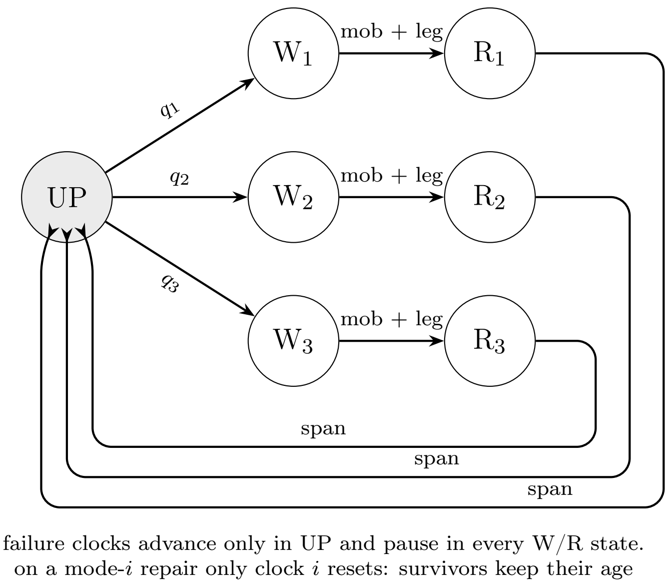

Figure 18: Common state-space topology of all four models, shown for $n = 3$ failure modes. The annotation states the operational-age partial-reset clock rules of Section 3.6.

Figure 18 shows the topology and the clock rules. Two composition rules extend it. Within an asset, the failure kernel is solved exactly and components are coupled by linking failure modes to a shared UP state, with all failure clocks pausing whenever the asset is down, as this is the regime the reference implements.

3.2.3 Assumptions

Table 12 collects the assumptions, their justification, and the models they bind, each probed by the campaign of Chapter 4.

Table 12: Modelling assumptions, their justification, and the models they bind.

ID	Assumption	Justification	Models
A-1	The finite state space Equation 2 summarises the operational condition	Smallest topology supporting the metric set and the reference's span decomposition	all
A-2	Failure behaviour follows the operational-age, partial-reset competing-Weibull regime: clocks advance only while operating, only the failed mode renews	Standard virtual-age regime	all
A-3	Weather-gated sojourns are summarised by their marginal law under stationary start-hour weighting, uncoupled from the concurrent weather state	Marginal-interruptibility formulation, exact for means by renewal–reward	TM/SS-S
A-4	Sojourn allocation: W_i holds mobilisation and the transit leg, R_i holds the weather-stretched span banking transfers and repair	Reproduces the reference's accounting split	all
A-5	The trailing crew transfer is weather-gated symmetrically with the leading one	Physically sequenced, marginally conservative against the reference, which appends it flat	TM/SS-S
A-6	Exponential moment match: the CTMC pair replaces every sojourn law by the exponential of equal mean	Defines the shape-isolation comparison, exact at $\beta = 1$	TM/SS-C
A-7	The transient semi-Markov kernel uses stationary mode shares and conditional first-passage laws	Keeps the transient model on the physical state space, exact for $\beta = 1$ or a single mode	TM-S

continues on next page

Table 12 (continued)

ID	Assumption	Justification	Models
A-8	Parameters are time-homogeneous, seasonality enters through the full-record weather operators, not time-varying rates	Scoping, long-run metrics unaffected, sub-annual transients approximate	all
A-9	Farm composition: identical, independent assets under dedicated servicing, component coupling by operational-rate superposition with shared downtime clock-pausing	Standard series and superposition structure, boundary exposed by design	all
A-10	Billing convention: all downtime billed at the vessel day rate except mobilisation, which carries a flat fee, materials flat per task	Matches the reference's verified billing identity (Section 4.2)	all
A-11	Energy accounting: net production equals gross production scaled by time-based availability	First-order, ignores the weather–production correlation	all

The structural group, A-1, A-8, and A-9, fixes what the framework can talk about. The state space is the smallest that supports the metric set, so relaxing A-1 towards finer states adds cost without adding a reportable quantity. Relaxing A-8 towards time-varying rates would sharpen sub-annual transients while leaving every long-run metric unchanged, since seasonality already enters through the full-record weather operators. Relaxing A-9 towards shared vessels or dependent assets introduces queueing interactions that lie outside the analytical scope of this thesis.

The mechanism group, A-2 to A-5, encodes how failures and weather actually behave. A-2 is the regime the reference implements and the standard virtual-age description of imperfect repair, and relaxing it towards full renewal is the known shortcut that the kernel of Section 3.6 exists to avoid. A-3 is the deepest modelling choice in the thesis: weather-gated durations are carried as marginal laws, exact in their means by the renewal–reward argument of Section A.5, while the coupling between the state of the weather and the instant of a transition is discarded. Relaxing it would require a weather-state-augmented kernel at a large increase in state-space size, and Chapter 5 measures what the marginal treatment leaves behind. A-4 mirrors the accounting split of the reference so that comparisons are like for like. A-5 keeps the one deliberate asymmetry, discussed with the allocation equations below.

A-6 and A-7 are the comparison-defining group, and A-6 is what makes the Markovian pair the priced baseline, since relaxing it is precisely the step from the C family to the S family. A-7 keeps the transient semi-Markov solver on the physical state space by averaging the age dependence of the UP row, and relaxing it would mean marching on an age-augmented space at a cost that defeats the purpose of an analytical model. Its error is localised to the early failure flow of an ageing fleet started fresh, and the falsifiers of Section 3.8 bound it.

A-10 and A-11 align accounting conventions. A-10 reproduces the billing identity verified inside the reference (Section 4.2), so cost comparisons carry no convention residue. A-11 is first-order and its bias is one-sided and common to all four models, so it cancels from every model-to-model comparison. Relaxing it means resolving production jointly with the state trajectory over the weather record, which the semi-Markov family supports through Equation 10, and Chapter 5 prices the difference as the gap between time-based and production-based availability.

3.2.4 Metrics

Every solver reports the same set of KPIs: availability, expected downtime, per-mode task frequency ω_i, state occupancies, and cost by Equation 9, with energy per A-11. The time-marching models add trajectories $A(t)$, $D(t)$, $C(t)$, and the sampler twin adds distributions over replications. The shared metric set is what makes every later comparison a comparison of models rather than of reporting conventions.

3.3 Input data and the parameter pipeline

Concept. The pipeline turns three data sources into the duration and cost statistics every model consumes, and it does so once, so that all six models eat from the same table.

Figure 19 draws the workflow: the failure and task table and the vessel parameters fix what a repair needs, the metocean record fixes when the weather allows it, the accessibility mask and the two weather operators convert that record into duration statistics, and the outputs split by consumer, means for every solver, full distribution laws for the time-marching semi-Markov model.

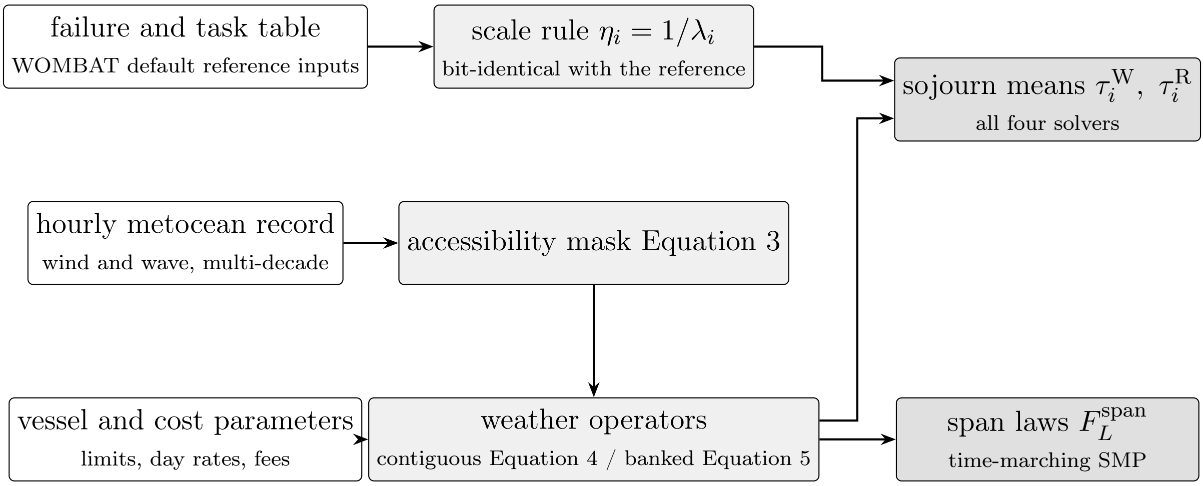

Figure 19: The input parameter pipeline. Three data sources enter, the accessibility mask and the two weather operators convert the metocean record into duration statistics, and the outputs split by consumer, with means feeding all four solvers, and the full span laws feeding the time-marching semi-Markov model. Billing and energy conventions (A-10, A-11) apply downstream of this diagram.

3.3.1 Failure and task data

Failure modes enter as one table, a row per (component, mode) pair: annual rate λ_i and Weibull shape β_i, active repair time H_i, vessel class, transfer time T_i, mobilisation duration and fee, transit leg ℓ_i, weather limits, day rate and materials cost. The values of the table are the default reference inputs of the WOMBAT library, adopted unchanged, which is the input-alignment decision recorded in Chapter 2 and audited in Chapter 4. Two conventions matter: The scale rule $\eta_i = 1/\lambda_i$ is that of the reference tool, not mean-preserving for $\beta_i \neq 1$ (the implied MTTF is $\eta_i\Gamma(1+1/\beta_i)$), but it keeps the inputs bit-identical with the reference. And H_i is active work content, with weather and logistics added by the operators below, never twice.

3.3.2 Metocean record and accessibility

The environment is an hourly wind and wave record spanning multiple decades, shared bit-identically with the reference. Per vessel class it induces a binary accessibility mask,

$$g(s) \;=\; \mathbf{1}\big\{v(s) \leq v^{\mathrm{lim}} \;\wedge\; h_s(s) \leq h_s^{\mathrm{lim}}\big\}, \qquad s = 1, \ldots, S_{\mathrm{rec}}, \tag{3}$$

tiled cyclically past the end of the record. Every piece of weather information the models use is a functional of Equation 3, and no synthetic weather is generated, so the models and the reference never disagree about the environment in any realisation.

3.3.3 The weather operators

Concept. Both operators answer the same plain question: starting at hour s, how long until a job of work content L is finished? They differ only in the reading of the weather.

The contiguous operator waits until one calm spell long enough for the whole job appears, and then works without interruption, which is the classical assumption of the access literature. The banked operator starts working in the first accessible hour, pauses whenever the weather closes, and keeps its progress, which is what vessels and the reference simulator actually do. Figure 20 shows both readings acting on the same mask segment before the equations state them.

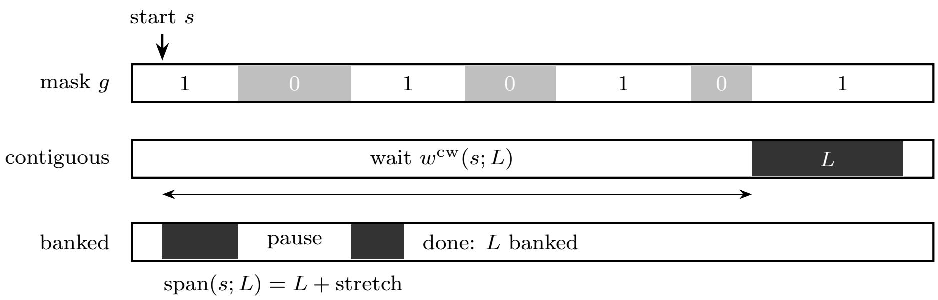

Figure 20: The two weather operators on one mask segment (accessible = 1, blocked = 0, in grey). The contiguous reading waits for the first uninterrupted window of length L and then works unbroken. The banked reading starts immediately, pauses through blocked hours with progress kept, and completes once L accessible hours are banked, so its span decomposes as the work content plus a stretch.

The core of the pipeline converts the mask into duration statistics through two operators. The contiguous-window operator is the classical reading of the access literature [6, 52]: a task of work content L waits for one uninterrupted accessible window,

$$w^{\mathrm{cw}}(s;L) \;=\; \min\Big\{\, u \geq 0 \;:\; g(s+u) = \cdots = g(s+u+\lceil L\rceil - 1) = 1 \,\Big\}, \tag{4}$$

then runs unbroken. It is kept as a configurable variant to measure the assumption of the analytical precedents.

The interruptible banked-stretch operator is the actual mechanism from the reference: work proceeds through accessible hours and pauses through inaccessible ones, keeping progress. The completed span from start hour s is the calendar time to bank L accessible hours,

$$\mathrm{span}(s;L) \;=\; \min\Big\{\, u \geq 0 \;:\; \textstyle\int_0^u g(s+x)\,\mathrm{d}x \;\geq\; L \,\Big\}, \tag{5}$$

which decomposes as L plus a stretch, the initial wait plus the interleaved pauses. Its start-hour mean under a weighting $\rho(s)$,

$$\bar{D}(L) \;=\; \sum_{s=1}^{S_{\text{rec}}} \rho(s)\,\text{span}(s;L), \tag{6}$$

is, for uniform ρ, the exact long-run mean restoration span of a stationary failure stream over this climate, a renewal–reward statement with no stochastic error, since the walk is deterministic given the record. The proof is given in Section A.5. Beyond the mean, the start-hour ensemble is itself the marginal, interruption-inclusive completion-time law

$$F_L^{\text{span}}(t) \;=\; \sum_{s} \rho(s)\,\mathbf{1}\big\{\text{span}(s;L) \le t\big\}, \tag{7}$$

This is the object A-3 names, consumed whole by the time-marching SMP.

Sojourn allocation follows. The wait state carries the weather-independent logistics, $\tau_i^{\text{W}} = \tau_i^{\text{mob}} + \ell_i$ (the contiguous variant adds $\bar{w}^{\text{cw}}$). The restoration state banks both transfers with the repair, $L_i = H_i + 2T_i$, giving

$$\tau_i^{\text{R}} \;=\; \begin{cases} \bar{D}\big(H_i + 2T_i\big) & \text{interruptible variant,} \\ H_i + 2T_i & \text{contiguous variant,} \end{cases} \tag{8}$$

with the full law Equation 7 at the same L_i retained for the transient model. Assumption A-5 records the one deliberate asymmetry: here the trailing transfer is weather-gated like the leading one, where the reference appends it flat, a small discrepancy on this side kept as a deliberate modelling choice. The episode figure of Chapter 2 draws the operational accounting of the reference, in which the turbine returns to service when the repair completes and the trailing transfer falls outside downtime, whereas the model banks that transfer inside R_i, the conservative side of the same convention, and Chapter 4 prices the difference.

3.3.4 Energy and cost

Gross energy is computed once from the power curve on the wind record, and net production is gross times availability (A-11). The known bias, storms that block access are also high-production hours, is one-sided and common to all four models, so it cancels from every model-to-model comparison. Costs follow the verified billing convention of the reference (A-10): every downtime hour billed at the vessel day rate except the mobilisation period, which carries only its flat fee, and materials flat per task. The annual expected costs per asset are

$$C^{\text{eq}} \;=\; 8766 \sum_{i=1}^{n} \omega_i \left[\frac{c_i^{\text{day}}}{24} \left(\tau_i^{\text{W}} - \tau_i^{\text{mob}} + \tau_i^{\text{R}} \right) + c_i^{\text{mob}} \right], \qquad C^{\text{mat}} \;=\; 8766 \sum_{i=1}^{n} \omega_i\, c_i^{\text{mat}}, \tag{9}$$

with ω_i the per-mode task frequency each solver delivers, evaluated cumulatively as $C(t)$ by the time-marching models.

Beyond the first-order convention of A-11, the pipeline supports the production-based reading of availability directly, because the span walks are calendar-anchored. Let $e(h)$ denote the potential power of record hour h, computed once from the power curve on the wind record, with long-run mean $\bar{e}$. For a mode-i episode whose restoration starts at record hour s, the down window is known exactly: the wait interval of length τ_i^{W} anchored immediately before s, followed by the restoration interval from s to $s+\mathrm{span}(s; L_i)$. Summing $e(h)$ over that window gives its potential-energy content $E_i^{\mathrm{down}}(s)$, and the start-weighted mean $\bar{E}_i^{\mathrm{down}} = \sum_s \rho(s)\, E_i^{\mathrm{down}}(s)$ is the same renewal–reward evaluation as Equation 6, with potential energy replacing calendar hours as the reward. The production-based availability of then follows as

$$A_E = 1 - \frac{\sum_{i=1}^{n} \omega_i\, \bar{E}_i^{\mathrm{down}}}{\bar{e}}, \tag{10}$$

with ω_i the task frequencies of the solving model.

This quantity is available to the semi-Markov family only. The Markovian description carries no calendar anchoring, so its sojourns are memoryless and therefore exchangeable over the record phase, the expected potential power during its down time equals the record mean, and the energy weighting collapses onto the time weighting, so A_E cannot be distinguished from A within the C family. The gap $A - A_E$ is consequently itself a shape-sensitive output, produced by the S family alongside A, and Chapter 5 reports it. A-11 remains the common first-order convention that every cross-model comparison shares, and Equation 10 is the refinement the S family adds on top of it.

Key outcome. Every duration a solver uses is a functional of one mask on one shared record, every mean is exact by construction, and the only quantities that differ between models downstream are the ones each model chooses to keep, means or full laws.

3.4 The CTMC pair

Concept. The Markovian pair forgets everything except which state the asset occupies, so every duration becomes exponential and the whole model reduces to linear algebra.

This is the priced baseline of the family: it keeps the correct interruption-inclusive means from the pipeline and discards only shape, which is exactly the simplification the analytical literature makes.

The Markovian pair replaces every sojourn law by the exponential of equal mean (A-6). Under memorylessness the reset and pausing rules stop mattering, so the failure rates are the tabulated λ_i, the wait rates $\alpha_i = 1/\tau_i^{\mathrm{W}}$, and the restoration rates $\mu_i = 1/\tau_i^{\mathrm{R}}$: the exponential keeps the correct interruption-inclusive mean and discards only shape. On the ordering $(\mathrm{UP}, \mathrm{W}_1, \mathrm{R}_1, \dots)$ the generator is

$$q_{\mathrm{UP},\mathrm{W}_i} = \lambda_i^h, \qquad q_{\mathrm{W}_i,\mathrm{R}_i} = \alpha_i, \qquad q_{\mathrm{R}_i,\mathrm{UP}} = \mu_i, \tag{11}$$

all other off-diagonals zero, diagonals fixed by row sums, and λ_i^h in per-hour units. A joint-component variant on the product space exists to price what the composition rules miss. The stationary solution is closed-form on this topology:

$$A = \pi_{\mathrm{UP}} = \frac{\tau_{\mathrm{UP}}}{\tau_{\mathrm{UP}} + \sum_{i=1}^{n} q_i \left(\tau_i^{\mathrm{W}} + \tau_i^{\mathrm{R}}\right)}, \qquad \tau_{\mathrm{UP}} = \frac{1}{\sum_i \lambda_i^h}, \quad q_i = \frac{\lambda_i^h}{\sum_j \lambda_j^h}, \tag{12}$$

mean up-time over mean cycle, with task frequencies $\omega_i = q_i / \left(\tau_{\mathrm{UP}} + \sum_j q_j (\tau_j^{\mathrm{W}} + \tau_j^{\mathrm{R}})\right)$. The derivation from the balance equations is given in Section A.1. The solve is one small linear system, effectively instantaneous.

3.4.1 The time-marching CTMC

The transient partner keeps the generator and changes only the operator: $\boldsymbol{p}(t) = \boldsymbol{p}(0)\, e^{\mathbf{Q}t}$ from an all-UP start, evaluated by scaling-and-squaring and cross-checked by uniformisation, both summarised in Section A.1. Availability is $A(t) = p_{\mathrm{UP}}(t)$, averages and downtime follow by quadrature, per-mode completion intensity is $p_{\mathrm{R}_i}(t)\, \mu_i$, and cost applies Equation 9 along the trajectory. Runtime is a handful of matrix operations per output time, the computational floor of the hierarchy.

3.4.2 The Markovian sampler twin

Alongside the deterministic solvers runs an exact stochastic sampler of the same chain, the first dashed twin of Figure 17, in the manner of [109]: an exponential sojourn at the current exit rate, then a categorical draw of the next state. By construction it describes the identical process, so agreement verifies both implementations with no reference in the loop, and its replication ensemble supplies what deterministic solvers cannot, like inter-annual spread and exceedance levels such as P90.

Key outcome. The complete Markovian half of the family: a closed-form steady state, a matrix-exponential transient, and a sampler of the identical process, all consuming the pipeline means unchanged. Whatever these three get wrong later must be attributable to the exponential shape assumption, not to inputs, states, or conventions.

3.5 Why a semi-Markov pair

The exponential assumption fails the physics twice. On the failure side, mechanical modes age ($\beta \neq 1$), and under partial reset the violation is structural: survivors keep their ages, up-times become dependent, no constant-rate superposition matches the aged fleet,

and the full-reset shortcut misestimates task frequency for $\beta > 1$. On the restoration side, the span law Equation 7 is multi-modal, and the moment match preserves its mean, hence long-run availability, but nothing that sees shape: transients, per-event span distributions, tail-driven dispersion.
Hence the following hypothesis pair.

- *H1*: mean-determined quantities are reproduced by the mean-matched CTMC whenever the failure side is near-exponential.
- *H2*: shape-sensitive quantities require the semi-Markov pair, with CTMC error growing in $|\beta - 1|$ and weather severity.

The design brief follows: keep state space, pipeline, metrics and costs, replace the failure superposition by the exact operational-age kernel and the restoration sojourns by the full span law, within the marginal-interruptibility scope of A-3.

3.6 The semi-Markov failure kernel

Concept. Each failure mode carries an odometer that runs only while the turbine runs, a repair resets one odometer and leaves the others where they stand, and the kernel computes exactly what this bookkeeping implies for which mode fails next and when.

Figure 21 shows the process: the age vector determines the competing-risks law of the next failure, the renewal update resets the failed clock, the loop is an embedded Markov chain on the age space, and solving that chain to stationarity yields the three quantities the downstream solvers consume.

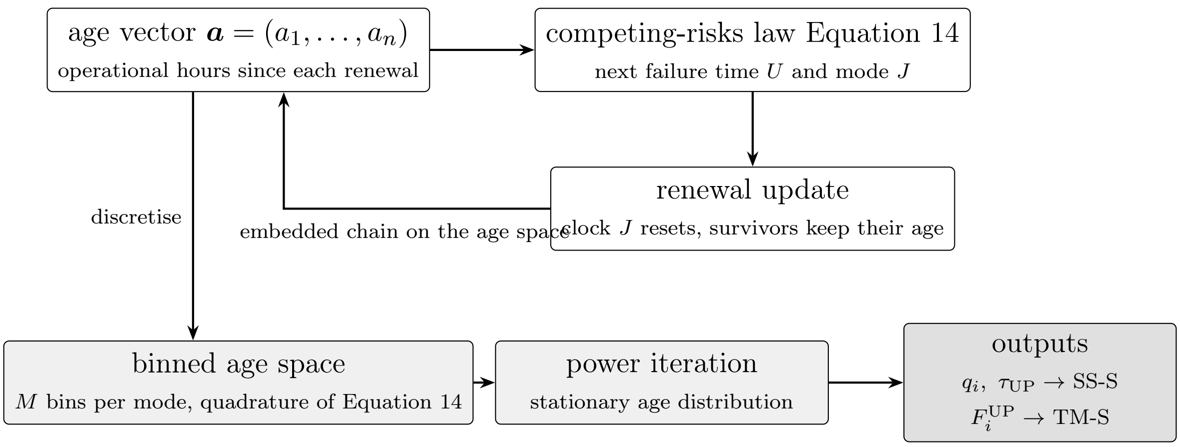

Figure 21: The failure kernel as a computation.

Both semi-Markov models rest on the exact treatment of n competing Weibull modes under the operational-age, partial-reset regime. Each mode carries an operational age a_i, the cumulative operating time since its last renewal. While the asset is up all ages advance together, with Weibull hazard

$$h_i(a) \;=\; \frac{\beta_i}{\eta_i}\left(\frac{a}{\eta_i}\right)^{\beta_i-1}, \qquad R_i(a) \;=\; \exp\!\big[-(a/\eta_i)^{\beta_i}\big]. \tag{13}$$

When any mode fires, all clocks pause, the failed mode resets to zero, and the survivors keep their ages: the partial-repair regime of [110], a virtual-age special case [102, 101], and the regime the reference implements. Figure 22 sketched the resulting clock behaviour. Consequently the post-repair age vector depends on the whole history, and no single-cycle closed form is exact. At $\beta_i \equiv 1$ everything collapses to the exponential superposition, the nesting point. Conditional on the age vector $\boldsymbol{a}$, the up-time U and failing mode J follow the standard competing-risks law

$$\Pr\big(U > u \mid \boldsymbol{a}\big) \;=\; \prod_{j=1}^{n} \frac{R_j(a_j+u)}{R_j(a_j)}, \qquad \Pr\big(J = i,\, U \in \mathrm{d}u \mid \boldsymbol{a}\big) \;=\; h_i(a_i{+}u)\,\Pr\big(U > u \mid \boldsymbol{a}\big)\,\mathrm{d}u, \tag{14}$$

after which ages update and the next up period begins: an embedded Markov chain on the age-vector space. The derivation of Equation 14 from the hazards, and the positioning of the regime within the virtual-age families, are given in Section A.4.

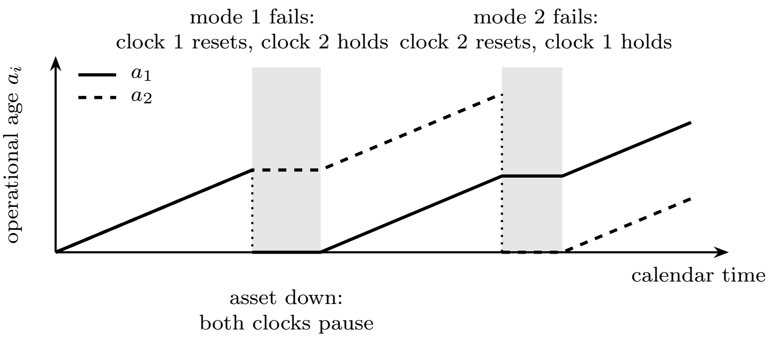

Figure 22: Sketch of the operational-age clock rules with two modes.

The kernel solves this chain to stationarity: the age space is binned (M bins per mode), Equation 14 is integrated by quadrature onto the binned space, and the stationary distribution is found by power iteration. Averaging over it delivers the three outputs downstream solvers consume: the stationary mode shares $q_i = \Pr(J = i)$, the mean operational up-time $\tau_{\mathrm{UP}} = \mathbb{E}[U]$, and the conditional first-passage laws

$$F_i^{\mathrm{UP}}(u) \;=\; \Pr\big(U \le u \mid J = i\big), \tag{15}$$

retained for the time-marching model. The construction is exact in the grid limit, with discretisation as its only error, the convergence argument is given in Section A.4, and it has been cross-checked against the independent operational-age sampler. For non-homogeneous

Weibull shapes in the failure modes, components are instead coupled at asset level by operational-rate superposition (A-9), sketched in Figure 23: each kernel yields rates $r_i = q_i/\tau_{\mathrm{UP}}$ per operating hour, the natural invariant of a clocks-pause regime, and the asset totals follow from $R_{\mathrm{tot}} = \sum_i r_i$. Within a component the competing-risk ageing is exact, and across components the age interaction is summarised by rates, exact when the interaction is memoryless or only one component ages, and it is the one approximation the failure side carries.

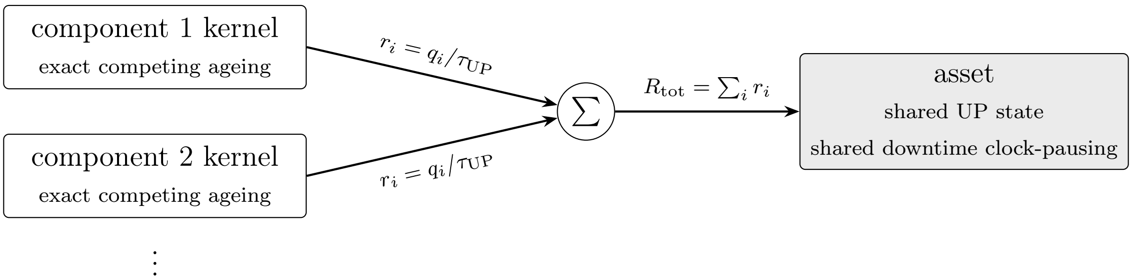

Figure 23: Component coupling under A-9. Note that within a component the ageing interaction is exact, and across components it is carried by rates.

Key outcome. The failure side of the semi-Markov family, delivered as three quantities: exact stationary mode shares, an exact mean operational up-time, and the per-mode first-passage laws. The first two feed the steady-state model, the third feeds the transient one, and the exponential case is recovered identically at $\beta_i \equiv 1$.

3.7 The steady-state SMP model

Concept. Long-run availability is bookkeeping how long an average cycle lasts and how much of it is spent up.

The kernel supplies the exact failure-side ingredients, the pipeline supplies the exact weather-side means, and the renewal–reward assembly divides one by the other.

The principal artefact of the thesis solves in two stages: the kernel above, then a renewal–reward assembly. The kernel delivers $(q_i, \tau_{\mathrm{UP}})$, the pipeline delivers $(\tau_i^{\mathrm{W}}, \tau_i^{\mathrm{R}})$, and with mean cycle length $D = \tau_{\mathrm{UP}} + \sum_i q_i(\tau_i^{\mathrm{W}} + \tau_i^{\mathrm{R}})$ the limiting occupancies read

$$P_{\mathrm{UP}} = \frac{\tau_{\mathrm{UP}}}{D}, \qquad P_{\mathrm{W}_i} = \frac{q_i\,\tau_i^{\mathrm{W}}}{D}, \qquad P_{\mathrm{R}_i} = \frac{q_i\,\tau_i^{\mathrm{R}}}{D}, \qquad \omega_i = \frac{q_i}{D}, \tag{16}$$

with $A = P_{\mathrm{UP}}$. The structure is identical to Equation 12, only the ingredients differ: ageing enters through the exact $(q_i, \tau_{\mathrm{UP}})$, interruptibility through τ_i^{R}. The whole solve is one kernel iteration plus arithmetic, so computation still remains at milliseconds.

As a consequence of Equation 16, the invariance theorem appears. The steady-state occupancies of a semi-Markov process depend on its holding-time laws only through their means, so any two duration models with identical means produce identical long-run availability, occupancies, and task frequencies, whatever their shapes. This is the invariance property of the steady-state solution referred to in the review, its proof is given in Section A.3, and its consequences cut both ways: the steady state depends on the weather operator only through the mean Equation 6, so span shape influences only the time-marching extension, and no steady-state agreement can ever certify distributional structure, which is why the validation campaign of Chapter 4 needs transient and distributional probes at all.

Key outcome. The principal artefact of the family is a steady-state model that is exact in its means on both the failure and the weather side, at millisecond cost, together with the named theorem that fixes what such a model can and cannot certify.

3.8 The time-marching SMP model

Concept. The transient model asks the steady-state question at every instant: starting from a fresh asset, what is the probability of being up, waiting, or under restoration at time t?

Answering it requires distributions where the steady state needed only means, and the pipeline already carries them.

The transient extension keeps the state space and pipeline outputs and replaces the operator with the discretised Markov renewal equations. The kernel rows come straight from the pipeline: from UP,

$$Q_{\mathrm{UP},\mathrm{W}_i}(t) \;=\; q_i\, F_i^{\mathrm{UP}}(t), \tag{17}$$

from W_i a unit step at the deterministic logistics content τ_i^{W}, and from R_i the full marginal span law at the work content of the mode,

$$Q_{\mathrm{R}_i,\mathrm{UP}}(t) \;=\; F^{\mathrm{span}}_{L_i}(t), \qquad L_i = H_i + 2T_i, \tag{18}$$

inserted as the empirical ensemble CDF directly, with no parametric fit between the weather record and the transient solution.

One structural approximation enters at Equation 17, assumption A-7: the true UP row is age-dependent across cycles, and the model replaces it by its stationary average, keeping the transient solver on the physical $1 + 2n$ state space. It is exact at the nesting point and for a single mode, agrees with the steady-state solver in the long run by construction, and should disagree only at the early failure flow of an ageing multi-mode asset started fresh. With the kernel assembled, the transition matrix solves the Markov renewal system by forward marching on a step fixed by refinement study, with substochasticity asserted as a running invariant, from an all-UP fresh-clock start, consistent with the reference runs it is compared against. The Markov renewal equations and their discretisation are summarised in Section A.2. Metrics mirror the CTMC case, with cumulative expected tasks accumulated by the convolution increments at no extra cost. It is the most expensive solver of the family, but still orders of magnitude below a converged reference study. Within it, there are two built-in falsifiers: at $\beta \equiv 1$ it must reproduce the time-marching CTMC, and as $t \to \infty$ it must converge onto Equation 16.

Key outcome. The most complex member of the family is this transient solver carrying the exact interruption-shaped span laws and the first-passage laws of the kernel on the physical state space, with its one approximation named, localised, and equipped with falsifiers. Its sampler twin is given below, closing the model family with its last sampler member.

3.8.1 The semi-Markov sampler twin

The S family carries its own sampler, the second dashed twin of Figure 17. One replication runs the exact model content realisation by realisation. At the current age vector it draws the next up-time and failing mode from the competing-risks law Equation 14, advances every age by the drawn up-time, resets the failed clock, and keeps the survivors, the regime of A-2. The wait sojourn is the deterministic logistics content τ_i^{W}. The restoration span is then drawn from the start-hour ensemble of Equation 7, realised by walking the recorded mask from a sampled start hour, so the sampled spans follow exactly the law the solvers consume. Because the twin samples precisely the processes the two semi-Markov solvers integrate, agreement between them verifies the implementations with no reference in the loop, and it is the independent cross-check already invoked for the kernel in Section 3.6. Its replication ensemble additionally supplies the dispersion outputs of the family, inter-annual spread and exceedance levels such as P90. At $\beta_i \equiv 1$ it must agree in law with the Markovian twin, one more expression of the nesting.

3.9 Expected computational cost

Since computational reach is one of the stated motivations of the framework, the expected cost of each member is worth fixing before any validation result. The steady-state CTMC is a closed-form expression, arithmetic in the number of modes n. The time-marching CTMC evaluates a matrix exponential on the $1+2n$ state space, with cost cubic in the state count per evaluation and effectively constant in practice at this size. The semi-Markov kernel is the one precomputation of the family: quadrature and power iteration on the binned age space, with cost growing in the bin count M and the number of competing modes per component, executed once per configuration and reused by both semi-Markov solvers. The steady-state SMP then adds only arithmetic, so its marginal cost is that of the kernel. The time-marching SMP is the most expensive member, as the discretised Markov renewal marching scales with the number of time steps and the support of the span laws, yet it remains orders of magnitude below a converged Monte Carlo study of the reference, whose cost multiplies a full lifetime simulation by the replication count that the square-root convergence of Subsection 2.3.2 demands. The samplers scale with events times replications and are used for verification and dispersion, not for expectation estimates. The a priori ordering is therefore SS-C, then TM-C and SS-S, then TM-S, all analytical members sub-second on the reference configuration, and Chapter 5 reports the measured wall times and the resulting speedup against the reference.

3.10 Summary and handoff

The chapter delivered what the bridge of Chapter 2 promised. One system definition and one pipeline feed four analytical models and two sampler twins, the family differs by one named assumption per edge of Figure 17, and every mechanism the gap statement demanded is now carried somewhere in the family: general durations and ageing in the kernel, interruptibility in the span laws, transients in the two time-marching operators. The internal falsifiers, exponential nesting, transient-to-stationary convergence, and sampler agreement, make the family self-checking before any external comparison, the web drawn in Figure 24. However, self-consistency is not validity. Whether these models reproduce the reference, mechanism by mechanism and under aligned conventions, is an empirical question, and Chapter 4 sets up the reference, the alignment, and the campaign that answers it.

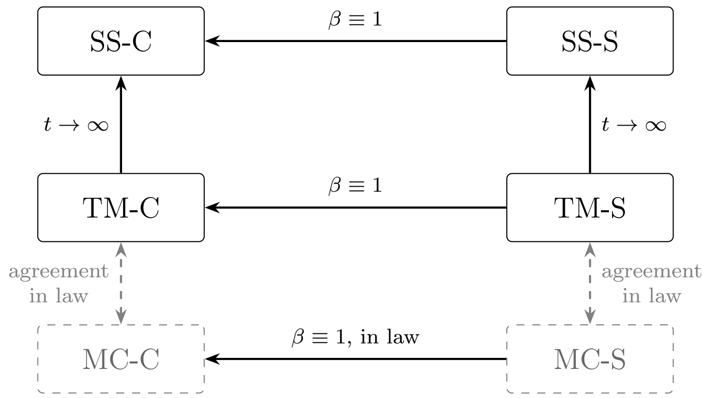

Figure 24: The internal consistency web of the family. Solid arrows are analytical falsifiers, as at $\beta \equiv 1$ every semi-Markov member must collapse onto its Markovian counterpart, and each time-marching solution must converge onto its steady state as $t \to \infty$. Dashed links are sampler agreements, which verify the implementations with no reference in the loop.

4 Model Validation

4.1 Introduction and scope

Chapter 3 established the model family, its assumptions, and its internal falsifiers. This chapter establishes that the models are correct, in two senses that are kept separate throughout. Verification asks whether each artefact solves its own equations: gates, identities, limits, and cross-checks between independently coded implementations of the same mathematics. Validation asks whether the model family reproduces the behaviour of an independent reference, the WOMBAT discrete-event simulator, under rigorously aligned inputs.

The chapter proceeds in six stages. Section 4.2 audits the reference, repairs its defects through a verified patch series, and freezes the build ladder and the aligned input baseline on which every comparison stands. Section 4.3 fixes the measurement instruments: the metric layers, their practical interpretation, and the pre-registered statistical discipline. Section 4.4 executes the mechanism-accretion ladder, which validates one mechanism per rung against sealed predictions. Section 4.5 verifies that the family solves its own mathematics, executing the internal consistency web of Figure 24. Section 4.6 presents the summit comparisons against the reference batches, and Section 4.7 measures the statistical resolution of the reference itself, which every z score in the chapter is expressed against. Section 4.8 closes with generalisability and the claims this programme licenses for Chapter 5. Table 13 maps the stages, and each results-bearing block below opens by naming the models under test and its purpose.

Table 13: Validation map. Each stage names the models under test, its purpose, and the basis on which it is accepted.

Stage	Models under test	Purpose	Acceptance basis
Reference audit and patches (Section 4.2)	the reference itself	establish per-event semantics, repair defects, separate repair from convention transplant	verified log signatures, exact
Alignment (Subsection 4.2.4)	reference and family jointly	one frozen input baseline, aligned clock, sequence, and billing conventions	input parity, ledger closure
Internal verification (Section 4.5)	all six family members	implementations solve their own mathematics, falsifier web executed	exact identities, twin gates
Mechanism ladder (Section 4.4)	SS-S kernel predictions against B3	attribute agreement mechanism by mechanism, one rung at a time	sealed predictions, paired deltas, $\lvert z\rvert \leq 2$
Reference batches (Section 4.6)	SS-S raw, whole steady-state family at $\beta = 1$	end-to-end validation, replication, ageing arm	per-batch gate sets
Reference statistics (Section 4.7)	the reference itself	measure the noise floor every z is expressed against	$1/\sqrt{n}$ law, bootstrap

4.2 The reference model

4.2.1 WOMBAT and how it schedules work

WOMBAT (version 0.13.3) is an open-source discrete-event simulator for offshore wind operations and maintenance developed at NREL [5]. Failures arrive per mode from Weibull draws, generate repair requests, and are serviced by vessels under a request-based dispatch strategy. Weather enters as an hourly time series against per-vessel operating limits.

The audit that preceded any comparison established the servicing mechanics from the source code and event logs from testing, and the facts matter because several depart from what a reader would assume. Mobilisation is a flat timeout of the configured days, ignores weather, and bills only the flat mobilisation cost, with the day rate not charged during mobilisation. Travel is weather-free, with deterministic legs set by port distance. For around-the-clock vessels there is no end-of-repair return leg, and demobilisation is instant and free. Active repair uses interrupted accumulation: work banks only during accessible hours and pauses otherwise, which is exactly the banked-stretch semantics of Subsection 3.3.3 and the reason the interruptible operator is the correct analytical counterpart. Billing charges every logged action hour (waits, travel, transfers, repair) at the day rate over 24, with mobilisation time free. Finally, the stock failure clocks implement a hybrid ageing convention: the clocks of a component freeze during its own repair but keep running through the repairs of other components and through waiting time. This hybrid coincides with the operational-clock convention established for the models of this thesis only in the memoryless case, a fact that becomes central at $\beta \neq 1$ and motivates Patch I below.

4.2.2 Defects found and the patch series

Making WOMBAT usable as a per-event reference exposed a series of defects and semantic quirks, each diagnosed at source level, patched minimally, and verified by its own signature in the event logs. Table 14 lists the frozen patch series. One long-standing edit predates the series and is present in every build: in the unscheduled in-situ routine, a vessel demobilises immediately when the request queue empties instead of billing idle two hour increments to charter end, which was the default behaviour.

Table 14: Modifications to WOMBAT (the patch series). Each patch modifies a concrete behaviour or assumption from the upstream source, is archived, and verified by the log signature in the last column.

Patch	Defect and fix / Assumption alignment	Verified signature
(demob edit)	Idle vessels billed for 2 h before charter end. Demobilise immediately on empty queue.	no idle vessel rows after repair completion

Patch	Defect and fix / Assumption alignment	Verified signature
A	Blind "weather unsuitable to transfer crew" waits and site-to-port retreats in the window search. Growing-horizon search replaces the blind wait.	no unsuitable-transfer waits, no site-to-port retreats
B	Request leak in the repair manager: in multi-turbine runs around 16% of requests failed to join the queue and were left ignored. Dormant on a single turbine but fixed from the start.	request count conserved
C	Rounding on departures added a mean 0.5 h to every travel leg. Legs made exact.	legs exactly 3.131749 h and 5.694649 h
D	Freshly dispatched vessels arrived location-less and teleported. The existing start-from-port branch is enabled for first dispatch, adding a real billed first leg.	zero teleport rows, first leg billed
E	Sometimes doubled repair hours on re-entry after interruption. Banked hours made exact.	per-mode repair-hour sums equal the yaml values, single-valued constants
F	Reduction of crew re-transfer cycles, made obsolete by G. Not part of the final patch-stack.	(redundant)
G	Window-exhaustion re-transfers: the forecast array held only the initially required clear hours, so an interior storm exhausted it with work remaining, forcing a crew-off, hop, crew-on cycle. The window is refilled in place, the loop exits only at zero remaining hours.	zero re-transfer pairs across all deployments
H	Stranded requests: dispatch fired only at submission, so requests arriving during an active charter could remain unserved indefinitely. Dispatch re-fires at every repair completion.	stranded requests extinct, residual waits are legitimate completion waits
I	Clock convention transplant: a system-wide operational gate replaces the own-subassembly term in both failure and maintenance clock gates, so every failure clock freezes for the entire system downtime, matching A-3.	only one simultaneous failure active

Two of the patches require further discussion due to their relevance in the argument.

Patch H repairs a genuine scheduler defect, not a modelling preference. The mechanism was isolated in a single-seed microscope (seed 17): a request submitted while the charter of the servicing vessel was still active never triggered dispatch again, leaving it stranded for the remainder of the simulation. Its verification proceeded in two stages. A 20 seed testing batch passed six mechanism-level gates: the request population was unchanged at 1,477, fourteen of the twenty seeds were bit-identical to their unpatched counterparts with the remaining six differing only through the intended releases. The 200 seed batch then established the population effect reported in Section 4.6.

Patch I is not a bug fix but a declared convention alignment, and it is the enabling condition for the whole $\beta \neq 1$ benchmarking programme. Since the stock hybrid clocks and the operational clocks of A-3 coincide only in the memoryless case, any comparison at $\beta \neq 1$ requires either building the hybrid convention into the models or building the operational convention into the simulator. The patching direction was chosen because the operational convention is the one with a preferred physical reading (components do not accumulate wear while the turbine is stopped) and a cleaner analytical structure. At $\beta = 1$ the modification is invisible, which is itself a verified prediction.

4.2.3 Build ladder

Modifying a benchmark invites an obvious objection: the reference was edited until it agreed with the model, so what did the comparison test? The answer is structural, and it is carried by the build ladder of Table 15. Four builds are established. B0 is the untouched upstream source, not used, but preserved so every modification is trackable. B1 adds the per-event fixes (the demobilisation edit and Patches A, B, C, D, E, G). B2 adds Patch H and acts as a more faithful reference, the original physics of the model with demonstrated defects removed. B3 adds Patch I and acts as the verification twin, the reference with the clock convention of the thesis models transplanted in.

Table 15: The four reference builds and their roles.

Build	Contents	Role	Batches
B0	upstream 0.13.3, untouched	the version a reviewer can directly install	none
B1	B0 + demob edit + A, B, C, D, E, G	repaired logging and accounting, stock scheduler intact	x10
B2	B1 + Patch H	"faithful" reference, WOMBAT original convention, major defect removed	x12
B3	B2 + Patch I	verification twin, clock convention implemented	x14-x18, thousand-seed set

4.2.4 Alignment conditions

A comparison is only meaningful if the two sides describe the same system. The alignment conditions were fixed before any batch was scored, and input parity verified against the frozen configuration. Table 16 states the frozen input baseline that every batch and every model run in this thesis shares.

Beyond input parity, three convention alignments sustain the comparison. The clock convention is aligned through the patching, the repair process sequence is aligned by construction, and the accounting is aligned by the shared billing rules (tested at every batch by ledger closure). One residual convention mismatch is known and carried openly: the models use 8,760 hours per year for rate conversion while the effective hours of the reference are not pinned, a 0.0685% effect on all rates, small against every tolerance in use but recorded in the discrepancy budget of Section 4.6.

Table 16: The frozen baseline configuration shared by models and reference.

Item	Value
Asset	one bottom-fixed 12 MW turbine, 22 corrective failure modes, all $\beta = 1$ at baseline
Horizon	20 years from 2000-01-01, 175,320 hours
Weather	hourly ERA5 (100 m wind, significant wave height)
Site	single reference offshore site, port distance 116 km
CTV	on site, no mobilisation, 24 h shift, wave limit 2 m, \$3,201.7 per day, leg 3.131749 h
HLV	chartered per request, 240 h unbilled mobilisation, 24 h shift, wind limit 10 m/s and wave 2 m, \$400,000 per day, leg 5.694649 h
Transfers	two crew transfers of 0.25 h per repair, weather-gated
Billing	downtime minus mobilisation at day rates, materials flat at completion
Dispatch	request-based strategy, one request per failure

Key outcome. A per-event-transparent reference: audited servicing semantics, a patch series with verified log signatures, a build ladder separating defect repair (B1, B2) from the declared convention transplant (B3), and one frozen, parity-checked input baseline. Every comparison in this chapter stands on this alignment.

4.3 The validation framework

4.3.1 Verification against validation

As previously mentioned, the two words must be used carefully. Verification establishes that an implementation solves its stated mathematics correctly. Validation establishes that the mathematics describe the reference system appropriately: statistical agreement with the simulator under the alignment conditions of Subsection 4.2.4. A comparison between two of the developed models can serve either purpose depending on configuration. Under the exponential assumption, the semi-Markov solvers must reproduce the Markov solvers exactly, which is verification. Free from it, their divergence is a model property to be characterised, which is explored in the results chapter.

4.3.2 Comparison metrics and error measures

Four measurement layers are used, from coarse to fine:

- Population statistics: For a KPI with reference ensemble mean $\bar{x}$ and standard error SE, a model value m is scored by $z = (m - \bar{x})/\mathrm{SE}$. The KPIs set are availability (operating-time basis, with the event log union as a cross-check), task counts, equipment cost, materials cost, and net energy.

- Paired statistics: Where two runs share seeds (adjacent patch builds, or common random number model pairs), per-seed differencing removes the common variance and resolves effects far below the population noise floor. Paired deltas carry their own standard errors.
- Per-event reconciliation: Event logs are decomposed request by request into mobilisation, travel leg, wait, transfer, and repair segments, matched against the per-event predictions of the model. Residuals at this layer are reported in hours per event and are expected at hour-rounding scale.
- Ledger closure: Every batch must reproduce its own cost accounting exactly (billed dollars equal duration times rate on every row), and batch totals must match to the cent. This is an integrity check, not a statistical test.

Structural counters supplement the four layers on the patched builds: stranded requests, re-transfer pairs, teleport rows, dispatch deferrals, and outage overlaps are counted exactly, but on B3 the mechanism is impossible by construction and a single occurrence indicates a coding error.

In practical terms, the layers read as follows. A z score expresses a model-minus-reference gap in units of the reference standard error, so $|z| \leq 2$ means the gap is indistinguishable from the sampling noise of the batch at the 95% level, and a z near zero on a 200 seed batch still permits absolute gaps of a few hundredths of a percentage point, which is why tolerances are always stated alongside scores. Paired designs change the question: by differencing shared seeds they remove the common fluctuation, so effects far below the population noise floor become resolvable, the 0.5 h transfer mechanism of Subsection 4.4.1 is resolved at ± 0.0003 percentage points this way. Per-event reconciliation is stronger still, as it compares individual episodes rather than averages, and its residuals are expected at hour-rounding scale, not at statistical scale. Ledger closure, finally, is not a statistical test at all but an integrity condition, and a single cent of imbalance fails a batch outright.

4.3.3 Statistical treatment

The reference is stochastic, so its side of every comparison carries uncertainty. Seeds are drawn once per batch and archived, replications are independent and identically distributed by construction, and each batch record stores per-seed rows so any statistic can be recomputed. Confidence intervals use normal theory with bootstrap cross-checks for skewed KPIs (rare-event counts and heavy-tailed costs). The developed samplers (MC-C, MC-S) report their own replication uncertainty and use common random numbers for paired designs.

Key outcome. Four measurement layers from population z scores to exact ledger closure, each with a stated practical reading.

4.4 Main campaign (mechanism-accretion ladder)

Concept. One mechanism per rung: each ladder step adds exactly one mechanism to the configuration below it, on the same seed list, so every paired increment isolates the contribution, and the blame, of one modelling choice.

Table 17: The mechanism-accretion ladder. Each rung changes one thing relative to the rung below. The ladder runs on B3 only, with the convention price carried by the measured summit bridge.

Rung	Configuration delta	Isolates	Acceptance
ML-0	one synthetic mode (exponential failure, deterministic repair), clear weather, zero distance, zero transfers, zero mobilisation	pure alternating renewal, the random-number and bookkeeping layer of the simulator	exact: $A = 1/(1 + T/\mathrm{MTTF})$, closed-form variance, ledger exact
ML-1	+ real travel legs	leg accounting, billed outbound and unbilled return protocol	per-event downtime delta equals the leg exactly, paired delta against ML-0 matches
ML-2	+ crew transfers (2×0.25 h)	transfer accounting	delta equals $2T$ exactly per event
ML-3	+ real ERA5 weather gating	the weather operator (initial wait, interleave, banked span), Patch I onset	per-event phase-conditioned replay residuals at hour-rounding scale, delta matches the expected stretch, the convention exposure is born here and is priced by the measured summit bridge
ML-4	+ the full 16 mode CTV set	superposition, structural no-concurrency under operational clocks, Patch H and G onset	structural zeros on B3, union equals sum, per-mode rates clean in z
ML-5	+ one heavy replacement mode with 240 h mobilisation	mobilisation, the unbilled-mobilisation identity	HLV dispatch gap equals 240 h plus quantisation, billed equals downtime minus mobilisation
ML-6	+ the remaining heavy modes (the full 22 mode default)	heavy-tail cost structure	the summit gates, already verified by the x13 and x14 batches ($z = +0.20$ on availability, ledger to the cent)
ML-7	fresh seeds 201 to 400 at the default	independence replication of the summit	executed as batch x15 (Subsection 4.6.2)

The campaign consists of a mechanism-accretion ladder (ML), with audits run on every batch. The ladder principle is for each rung to add exactly one mechanism to the configuration before it, on the same 50 seed list, so paired increments isolate the contribution of that specific mechanism easily, the mechanism-isolating discipline whose code-comparison lineage the review traced to its current endpoint [63]. Rungs ML-0 to ML-2 have almost exact analytic answers, but from ML-3 upward acceptance is statistical, with per-rung noise re-estimated rather than inherited from the baseline.

The ladder runs on B3, the verification twin, because patch necessity and correctness were already established at the summit by the batch record of Section 4.6 (the seed-17 microscope, the paired deferral census between x10 and x12, and the structural zeros of x13 and x14).

Standing audits run on every batch regardless of rung: the signature audit (all patch signatures of Table 14 re-checked), ledger closure, the request-population census, and the basis-coherence check between the two availability accountings introduced in Section 4.7.

4.4.1 Ladder execution and outcomes

Models under test: the SS-S kernel predictions, rung by rung, against build B3.

The rungs were executed all on B3, all on seeds 1 to 50, with every prediction sealed before its data existed. A pre-launch amendment is on record as the ML-0 configuration was amended after an aborted diagnostic: at exactly zero port distance the simulator computes the travel duration by subtraction and returns a negative epsilon of order 10^{-17} hours, which the event queue rejects. The rung therefore runs at 0.03704 km, a leg of exactly 0.001 h, the residual is priced into the sealed prediction, and a minimal reproduction of the upstream defect is archived with the campaign records. The same amendment moved the dispatch delay g (the request-to-mobilisation gap, at most 0.0186 h throughout, under its standing 0.02 h gate) out of the per-event identity, which is stated as $D - g$.

Table 18: Ladder outcomes at 50 seeds per rung. Predicted availabilities are the closed forms (ML-0 to ML-2) and the numbers from the kernel registered before launch (ML-3 to ML-5). Paired deltas are measured against the rung below on the shared seed list, predicted values in parentheses. ML-4 pairs across a mode-set swap and is reported standalone.

Rung	predicted A (%)	measured A (%)	z	paired ΔA (pp)
ML-0	98.8713	98.8602	−0.44	—
ML-1	98.8015	98.7901	−0.42	−0.0700 (−0.0698)
ML-2	98.7903	98.7790	−0.42	−0.0111 (−0.0111)
ML-3	98.3189	98.3222	+0.08	−0.4568 (−0.4715)
ML-4	98.6179	98.5563	−2.03	mode-set swap
ML-5	98.1423	98.0581	−1.27	−0.4982 (−0.4756)

The identity rungs correctly behave as identities. Across 1,998, 1,996 and 1,996 events the per-event $D - g$ matches its sealed constant to 4×10^{-11} hours, the float unit at timestamp magnitude, travel legs match 116/37.04 to 4×10^{-16}, all 3,992 crew transfers read 0.25

h with zero deviation, and repair segments conserve their work content to 3×10^{-14} hours per request, including through 1,701 weather interruptions at ML-3. Ledgers close to at most 1.9×10^{-9} USD on every rung. The ML-0 variance ratio, the closed-form test of the random-number layer of the simulator, lands at 1.0037 against a $[0.80, 1.20]$ band, and the ratios at ML-1 to ML-3 read 1.0043, 1.0048 and 0.9888.

The paired design delivers the resolution it was built for. The standalone z is -0.42 on all three exact rungs, the shared draw sequence carrying one common fluctuation through every configuration, while the paired increments read $z = -0.12$, -0.01, $+0.73$ and -0.32: the 0.5 h transfer mechanism is resolved at ± 0.0003 percentage points. Per-seed count differences confirm the replay, with 48 of 50 seeds identical between ML-0 and ML-1 with two boundary events dropped, and 50 of 50 identical between ML-1 and ML-2.

At ML-4 the mode superposition is structurally exact, zero overlapping excursions among 3,790 requests across 16 competing modes, and the per-mode z set has a maximum of 2.05 with one mode of sixteen past 2. The standalone availability of the rung sits at $z = -2.03$ with the request count at $z = +1.94$, one fluctuation counted twice, since the count surplus decomposes broadly across twelve modes rather than concentrating, it is recorded here as the single borderline pass, at the rate one expects from six standalone tests. At ML-5 the mobilisation identity holds on the full census, as all 92 heavy-vessel events show a request-to-travel gap in $[240.0019, 240.0161]$ hours, exactly the sealed 240 h plus the dispatch delay, billed cost equals downtime minus $(240 + g)$ to a maximum residual of 10^{-6} USD on events of 5.3 million, and the 17 mode z set peaks at 1.59.

Key outcome. Six rungs executed against sealed predictions: the identity rungs exact to floating-point scale, the weather operator validated per event at hour-rounding scale through 1,701 interruptions, superposition structurally exact, mobilisation exact on the full census, and one borderline standalone pass (ML-4, $z = -2.03$) at the rate expected from six standalone tests. Each mechanism of the analytical description is individually validated before the summit.

4.5 Code verification

Concept. Before the reference is consulted, the family must agree with itself: independent implementations of the same mathematics, meeting on exact identities. This block executes the internal consistency web of Figure 24.

Models under test: all six family members, against each other.

The steady-state layer is verified by exact identities: the shared steady state of 97.9903% is recovered by four independent computations that share no solver code (the SS-S kernel, the TM-C eigendecomposition limit, the TM-S Markov-renewal limit, and the MC-C serialised arm). TM-C passes its full check set, including the exact 45 against 23 state lumping (Subsection 3.4.1). TM-S passes limits, the transient area identity, and the $\beta \to 1$ collapse onto its Markov part.

However, one genuine deficiency was found by stress testing. Under a uniform $\beta = 0.8$ configuration the two-tier checkpoint schedule failed its refinement gate at 2.9×10^{-3}. The schedule was extended to three tiers, the same configuration now passes at 7.3×10^{-4}, and the three-tier rule fires automatically whenever any mode has $\beta < 1$.

The samplers are verified against their deterministic twins. MC-S matches the SS-S occupancy vector in all 45 dimensions (largest $|z| = 2.09$) and its replication mean brackets the closed form. The serialised arm of MC-C converges on the shared steady state, and its union arm reproduces the historical targets of the notebook of 97.9725%.

Key outcome. The family solves its own mathematics: four independent routes reach one steady state, the falsifier web of Figure 24 passes in full, the samplers match their deterministic twins, the SS-S construction ladder is audited rung by rung, and the one deficiency found by stress testing is repaired and promoted to method.

4.6 Validation against the reference batches

Concept. The summit comparisons: the full 22-mode configuration, the analytical family against the patched reference, raw. Five 200-seed batches on build B3 (x14 to x18) form the record, x14 and the fresh-seed replication x15 are reported here, and the pooled thousand seeds constitute the reference statistics of Section 4.7.

4.6.1 The verification twin (x14)

Models under test: SS-S raw against B3, and through the invariance theorem the whole steady-state family at $\beta = 1$.

Table 19: Raw validation of the thesis model family against the verification twin (batch x14, 200 seeds, build B3). The z scores are in reference standard-error units.

metric	x14 (200 seeds)	SS-S	Δ	z
availability [%]	97.9832 ± 0.0361	97.9903	+0.0071 pp	+0.20
tasks [events]	14,923	15,019 expected	−96	−0.78
equipment [$/yr]	$581{,}899 \pm 31{,}617$	557,585	−24,314	−0.77
materials [$/yr]	$20{,}178 \pm 694$	18,808	−1,370	−1.98

Batch x14 (200 seeds, build B3) is the first raw validation batch of the thesis family, with the headline residuals in Table 19. Availability agrees at $z = +0.20$, and every remaining gap traces to a single cause, a Poisson count draw in the rare heavy replacement modes riding through x10, x12 and x14 on the shared seeds. The materials and task gaps are that draw seen from two sides, decomposed and closed on independent data, and once counts and mix are removed the equipment residual is −$3,915 per year, 0.70% of the result and inside per-event weather noise.

4.6.2 Replication (x15)

Models under test: the same predictions as above, on data that did not exist when they were sealed.

The replication rung ML-7 re-runs the default configuration on fresh seeds 201 to 400. The headline hypothesis passed, with 66 heavy-mode events against the predicted band of $[54, 87]$, with x14's outlying 94 regressing onto 70.6. Five secondary gates failed together and were traced to a single LCN count draw at $z = -2.26$, the same rare-mode mechanism already priced, and the pooled 400 seed statistics close at $z = -1.81$ on availability and $z = +0.18$ on materials.

With both conventions validated raw on their own builds, every entry is either in agreement within statistical uncertainty or a named, measured difference.

4.6.3 $\beta = 2$ (x19)

Models under test: the semi-Markov kernel proper, since at $\beta = 2$ the ageing mechanism is live and no exponential shortcut exists.

Batch x19 (200 seeds, 1001 to 1200, build B3) reruns the full configuration with every failure shape set to $\beta = 2$ at fixed MTTF, testing the ageing mechanism. Twenty-year availability: 98.3434% predicted, 98.3665 ± 0.0223 measured, $z = +1.04$. Starts per life: 70.001 predicted, 70.005 measured. The per-mode calibration exercises all 22 ageing regimes at once and peaks at $|z| = 2.27$, one mode of 22 past 2, and the variance falls in the registered direction ($p = 8 \times 10^{-14}$). Every check passed, and the logs carry the ageing signature directly, having no repair requests before the 500th hour where the exponential configuration expects 4.4.

Key outcome. Raw validation at the summit: availability at $z = +0.20$ on the first batch, the replication holding on seeds that did not exist at sealing, and the ageing arm passing every gate at $\beta = 2$ with the ageing signature visible in the raw logs. Every remaining gap is a named, measured mechanism, dominated by one rare-mode count draw.

4.7 Reference seed convergence

Concept. The reference is stochastic, so its noise floor must be measured before any z score in this chapter means anything. This section measures it.

Models under test: none, the reference itself is the subject.

1000 independent seeds were run on build B3. Every KPI converges by the $1/\sqrt{n}$ law, confirmed by bootstrap (Figure 25), and what separates the KPIs is per-seed dispersion, spanning two orders of magnitude from net energy (coefficient of variation 0.58%) to equipment cost (75.4%). Inverting $n = (1.96\,\sigma/\varepsilon)^2$ gives a practical consequence: availability to ± 0.05 percentage points needs several hundred seeds, while the same relative precision on cost is beyond any practical batch size. The healthy-population reference is 98.0063 ± 0.4946 percent per seed, and the thesis models sit at $z = -1.02$ against the pooled mean.

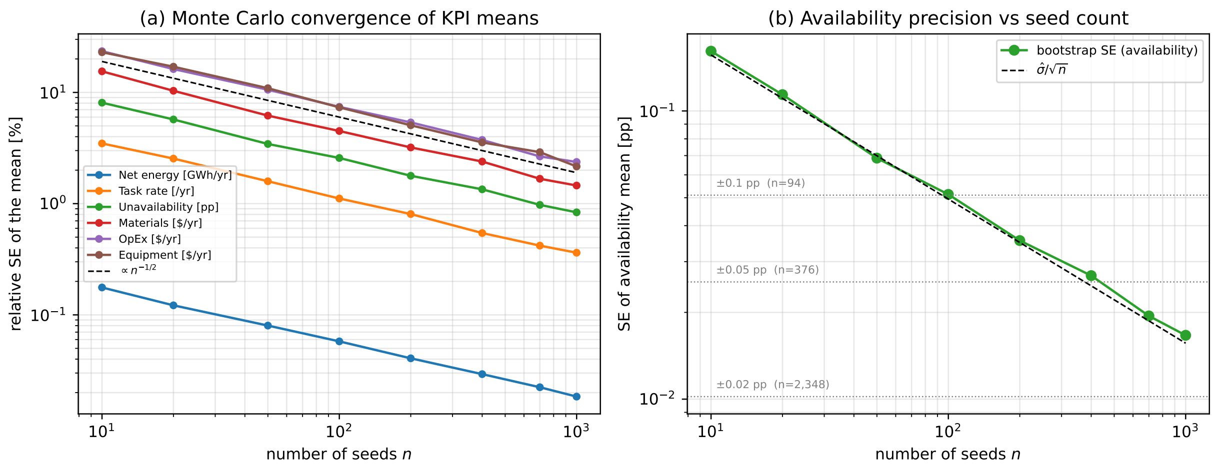

Figure 25: Seed convergence of the 1,000 seeds batch (build B3). Left, bootstrap standard errors of each KPI mean against ensemble size on log-log axes, with the fitted slopes confirming the $1/\sqrt{n}$ law. Right, availability precision against seed count with the seed requirements for common tolerance targets.

4.8 Generalisability and licensed claims

The campaign validates the family on one representative case, a single 12 MW bottom-fixed turbine at one North Sea class site, so the reach of its conclusions deserves a statement. Three layers transfer differently. The mechanism conclusions, that the weather operators reproduce the per-event episode structure, that the superposition and mobilisation identities hold, and that the kernel carries the ageing regime, are properties of the model mathematics against the simulator mechanics, and they transfer to any configuration built from the same mechanisms. On the other hand, the statistical conclusions, the z scores and tolerances, are tied to this input baseline, and a different turbine, site, or maintenance strategy changes the numbers but not the validation procedure, since the ladder and the audits rerun unchanged on any configuration the pipeline accepts. However, what does not transfer without new work is the scope fixed by the assumption register: shared-vessel fleets and dependent assets sit outside A-9, and time-varying rates outside A-8, as Chapter 3 discusses. The validated envelope is therefore the corrective single-asset setting.

Within that envelope, the programme licenses the following for Chapter 5: steady-state availability, task, and cost estimates of the family carry the accuracy measured at the summit, effects below the reference noise floor may be computed analytically and trusted to the extent the mechanism ladder validates their ingredients, the ageing arm is validated end to end at $\beta = 2$, and every known residual, the rare-mode count mechanism, the 8,760-hour rate convention at 0.0685%, and the trailing-transfer asymmetry of A-5, is named, bounded, and carried into the interpretation of results.

Key outcome. The family is validated mechanism by mechanism and at the summit, on sealed predictions, against a per-event-transparent reference whose own noise floor is measured. Chapter 5 may now present analytical results as validated estimates within the stated envelope.

5 Results and Discussion

Chapter 4 closed by licensing the claims of this chapter (Section 4.8): the steady-state models are exact implementations of their mathematics and empirically indistinguishable from the aligned reference at the baseline, the transient models are verified solvers anchored externally at $\beta = 1$ and $\beta = 2$, and any model-reference discrepancy is routed to a named cause. This chapter makes use of that licence. It does not re-argue fidelity, and the certified headline figures are quoted as established where they are first used. What the chapter adds is behaviour analysis: what the validated family says about the reference system that the reference simulator cannot say cheaply, or other models cannot say at all.

The chapter runs in three sections, ordered by research question, and Table 20 restates each question in one line so that no section needs to be read with Chapter 2 open. Section 5.1 characterises the transient behaviour of the family, the time-domain half of RQ2, with the $\beta = 2$ replacement waves anchored externally by batch x19. Section 5.2 answers RQ3 by pricing the interruptibility assumption against a contiguous-window alternative on the same masks. Section 5.3 answers RQ4 by demonstration, propagating joint input uncertainty, pricing vessel capability, lifting envelopes, mobilisation lead time and component reliability, and resolving structure below the noise floor of any practical simulation ensemble. Section 5.4 closes the chapter with its take-home messages, and Chapter 6 states the verbatim answers to the research questions.

Unless explicitly stated otherwise, all steady-state results are computed on the $\beta = 1$ build and, by the invariance theorem of Section 3.7, speak for all shapes at once under fixed MTTF, with Section 5.1 marking exactly where that behaviour ends in time.

Table 20: The research questions restated, with the sections of this chapter that answer them. RQ1 is answered by Chapter 4, and Chapter 6 states all four answers verbatim.

	Question, in one line	Answered in
RQ1	Can the analytical family reproduce the certified reference within its own statistical resolution?	Chapter 4
RQ2	Where is the memoryless simplification adequate, and where does the semi-Markov description become necessary?	Section 5.1, steady state per the invariance theorem
RQ3	How large is the interruptibility effect, and under which conditions does the execution choice matter?	Section 5.2
RQ4	What does the analytical route resolve and afford that simulation cannot, and what studies does that enable?	Section 5.3

5.1 Transient behaviour: the boundary ramp and the replacement waves

This section serves RQ2.

Before any curve, the engineering stake deserves stating. Availability guarantees and warranty periods are written over the first years of operation, exactly where a cold-started fleet is still relaxing towards its long-run behaviour. Financing and O&M budgeting consume time-resolved expectations over a debt tenor, not limits. And end-of-life planning is a question about replacement waves, not about steady states. Whether a model can represent the transient at all, and how accurately, is therefore an engineering question before it is a mathematical one.

Chapter 4 verified the two transient solvers and deliberately stopped there, since away from the exponential nesting point their divergence is a model property, and characterising it belongs to this chapter. In this section, both models consume the same 22-mode kernel, the same weather operators and, by construction, the same per-mode means, so their stationary limits coincide exactly, at 97.9903 % availability, and at $\beta = 1$ the TM-S solution collapses onto the TM-C solution to within 4×10^{-5} percentage points at every probed time. Everything reported below is therefore a pure transient contrast between two descriptions calibrated to the same steady state, which is precisely the regime in which the memoryless simplification is usually defended.

5.1.1 The boundary ramp at $\beta = 1$

A cold-started system, all components new and the turbine up at $t = 0$, begins at unit availability and relaxes towards the steady state, and this boundary ramp is the entire transient content of the $\beta = 1$ problem. Figure 26 shows the first 3,000 hours of both solutions and Table 21 collects the metrics. The two models agree on where they are going and disagree on how they get there: the TM-S trajectory last leaves the 0.1 percentage-point band around the steady state after 498 hours, the TM-C trajectory after 798 hours, which is a sixty percent later, and the gap widens down the tolerance ladder, reaching 1,335 against 2,434 hours at the 0.005 point band.

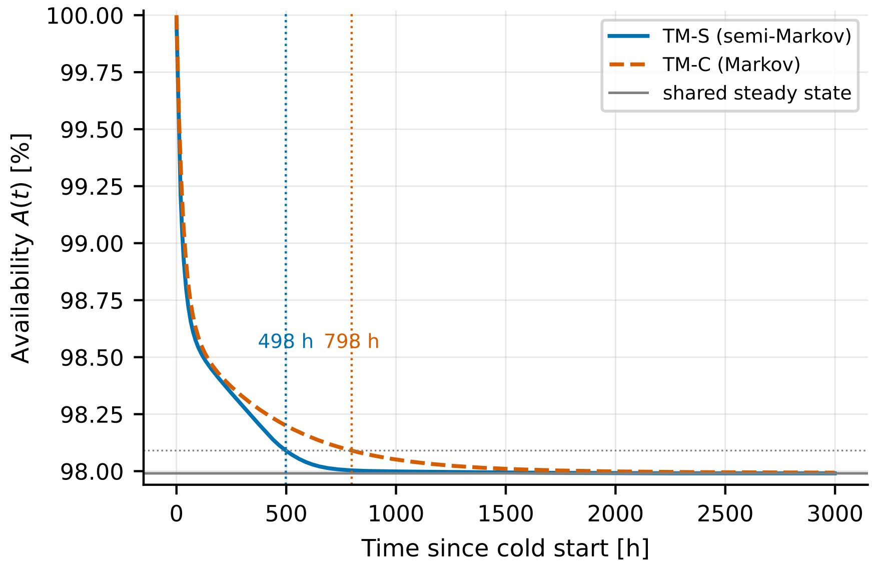

Figure 26: The cold-start boundary ramp at $\beta = 1$. Both models are calibrated to identical means and share the steady state exactly. The dotted vertical lines mark the last exit from the 0.1 pp band, 498 h for TM-S and 798 h for TM-C.

Table 21: Transient metrics of the two solvers at $\beta = 1$ under identical inputs and a shared steady state of 97.9903 %. The boundary credit is $\int_0^T (A(t) - A_\infty)\,dt$ over the 20-year horizon.

Metric	TM-S	TM-C
settling $\tau(0.1\,\text{pp})$ [h]	498	798
settling $\tau(0.05\,\text{pp})$ [h]	589	1,081
settling $\tau(0.02\,\text{pp})$ [h]	713	1,494
settling $\tau(0.01\,\text{pp})$ [h]	892	1,883
settling $\tau(0.005\,\text{pp})$ [h]	1,335	2,434
bulk relaxation constant [h]	356	1,541
boundary credit [avail-h]	+2.29	+3.35
availability, year 1 [%]	98.0166	98.0287

The relaxation structure behind the numbers is visible on the logarithmic axis of Figure 27: the semi-Markov solution sheds its bulk transient at a fitted rate of 2.807×10^{-3} per hour, an effective 356-hour constant, because its nearly deterministic repair spans concentrate the renewal cycle, whereas the exponential chain relaxes at its spectral gap of 6.491×10^{-4} per hour, a constant of 1541 hours from the outset. It is worth noting that the ordering reverses in the far tail, where the TM-S lattice relaxes at 5.598×10^{-4} per hour, marginally slower than the TM-C gap, but by then both trajectories sit within thousandths of a point of the limit and the distinction carries no practical content.

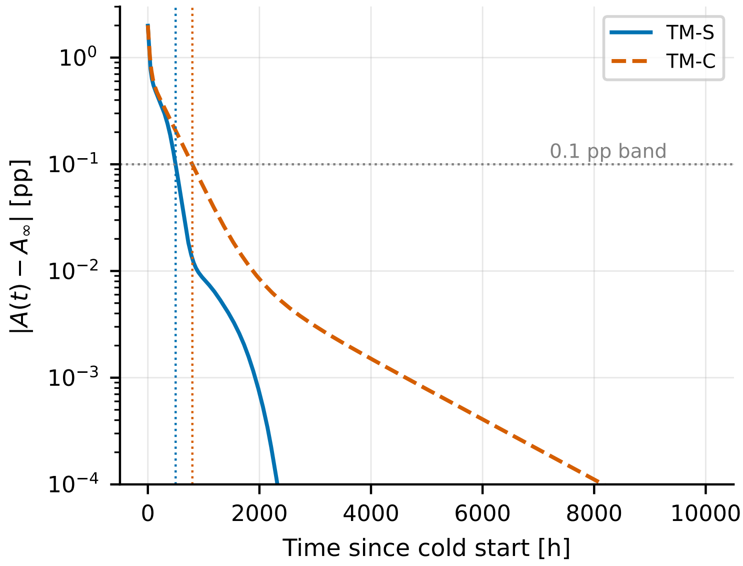

Figure 27: Distance to the steady state on a logarithmic axis. The semi-Markov solution relaxes at an effective 356-hour constant in the bulk, and the exponential chain at 1,541.

The integrated version of the contrast is the boundary credit, the excess availability the cold start buys over a system born in steady state, shown cumulatively in Figure 28. The semi-Markov value is +2.29 availability-hours over the twenty-year horizon, and the exponential chain awards +3.35, an overstatement of 46.0 per cent. The number is not an accident of this configuration: Gate 4 of the model proves the $\beta = 1$ credit equal to the closed form $A_\infty\,(m_{2c}/2\mu_c - 1/\Lambda)$, so the entire transient is carried by the first two moments of the renewal cycle, and the second moment is exactly the quantity an exponential holding-time model is not free to choose. Therefore, the C-family misprices the ramp because matching the means pins its second moment at the exponential value, and the 46% is the price of that constraint, measured. However, at an absolute scale, 2.29 availability-hours across 175,320 is a genuinely small effect. Nevertheless, the next regime is more interesting.

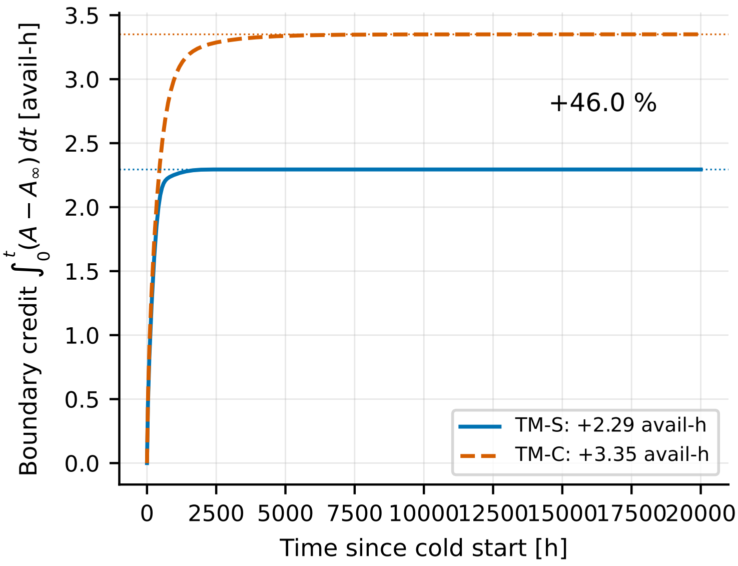

Figure 28: Cumulative boundary credit of the cold start. Both curves converge within the first two years. The exponential chain overstates the converged credit by 46%.

Takeaway. At $\beta = 1$ the two descriptions share the steady state and still disagree about the approach to it: the exponential chain overstates the settling time by sixty per cent (798 against 498 hours at the 0.1 pp band) and the cold-start boundary credit by 46% (3.35 against 2.29 availability-hours), because matching the means pins its second cycle moment at the exponential value. The absolute effect is small at this baseline, so for long-run answers the simplification stands, but any commitment written over the first months of operation, such as a ramp-up availability guarantee, is systematically mispriced by the memoryless description.

5.1.2 $\beta = 2$ at a fixed MTTF

Under a fixed-MTTF change of shape the steady state is invariant, since the closed form uses the means alone, and the inputs of the C-family are exactly that, therefore its entire trajectory, not merely its limit, is unchanged at any β. Whatever a wearing fleet does in calendar time is structurally invisible to the exponential description. Figure 29 shows what it misses. With every mode reparametrised to $\beta = 2$ at fixed MTTF, the cold-started year-one availability is 99.5030%, 1.47 percentage points above the shape-invariant TM-C prediction of 98.0287, because new Weibull components with increasing hazard barely fail in their first year. The trajectory then descends through replacement waves, the first pitch-system wave peaking at 1.59 years and the generator and converter waves at 2.39 years, reaching 98.4689% over years three to five and 98.1496 over the second decade, and its deepest point, 98.1238% at 19.93 years, still sits 0.13 points above the steady state. The horizon simply ends before the component ages decorrelate from the common installation date. The amplitude-weighted settling formula of the model places the binding contribution on the 13-year converter replacement, whose waves cannot decay within twenty years, and the measured tail offset of +0.1368 percentage points over years seventeen to twenty closes that account quantitatively. A twenty-year design life at $\beta = 2$ is, in other words, one long transient, and the steady state to which every model in this family is calibrated is never actually reached.

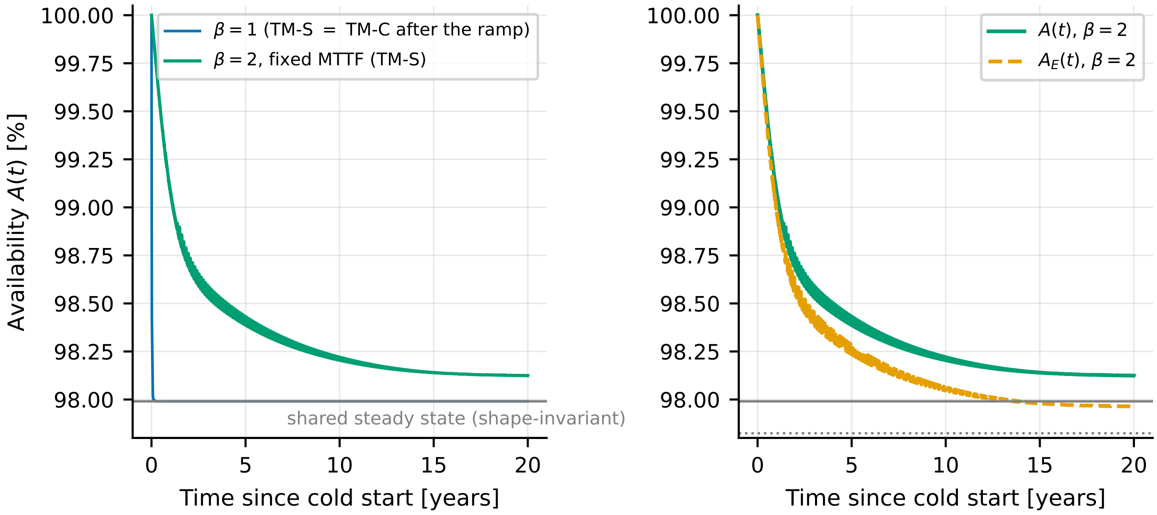

Figure 29: Where the transient becomes the phenomenon. Left, cold-start availability at $\beta = 1$ against $\beta = 2$ at fixed MTTF. The steady state is shape-invariant and the exponential model therefore predicts the $\beta = 1$ curve at any shape. Right, time and production-based availability at $\beta = 2$, with neither settling within the horizon.

Two further properties of the $\beta = 2$ solution matter for the argument of this chapter. First, it is externally anchored: the window-weighted twenty-year availability of the regenerated trajectory is 98.3434%, the comparator against which the cold-start simulator batch x19 of Subsection 4.6.3 measured 98.3665 ± 0.0223, agreement at $z = +1.04$, together with the per-mode start calibration across all 22 aging regimes. The replacement waves in Figure 29 are therefore not a model curiosity but a validated prediction of the aligned discrete-event reference. Second, the production-based estimand carries through the transient: $A_E(t)$ tracks the waves 0.16 points below $A(t)$ in the tail, against the steady gap of 0.166, so the energy correction of the S-family is available at every instant of the trajectory rather than only in the limit, a structure no constant correction factor applied afterwards could reproduce.

Takeaway. At $\beta = 2$ the transient is the phenomenon: year-one availability sits 1.47 pp above the shape-invariant exponential prediction, the trajectory descends through validated replacement waves to its deepest point of 98.12% at 19.93 years, and a twenty-year design life never reaches the steady state to which every model is calibrated. The exponential description misses the required capabilities, so any calendar-indexed question on an ageing fleet, warranty energy budgets, expenditure phasing, end-of-life planning, requires the semi-Markov transient, whose $\beta = 2$ predictions the reference confirmed at $z = +1.04$ (Subsection 4.6.3), with the production-based correction available at every instant.

5.2 Interruptibility

This section serves RQ3.

The model family assumes interruptible repair execution: work progresses during accessible weather, pauses when conditions deteriorate, and resumes later without losing the progress already made. Chapter 4 established that the reference simulator implements exactly this execution, so the assumption itself is validated by construction. This section therefore turns to the comparison the validation deliberately left open: interruptible execution against a contiguous-window treatment, in which the full task must fit inside one uninterrupted weather window. It quantifies how strongly this single modelling choice moves the predicted O&M performance, and it locates the conditions under which the choice changes what a model concludes.

5.2.1 Two approaches and where the literature stands

Two idealised execution semantics bracket how a weather-limited task can be carried out. Under the banked, or resume, semantics, work accumulates during every accessible hour and pauses through inaccessible ones, so a task requiring L active hours completes once L accessible hours have elapsed, however fragmented. Under the contiguous-window approach the task requires one uninterrupted accessible run of at least L hours, and nothing counts until such a run arrives. The developed models implement the former, and so does the reference simulator, whose interrupted accumulation was established from source in Section 4.2. The discrete-event tool NOWIcob likewise resumes started tasks across shifts, giving them the highest scheduling priority [58]. The contiguous semantics, on the other hand, is the foundation of the closed-form access literature, where a maintenance delay is modelled as the wait for a weather window at least as long as the task [6].

In reality, the process physics of turbine repair does not sit uniformly on either side. Crane lifts must complete once begun, and are planned as non-interruptible marine operations [111, 112], but the lift itself occupies hours within replacement tasks that last weeks. Composite blade repairs contain adhesive and laminate cure cycles that cannot be paused, since an interrupted cure is a defective bond, and the literature treats curing under temperature and humidity constraints as the governing step [113]. Most repair labour, by contrast, is interruptible in the plain sense: technicians operate by shifts and return the next weather window to continue. A whole task is therefore a sequence of interruptible labour punctuated by short non-interruptible sub-operations. One further variant deserves a remark because it collapses onto one of the two: a no-forecast restart process, in which work is attempted in every accessible hour but progress is lost on interruption, completes at exactly the same moment as the perfect-forecast contiguous operator, since both finish at the first accessible run of length L, and the variants differ only in the wasted hours worked during failed attempts. Those wasted hours are computed below, and they quantify what forecasting is worth under non-interruptible execution.

5.2.2 Two operators on one mask

Both operators are evaluated on the same hourly accessibility masks, with the same cyclic phase convention: the reported span is the mean over all 175,320 possible start hours of the record. For the banked operator this is the locked kernel walk. For the contiguous operator the span from a start phase is the wait until the first opportunity to begin L uninterrupted accessible hours, either the remainder of the current run if it suffices or the start of the next qualifying run, plus L itself, and the operator is declared infeasible when the record contains no accessible run of length L, in which case the task never completes. The energy loss over the span is walked jointly with the power series, exactly as for the banked operator, so production-based availability carries through both approaches. Costs use the unchanged billing rules, billed hours equal to travel plus span at the day rate, so the comparison isolates the waiting behaviour and holds every other convention fixed.

5.2.3 Window supply of the reference site

Everything in this section is governed by one empirical object, the run-length distribution of the accessibility masks, and it is worth stating before any model output. The CTV mask (waves at most 2 m) is accessible 77.8 % of the time, in 1,554 runs of mean length 87.7 h, median 40 h and 90th percentile 189 h, and its longest run over the twenty-year record is 2,322 h. The HLV mask, which adds the 10 m/s lifting wind limit, is accessible 59.0 % of the time in 3,970 runs of mean length 26.0 h, median 15 h and 90th percentile 66 h, and its longest run is 339 h (Figure 30). The contrast between the masks is not the accessible fraction, which differs by a modest factor of 1.3, but the fragmentation. The wind limit cuts the accessible time into runs a third as long, and the longest window the site ever offers a lifting operation is two weeks. Since the largest repair job in the configuration (blade replacement) requires 864.5 active hours (five weeks of continuous access), one conclusion is available before any solve: whole-task contiguous execution of the blade replacement is not rare at this site, it is impossible.

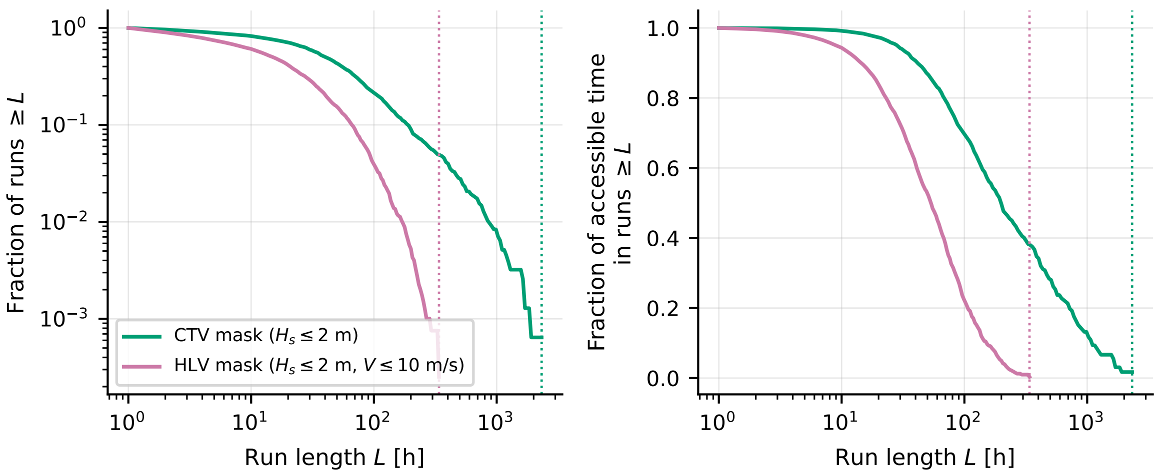

Figure 30: Window supply of the reference site. Left, survival function of accessible-run lengths for the two vessel masks, with the longest run of each marked (2,322 h CTV, 339 h HLV). Right, fraction of all accessible time contained in runs of at least L hours.

Takeaway. The masks decide the argument before any solve. The two vessels differ modestly in accessible fraction (77.8% against 59.0%) but drastically in fragmentation, since the wind limit cuts the heavy-vessel runs to a mean of 26 hours and caps the longest at 339 (two weeks) while the largest job needs 864.5 active hours. At this site, whole-task contiguous execution of the blade replacement is impossible, so any whole-window model is structurally unable to describe the heaviest intervention.

5.2.4 Price of requiring the window

Figure 31 compares the mean span per event under the two operators as the active requirement L grows, on each mask. The two approaches are practically interchangeable for short tasks: at $L = 10\,\text{h}$ the contiguous span exceeds the banked one by 4 % on the CTV mask and 13 % on the HLV mask, which is small against every other convention in the model. The divergence sets in once L becomes comparable to the typical run: at $L \approx 100\,\text{h}$ the ratio is already 2.4 on the CTV mask and 9.1 on the HLV mask, and as L approaches the longest run of the record the contiguous span diverges, with a vertical asymptote at 2,322 h and 339 h respectively, beyond which the operation is not feasible. The banked span, by contrast, grows essentially linearly with slope one over the accessible fraction at every scale, which is why its stretch factors remain bounded.

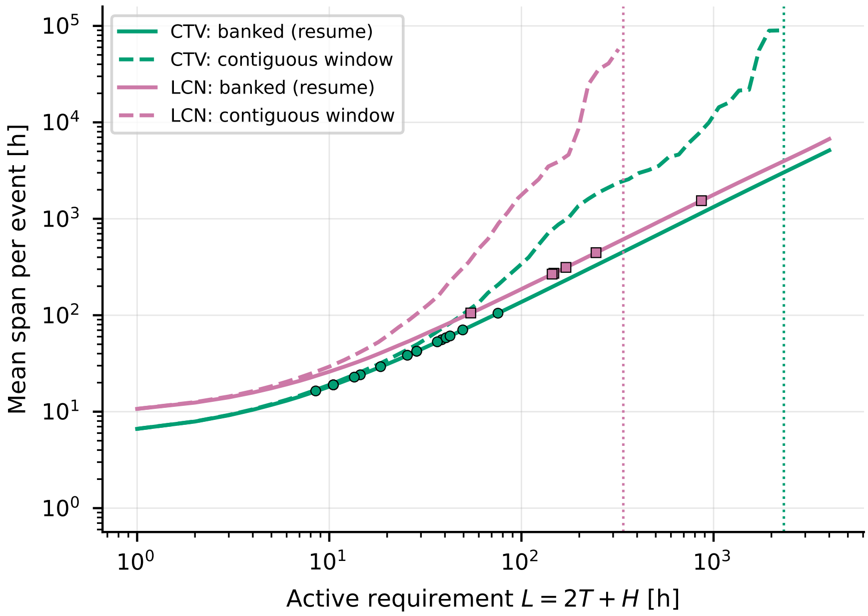

Figure 31: Mean span per event against the active requirement L under the two approaches, per vessel mask. The filled circles and squares mark the 22 configured modes at their active requirements on the banked curves, circles for the CTV-serviced modes and squares for the heavy-vessel modes, and the legend label LCN denotes the heavy-lift (HLV) accessibility mask. The dotted vertical lines are the longest accessible runs of each mask, the asymptotes of the contiguous operator.

Table 22 and Figure 32 give the per-mode census. Under the banked operator every one of the 22 modes has a stretch factor (span over active hours) between 1.39 and 1.94, with the 864.5 h blade replacement at 1.78. Under the contiguous operator the sixteen CTV modes stretch between 1.74 and 2.76, a noticeable but bounded penalty, whereas the heavy-vessel modes leave the physically plausible range entirely: the 54.5 h electrical replacement stretches by 7.4, the 144.5 to 170.5 h replacements by 25 to 26, meaning mean spans around 3,700 to 4,500 hours, the 244.5 h generator replacement by 143, a mean span of 34,935 hours or very nearly four years, and the blade replacement never completes. The wasted-hours column prices the no-forecast restart variant: a generator replacement executed by blind restart would burn 20,215 accessible working hours in failed attempts, whereas the CTV minors waste between half an hour and fifteen hours, so perfect forecasting is worth essentially nothing for short interruptible-adjacent work and everything for long non-interruptible work.

Table 22: Per-mode contrast at baseline. $L = 2T + H$ is the active requirement, stretch is mean span over L. Wasted hours are the accessible hours consumed by failed attempts under restart execution.

Mode	Vessel	L [h]	Stretch, banked	Stretch, contig.	Wasted [h]
minor ballast pump repair	CTV	8.5	1.93	2.01	0.5
power electrical minor repair	CTV	10.5	1.81	1.90	0.7
yaw system minor repair	CTV	10.5	1.81	1.90	0.7
main shaft minor repair	CTV	10.5	1.81	1.90	0.7
generator minor repair	CTV	13.5	1.69	1.80	1.2
power converter minor repair	CTV	14.5	1.66	1.79	1.4
minor pitch system repair	CTV	18.5	1.58	1.74	2.2
blades minor repair	CTV	25.5	1.51	1.74	4.5
power converter major repair	CTV	28.5	1.49	1.76	5.7
power electrical major repair	CTV	28.5	1.49	1.76	5.7
main shaft major repair	CTV	36.5	1.45	1.88	10.7
major pitch system repair	CTV	38.5	1.45	1.90	12.0
yaw system major repair	CTV	40.5	1.44	1.93	13.5
blades major repair	CTV	42.5	1.43	1.95	14.9
generator major repair	CTV	49.5	1.42	2.07	21.8
major pitch system replacement	CTV	75.5	1.39	2.76	67.6
power electrical major replacement	HLV	54.5	1.94	7.37	145.1
main shaft replacement	HLV	144.5	1.85	25.6	2,006
yaw system major replacement	HLV	147.5	1.85	25.1	2,008
power converter replacement	HLV	170.5	1.84	26.3	2,462
generator major replacement	HLV	244.5	1.82	142.9	20,215
blades major replacement	HLV	864.5	1.78	∞	∞

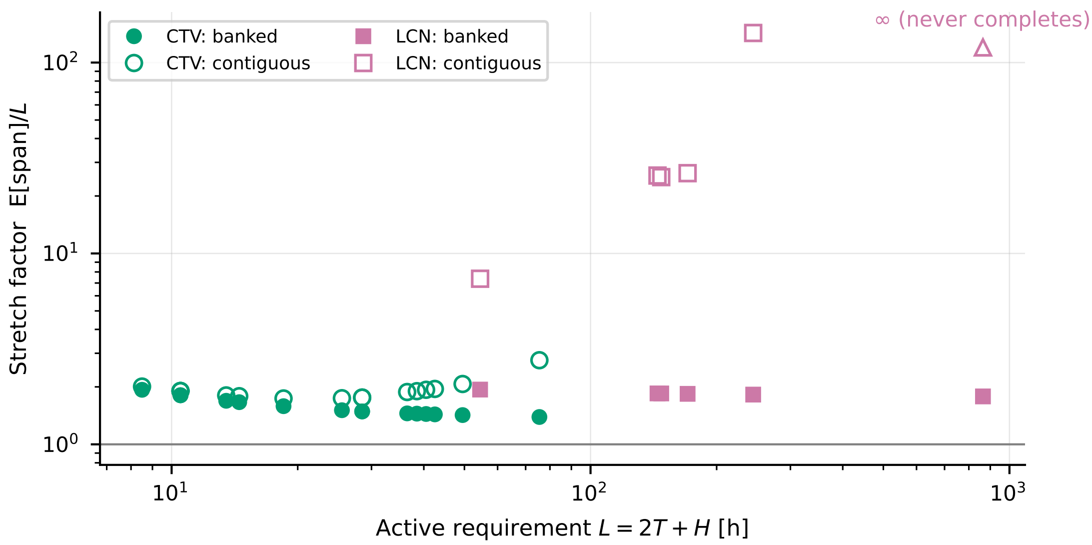

Figure 32: Stretch factors of the 22 modes under both operators against the active requirement. Filled markers show the banked operator and open markers the contiguous one, and the legend label LCN denotes the heavy-lift (HLV) accessibility mask. The ∞ annotation on the blade replacement means that no accessible run of the required 864.5 hours exists anywhere in the twenty-year record, so its contiguous span is unbounded and the mode never completes under that operator.

Takeaway. Interruptibility is a two-regime convention governed by the ratio of the task to the runs of its mask. Every banked stretch factor sits between 1.39 and 1.94, while the contiguous treatment stretches the CTV modes by a bounded 1.74 to 2.76 and drives the heavy modes out of physical plausibility, a factor 143 (four years of mean span) for the generator replacement and infeasibility for the blades. The wasted-hours census prices forecasting, being worth almost nothing for short interruptible work, and 20,215 accessible hours per generator replacement under blind restart, so forecast value concentrates exactly where execution is non-interruptible.

5.2.5 System consequences

Table 23 propagates the operators through the full closed-form chain. Imposing contiguous execution on every mode is not viable: because one mode can never complete, its mean downtime is infinite and thus the steady-state availability is zero. Restricting the contiguous requirement to the modes where it is feasible at all still costs 6.95 percentage points of availability and multiplies the annual direct cost by 18.5, to 10.66 MUSD, driven entirely by the multi-thousand-hour billed spans of the heavy modes. Imposing it only on the sixteen CTV modes, the regime where the semantics question is genuinely open, costs 0.167 percentage points of availability, 0.182 points of production-based availability, 102.9 MWh of annual production and about 1,000 USD of direct cost. Priced with the economic layer of Section 5.3, the whole-window treatment of the short repairs amounts to roughly 11,300 USD per turbine-year at the central electricity price, real but second order. It is worth noting that the time-versus-energy gap widens from 0.166 to 0.180 percentage points under the contiguous CTV scenario, since the added downtime is pure storm waiting, so the interruptibility choice also perturbs the estimand distinction of the S-family, though mildly.

Table 23: System KPIs under the execution-semantics scenarios. The banked row is the certified result headline, while the contiguous rows re-solve the chain with the whole-window operator on the stated mode set.

Scenario	A [%]	A_E [%]	Cost [USD/yr]	Net AEP [MWh/yr]
banked, all modes	97.9903	97.8243	576,394	55,444.1
contiguous, CTV modes only	97.8230	97.6428	577,407	55,341.3
contiguous, all feasible modes	91.0431	90.7875	10,663,281	51,455.9
contiguous, all modes	0	0	undefined	0

The severity sweep of Figure 33 locates the conditions under which the CTV-side contrast stops being second order. As the wave limits tighten, the run-length distribution compresses, the longest CTV run falls from 2,322 h at the baseline 2 m to 254 h at 1 m, and the availability penalty of the contiguous treatment grows from 0.014 percentage points at 4 m through 0.167 at the baseline to 3.23 percentage points at 1 m, a factor of 230 across the sweep. Every CTV mode remains feasible throughout, since even the compressed record keeps one run above the largest 75.5 h job, so the growth is entirely the window-wait mechanism rather than an infeasibility cliff. On the HLV mask the longest run saturates at 339 h for all wave limits above 1.4 m because the 10 m/s wind limit, not the waves, terminates the long runs, which is worth stating: relaxing wave limits (without relaxing wind limits too) cannot buy contiguous feasibility for the heavy jobs at this site, only interruptibility can. Figure 34 makes the same point in the job-length dimension for the two largest jobs of each vessel: the banked and contiguous spans of the pitch replacement separate smoothly with growing L, while the contiguous span of the blade replacement diverges at $L = 339$ h, at barely 39% of its actual requirement.

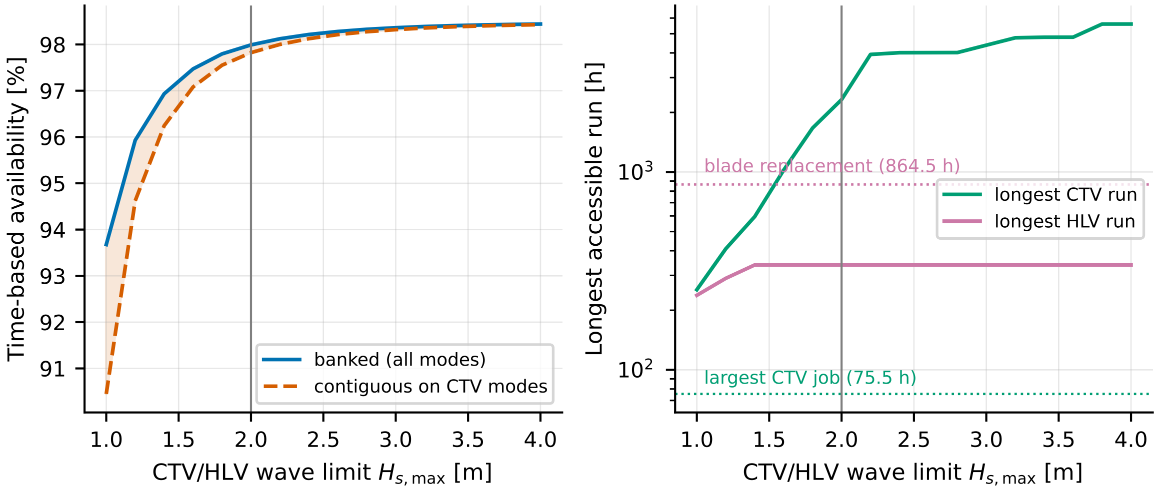

Figure 33: Severity dependence of the interruptibility contrast. Left, availability under the banked operator and under contiguous execution of the CTV modes, against the wave limit. Right, the longest accessible run per mask, with the largest CTV job and the blade replacement marked. The HLV curve saturates at 339 h because the wind limit terminates its long runs.

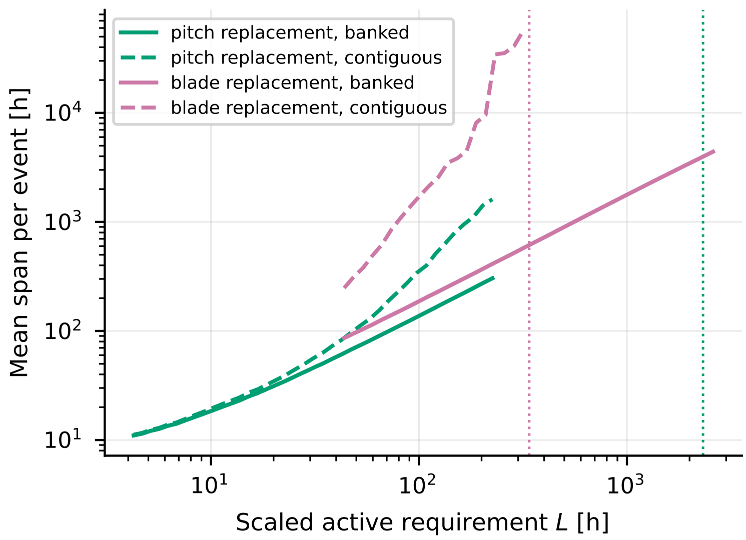

Figure 34: Mean span against a scaled active requirement for the largest CTV job (pitch replacement) and the largest HLV job (blade replacement) under both operators. The contiguous span of the blade replacement diverges at the 339 h asymptote, 39% of the actual requirement.

Takeaway. At system level the execution choice spans four orders of importance, as on the sixteen CTV modes it is a second-order convention worth 0.167 pp of availability and roughly 11,300 USD per turbine-year, on the feasible heavy modes it costs 6.95 pp and multiplies direct cost by 18.5, and on the full mode set it is degenerate, since one infeasible mode sends steady-state availability to zero. The severity sweep locates the boundary, the CTV-side penalty grows by a factor 230 as wave limits tighten from 4 m to 1 m, and the wind limit caps the heavy windows at 339 h regardless of waves, so relaxing wave limits alone cannot rescue contiguous execution. For O&M modelling the recommendation is that execution treatment should be declared as explicitly as the failure rates and considered carefully, because at the heavy end it outweighs every other convention in the model.

5.3 Uncertainty propagation and design support

This section serves RQ4.

Of the studies below, three carry the primary findings, the sensitivity structure of Subsection 5.3.1, the dependence result of Subsection 5.3.2, and the converter concentration of Subsection 5.3.6, while the three procurement studies between them are demonstrations of the same machinery on decisions an operator genuinely faces.

Chapter 4 established that the analytical hierarchy reproduces the certified reference behaviour, and the two sections above characterised its transient and execution-semantics content. This section is concerned with what that fidelity buys. A steady-state solve of the SS-S model completes in approximately 13 ms, consistent with the a priori cost ordering of Section 3.9, and the closed form $A = 1/(1 + \sum_i \lambda_i D_i)$ admits exact derivatives, therefore entire classes of studies that are impractical for a discrete-event simulator become routine: dense one-at-a-time sweeps, exact elasticities, joint propagation of input uncertainty through tens of thousands of samples, and the systematic enumeration of design alternatives. For perspective, the single sweep presented in Subsection 5.3.1 comprises 2223 operating points. Resolving each point in WOMBAT to the precision at which the effects discussed below are visible would require on the order of at least 200 seeds per point, that is, roughly 4.4×10^5 lifetime simulations for one figure, and the uncertainty study of Subsection 5.3.2 would require twenty times as many. The analytical model performs the former in minutes and the latter in under a second.

The questions addressed here are not artificial. The selection of O&M vessel fleets has been identified as one of the key strategic decision problems in offshore wind [55], and a six-tool benchmarking study found that the resulting vessel rankings are particularly sensitive the limiting significant wave height for turbine access [44]. With maintenance responsible for up to a third of the cost of energy [25], and with failure rates and accessibility repeatedly identified as the dominant drivers of O&M cost and availability [114, 8], real practical questions a wind farm operator faces are of the following kind: Which crew transfer vessel should be chartered when a conventional catamaran is limited to 1.5 m significant wave height but a surface effect ship offers 2.5 m at a fifty per cent day-rate premium? How large a premium is a heavy lift vessel with a wider lifting envelope worth? Does a shorter mobilisation lead time justify a more expensive charter arrangement? And how much capital expenditure does a ten per cent reliability improvement of a single component justify? The remainder of this section answers these questions with the developed SS-S model.

To do so, one economic layer is added on top of the certified cost convention. The direct cost C reported throughout this thesis excludes the value of lost production, therefore we define the total annual cost at electricity price p as

$$\mathrm{TC}(p) \;=\; C \;+\; p\left(E_{\mathrm{gross}} - E_{\mathrm{net}}\right), \tag{19}$$

where $E_{\mathrm{net}} = A_E \cdot E_{\mathrm{gross}}$ is the production-based net AEP, with A_E the production-based availability of Equation 10, so that the lost-production term is priced on the energy-weighted availability rather than the time-based one. This is one of the places where the A–A_E distinction of the S-family carries direct monetary meaning, since pricing the time-based downtime would understate the loss by construction. A fixed price is used, following the convention adopted for revenue-loss valuation in the O&M literature [115, 51]. Because market prices vary widely, from strike prices near 60 USD/MWh in recent European auctions to substantially higher merchant exposure, results are reported at $p \in \{60, 100, 150\}$ USD/MWh with 100 USD/MWh as the central case. At the baseline the lost production is $56\,677.3 - 55\,444.1 = 1\,233.2$ MWh/yr, which at the central price adds 123 315 USD/yr of lost revenue to the direct 576 393.60 USD/yr, giving $\mathrm{TC}(100) = 699\,709$ USD/yr. Where a one-off investment must be weighed against an annual saving, the saving is capitalised with a 20-year annuity at a 6 % real discount rate, $a = 11.47$.

Although the studies below vary inputs over wide ranges, the maintenance strategy itself is held fixed, therefore the far-from-baseline points should be read as stress tests of the modelled strategy rather than as forecasts of what an adaptive operator would experience. This caveat is returned to in Subsection 5.3.7.

5.3.1 Local sensitivity structure of the reference system

The foundation for everything that follows is a registered one-at-a-time sensitivity study. Thirteen parameter groups, covering failure rates, active repair durations, material costs, the two day rates jointly and separately, vessel transit speed, the wave limits jointly and per vessel, the HLV wind limit, the HLV mobilisation time and the crew transfer time, were swept with multipliers from 0.30 to 2.00 in steps of 1 %, that is, 171 points per parameter with all other inputs at baseline. The same grid was applied to the failure rate of each of the 22 modes individually, and one two-parameter interaction grid (failure rates against wave limits, 35×35 at 5 % steps) was added. Before the sweep was run, six elasticities with closed forms at the baseline point were registered as predictions, including $\varepsilon_A(\lambda) = -(1-A)$, $\varepsilon_C(\lambda) = A$ and the mobilisation elasticity $\varepsilon_A(\mathrm{mob}) = \varepsilon_C(\mathrm{mob}) = -A\sum_{\mathrm{HLV}} \lambda_{h,i}\,\mathrm{mob}_h = -0.002657$. All six were subsequently reproduced by the numerical central differences to within 4×10^{-8}, which can be considered a strong internal consistency check: the sweep machinery demonstrably computes exact derivatives of the certified closed form rather than regression slopes through simulator noise. Table 24 collects the resulting elasticities, and the sweeps themselves are presented alongside the observations that follow.

Table 24: Baseline elasticities $\partial \ln(\text{KPI})/\partial \ln(\text{parameter})$ from central differences at $\pm 1\,\%$.

Parameter	Availability A	Direct cost C
Failure rates (all modes)	-0.0201	$+0.9799$
Repair durations (all modes)	-0.0113	$+0.9225$
Material costs	0	$+0.0326$
Vessel day rates (both)	0	$+0.9674$
CTV day rate	0	$+0.0279$
HLV day rate	0	$+0.9395$
Vessel transit speed	$+0.0014$	-0.0173
Wave limits $H_{s,\max}$ (both)	$+0.0157$	-0.3167
CTV wave limit	$+0.0145$	-0.0150
HLV wave limit	$+0.0012$	-0.3018
HLV wind limit	$+0.0048$	-1.2064
Mobilisation time (HLV)	-0.0027	-0.0027
Crew transfer time	-0.0002	$+0.0016$

Four structural observations organise the results. First, availability and cost answer to essentially disjoint sets of levers, as the elasticity matrix of Figure 35 makes visible at a glance. Day rates and material costs have exactly zero availability elasticity, since they enter the model only after the state probabilities have been determined, whereas the weather limits, which dominate availability, have comparatively modest direct-cost elasticities. The one large exception is deliberate: failure rates and repair durations move both KPIs, and their cost elasticities of $+0.9799$ and $+0.9225$ are close to unity because almost every dollar in the system is billed per task. It is worth noting that the per-mode cost elasticities sum exactly to A and the per-mode availability elasticities to $-(1-A)$, two identities of the closed form that the numerical sweep verifies to machine precision.

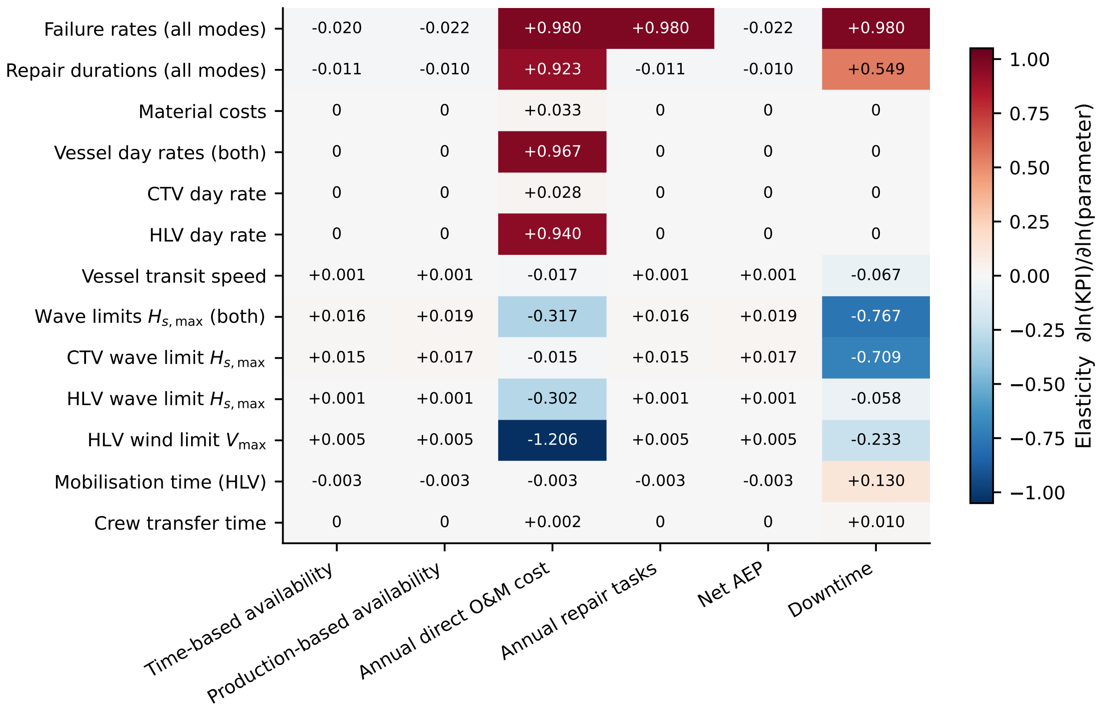

Figure 35: Elasticity matrix of the reference system at the baseline point, across six KPIs. Cost levers and availability levers form almost disjoint sets.

Second, the response to the weather limits is strongly nonlinear and asymmetric, as the full sweeps of Figure 36 show. Tightening both wave limits to 0.6 m (multiplier 0.3) collapses availability from 97.99 % to 63.8 % and multiplies the direct cost by 13, to 7.69 MUSD/yr, whereas relaxing them to 4.0 m recovers only +0.45 pp and −7.3 % cost. Practically all of the curvature therefore lies on the tightening side, and it originates in the accessibility-to-stretch mapping of the weather walks rather than in the Markov chain itself. The per-vessel split is equally clean: at multiplier 0.3 the CTV wave limit alone drives availability down to 67.0 % while the cost rises only to 763.8 kUSD/yr, whereas the HLV wave limit alone leaves availability at 91.3 % but drives the cost to 10.5 MUSD/yr. The CTV gates the frequent minor repairs and therefore owns availability, while the HLV day rate bills the weather-stretched major interventions and therefore owns cost.

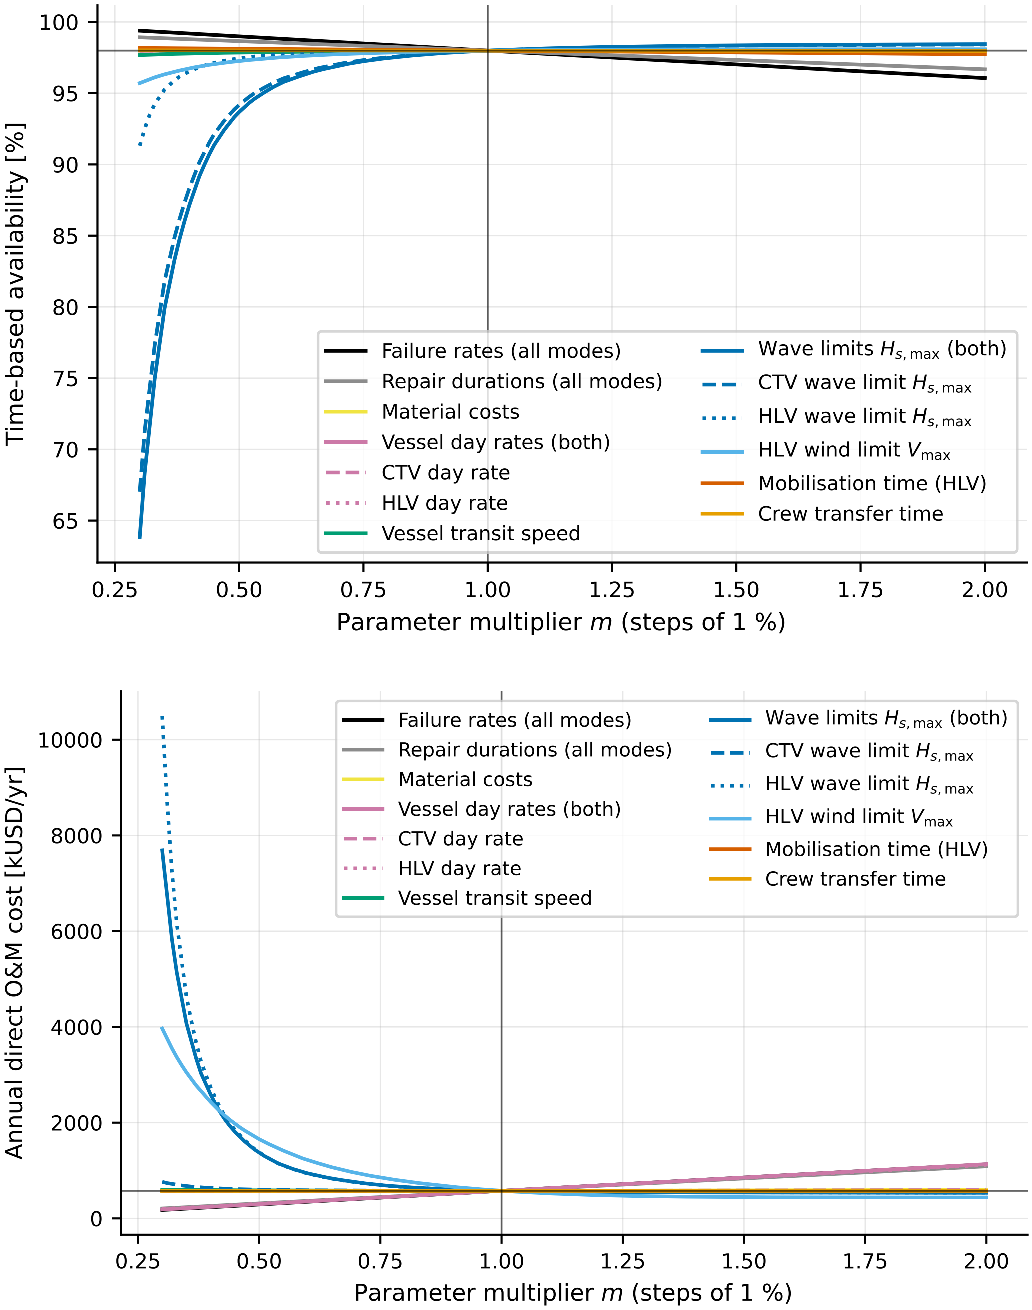

Figure 36: One-at-a-time sweeps of the thirteen parameter groups over multipliers 0.30–2.00 in 1 % steps: time-based availability (top) and annual direct cost (bottom).

Third, at the baseline the HLV wind limit is the locally strongest cost lever in the entire system, with $\varepsilon_C = -1.21$, super-elastic and larger in magnitude than the failure rates (+0.98) or the day rates (+0.97). This single number motivates the lifting-envelope study of Subsection 5.3.4. Related to it is the mobilisation paradox: lengthening the HLV mobilisation lowers availability and simultaneously lowers the direct cost, both with elasticity

-0.00266, because mobilisation hours are unbilled and the completed-task rate $f_i = A\lambda_i$ falls with availability. Direct cost is therefore demonstrably the wrong objective for mobilisation decisions, a point developed quantitatively in Subsection 5.3.5.

Fourth, the gap between time-based and production-based availability is itself weather-sensitive and, over the widened range, reveals a structure that the earlier $\pm 50\,\%$ window had hidden: it is non-monotonic (Figure 37). From 0.034 pp at wave multiplier 2.0 it grows through the baseline value of 0.166 pp to a peak of 0.70 pp near multiplier 0.36, where $H_{s,\max} \approx 0.72$ m and $A \approx 81.6\,\%$, before falling again towards multiplier 0.3. The mechanism seems clear on reflection. Moderate weather gating concentrates downtime in storm periods, during which the lost power exceeds its unconditional mean, whereas extreme inaccessibility stretches single repairs across whole seasons, so the power averaged over the span reverts towards the site mean and the covariance dilutes. Production-based availability is therefore most distinct from time-based availability precisely in the mid-severity regime in which most real sites operate, which is arguably the strongest argument this thesis offers for carrying the energy reward through the semi-Markov chain rather than reweighting time-based results after the fact.

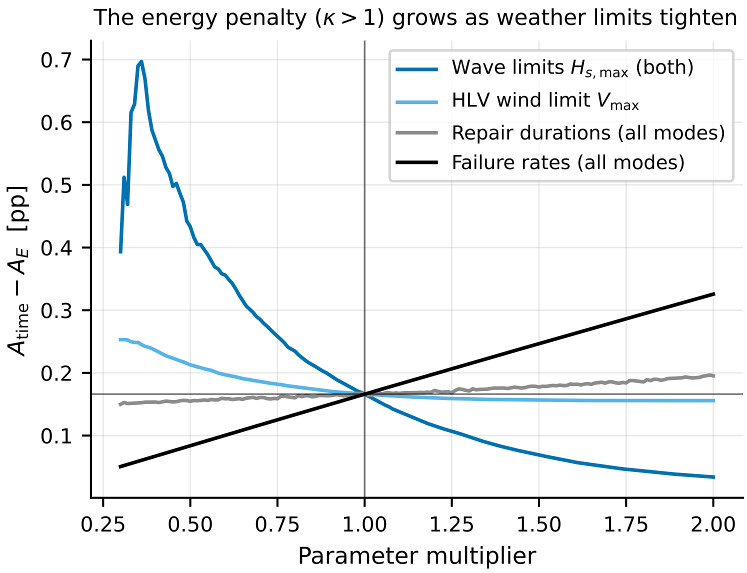

Figure 37: The time-based versus production-based availability gap $A - A_E$ along four sweeps. Along the wave-limit axis the gap is non-monotonic, peaking at 0.70 pp near multiplier 0.36 ($H_{s,\max} \approx 0.72$ m).

The two-parameter grid of Figure 38 completes the local picture by mapping the joint response of failure rates and wave limits. Over most of the plane the two effects compose almost separably, with the cost contours running nearly parallel to the wave axis under benign weather, where the failure rates decide, and bending sharply once the wave multiplier falls below roughly 0.6, where accessibility collapses and the weather term overwhelms everything else. It can also be observed that the absolute value of a reliability improvement grows steeply as the weather window tightens, since each avoided task then also avoids a heavily stretched, fully billed excursion, whereas under benign weather the same relative improvement is worth far less in absolute terms. Site severity and component reliability are therefore not interchangeable levers, and their joint treatment is likely to matter for any siting or procurement decision made near the steep region of the map.

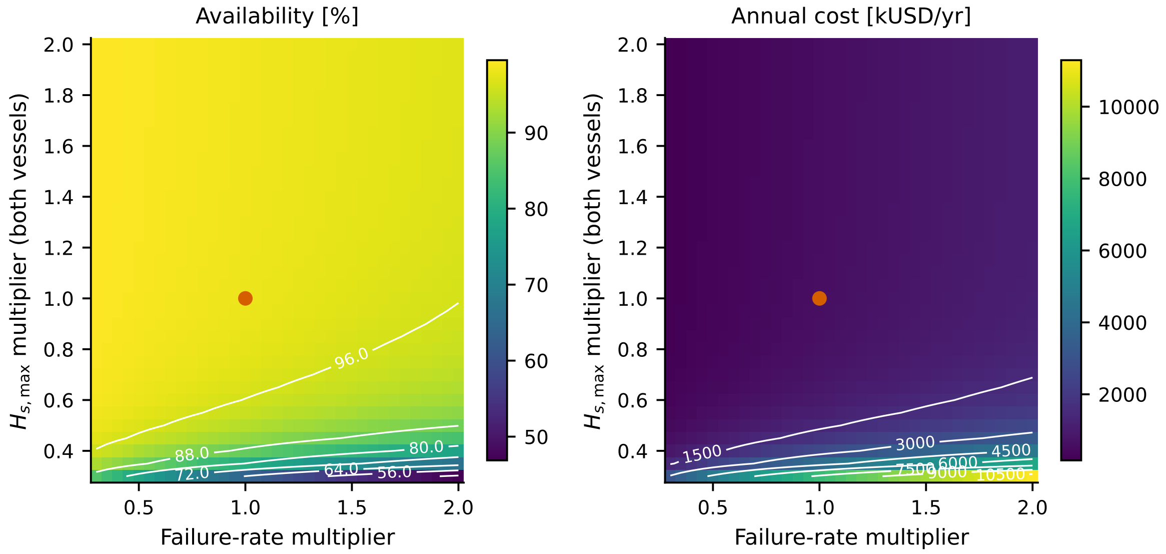

Figure 38: Two-parameter interaction of failure rates and wave limits (35 × 35 grid, 5 % steps): availability (left) and annual direct cost (right). The red marker is the baseline.

Takeaway. The reference system answers to nearly disjoint lever sets: weather limits own availability, day rates own cost, and only reliability and repair durations move both. For design and monitoring priority, the HLV wind limit is the single strongest cost lever in the system (super-elastic at −1.21), failure rates dominate both KPIs at near-unit cost elasticity, and the weather response is violently asymmetric, with the tightening side collapsing availability to 63.8% and multiplying cost thirteenfold at multiplier 0.3. Direct cost is demonstrably the wrong objective for mobilisation decisions, and the availability-energy gap peaks at 0.70 pp exactly in the mid-severity regime where real sites operate, the strongest case for carrying the energy reward through the chain.

5.3.2 Propagation of input uncertainty

Elasticities describe the response to one input at a time, yet the inputs of an O&M model are uncertain jointly, and the literature is explicit both about the size of that uncertainty and about its consequences: failure rates are the dominant uncertain input [114], the population statistics behind them mix turbine generations and sites [7], and charter rates vary with market conditions and season [116]. This subsection therefore propagates a registered joint input distribution through the model and examines the full output distributions, something a discrete-event simulator can approach only at extraordinary expense.

The registered input model is as follows. Each of the 22 per-mode failure-rate multipliers is lognormal with median 1 and a 90 % interval of $[0.5, 2.0]$, that is $\sigma_{\ln} = 0.4214$, a width consistent with the spread reported across populations and configurations by Carroll et al. [7]. Each per-mode active-repair duration multiplier is lognormal with a 90 % interval of $[0.7, 1.43]$. The two day-rate multipliers are uniform on $[0.8, 1.2]$ and the material multiplier uniform on $[0.85, 1.15]$. Because the median multiplier is 1, the mean is $\exp(\sigma_{\ln}^2/2) = 1.093$, therefore the input model itself shifts the expected task arrival rate upwards by 9.3 %, and this must be kept in mind when reading the output means. Two dependence cases are run with identical marginals: fully independent failure-rate multipliers, and comonotone multipliers driven by a single shared factor. The truth for a real fleet presumably lies between the two, since systematic effects such as serial design defects or site severity act fleet-wide while unit-to-unit variation acts per mode, and running both extremes brackets the answer.

The input model is a choice, and its influence deserves stating. The lognormal family for rates and durations encodes multiplicative, right-skewed uncertainty, which is what population spreads across turbine generations and sites look like, and the registered widths tie the 90% intervals to the spread reported by Carroll et al. rather than to convenience. Three consequences follow for reading the outputs. Firstly, the output means inherit the +9.3% mean shift of a median-one lognormal, which is a property of the input model and not of the system. Secondly, the output spreads scale approximately with the input widths, since the propagated response is close to linear over the sampled range, so halving the rate interval roughly halves the availability spread while leaving the dependence conclusion intact. Thirdly, the output skew is inherited from the lognormal choice, and a symmetric input model would understate the upper cost tail that drives contingency. The two dependence cases bracket the unknown correlation structure by construction, and because one full propagation costs under a second, re-running the study under alternative distributional assumptions is a routine robustness exercise rather than a campaign, which is itself part of the RQ4 answer.

Propagation uses an exact per-mode reconstruction of the closed-form solve in which the two weather walks, E[span] and E[L_R] as functions of the banked active-hours target, are replaced by monotone splines fitted per mode on 49 exactly solved nodes over duration multipliers 0.25–2.20. At the nodes the reconstruction matches the full model to machine precision ($|\Delta A| \sim 10^{-14}$ pp), so the only approximation is spline interpolation between nodes, and a validation against the full model on 30 random joint samples bounds that error at 0.0061 pp in A, 0.010 pp in A_E and 0.20 % in cost, negligible against the output spreads reported below. With the surrogate in place, $2 \times 20\,000$ joint samples propagate in 0.9 s. The equivalent simulation campaign, at roughly 200 seeds per sample, would be of the order of 8×10^6 lifetime simulations, which is why input uncertainty quantification of this kind is essentially never attempted with discrete-event O&M models. Table 25 summarises the propagated distributions, from which three findings stand out, each illustrated by the figure accompanying it.

Table 25: Propagated output distributions ($N = 20\,000$ joint samples per dependence case). Costs in kUSD/yr, TC at 100 USD/MWh.

Case	KPI	Mean	SD	P5	P50	P95
independent	A [%]	97.77	0.35	97.15	97.81	98.28
independent	C [kUSD/yr]	640.0	234.9	351.1	594.5	1 078.1
independent	TC(100)	776.5	248.7	468.9	729.1	1 238.0
comonotone	A [%]	97.78	0.95	95.98	97.95	98.96
comonotone	C [kUSD/yr]	641.5	300.2	277.9	581.1	1 214.7
comonotone	TC(100)	777.9	354.5	344.5	708.5	1 449.6

First, the cost distribution is strongly right-skewed (Figure 39): in the independent case the upper tail distance P95 − P50 of the total cost is 508.9 kUSD against a lower distance P50 − P5 of 260.2 kUSD, a ratio of almost two. The first-order Gaussian constructed from the analytic elasticities predicts an availability standard deviation of 0.29 pp against the observed 0.35 pp and, by construction, none of the skew, therefore the delta method is adequate for quick availability error bars but materially understates the upside cost risk. An operator budgeting the P95 rather than the mean would underestimate the required contingency by roughly a third if a linearised propagation were used.

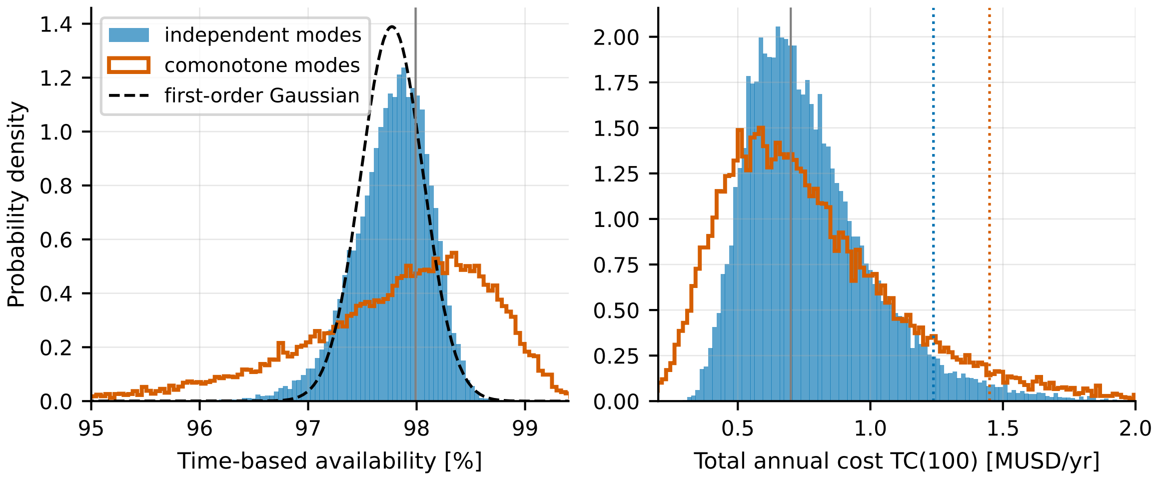

Figure 39: Propagated distributions of availability (left) and total annual cost at 100 USD/MWh (right) under independent and comonotone failure-rate uncertainty. The dashed curve is the first-order Gaussian implied by the analytic failure-rate elasticities, and the dotted vertical lines mark the two P95 values.

Second, and perhaps most importantly, the dependence structure of the failure-rate uncertainty matters more for risk than the marginals themselves, as the distribution and exceedance curves of Figure 40 show. Moving from independent to comonotone multipliers leaves the means essentially unchanged but multiplies the availability standard deviation by 2.7, from 0.35 pp to 0.95 pp, because independent per-mode errors partially cancel inside

$\sum_i \lambda_i D_i$ while a shared factor does not. The probability of falling below 97 % availability rises from 2.7 % to 17.9 %, almost sevenfold, and the probability that the total annual cost exceeds 1 MUSD rises from 15.9 % to 21.5 %. Reliability databases report marginal uncertainty but rarely its correlation structure, and these results suggest that, for availability risk, knowing the correlation would be worth more than narrowing the marginals.

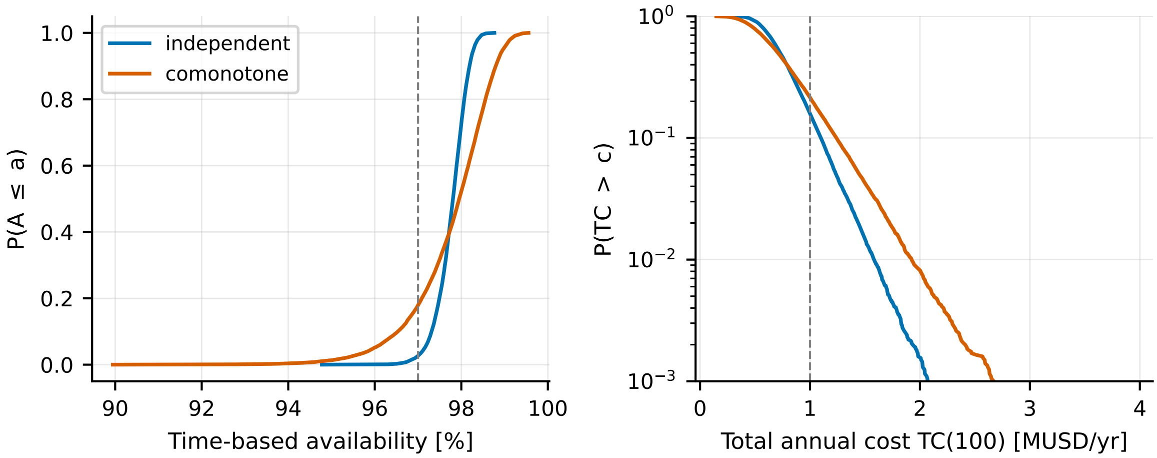

Figure 40: Distribution functions of the propagated availability (left) and exceedance probabilities of the total annual cost (right, logarithmic).

Third, the variance decomposition of Figure 41 confirms and quantifies the expectation from the literature: failure-rate uncertainty contributes 86 % of the availability variance and 68–70 % of the cost variance in the independent case, repair durations contribute 8 % and 13–14 % respectively, day rates contribute nothing to availability and 6–7 % to cost, and materials are negligible throughout. Efforts to reduce input uncertainty are therefore best spent on the failure rates, which is consistent with the conclusions of Martin et al. [114] and Dinwoodie et al. [8], here obtained with exact propagation rather than scenario comparison.

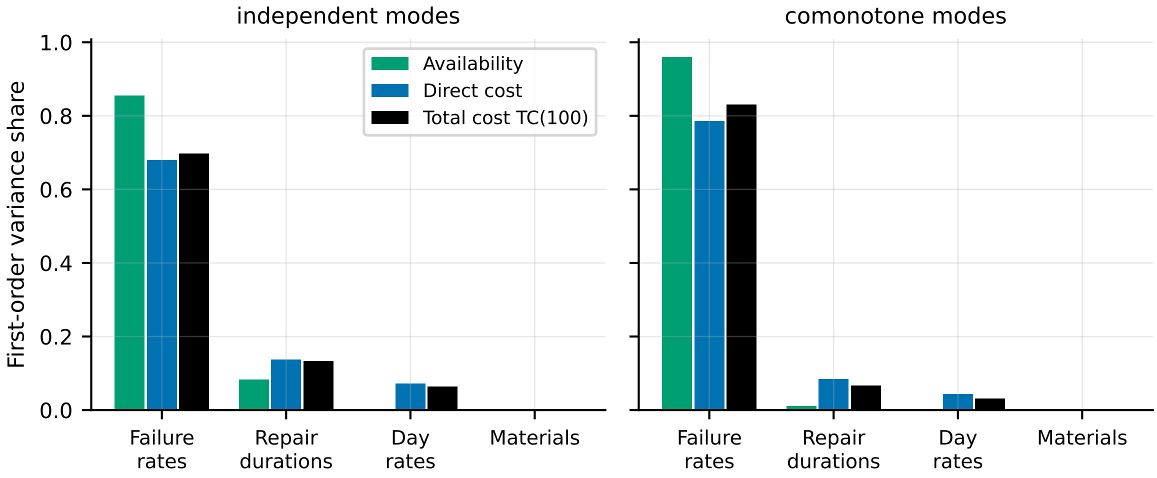

Figure 41: First-order variance shares by input group, obtained by propagating one group at a time with the others frozen at their medians.

Takeaway. Three results for the risk picture: the cost distribution is strongly right-skewed, so linearised propagation understates the P95 contingency by roughly a third, the dependence structure of failure-rate uncertainty matters more than the marginals, multiplying the availability spread by 2.7 and the probability of sub-97% availability sevenfold at identical marginals, and 86% of the availability variance traces to the failure rates. For an operator this reorders the data agenda, as learning the correlation structure of failures is worth more than narrowing their individual estimates, a quantity reliability databases do not currently report.

The same machinery can answer procurement questions directly, since a design alternative is nothing more than a point, or a curve, in the space that the model evaluates in milliseconds. Four studies follow, each built around a decision an operator genuinely faces, and each resolved to an explicit recommendation under the stated conventions.

5.3.3 Access capability versus charter rate: CTV selection

Crew transfer vessels span a well-documented capability ladder: conventional catamarans transfer safely up to roughly 1.5–2.0 m significant wave height, while surface effect ships and similar advanced designs offer 2.0–2.5 m at a higher charter cost, and the selection between them has been framed in the literature as an explicit cost-versus-capability decision [117, 44]. Three candidates are therefore defined around the reference vessel. Candidate one is a conventional catamaran with a 1.5 m limit at 0.80 times the reference day rate (2561 USD/day), candidate two is the reference vessel itself with 2.0 m at 3201.7 USD/day, and candidate three is a surface-effect-ship class vessel with 2.5 m at 1.50 times the rate (4803 USD/day). Table 26 reports the outcome at the central price.

Table 26: CTV candidates at $p = 100$ USD/MWh. Direct cost, lost revenue and total cost in kUSD/yr. The direct cost barely distinguishes the vessels, so the lost-production channel decides.

Candidate	A [%]	A_E [%]	Direct	Lost rev.	TC(100)
Catamaran 1.5 m, 0.80×	97.32	97.07	575.7	166.0	741.6
Reference 2.0 m, 1.00×	97.99	97.82	576.4	123.3	699.7
SES 2.5 m, 1.50×	98.23	98.12	581.6	106.5	688.0

The three candidates are practically indistinguishable in direct cost, spanning barely 6 kUSD/yr, because the CTV carries only 2.8 % of the direct-cost elasticity, yet they are separated by 0.91 pp of availability and therefore by 59.5 kUSD/yr of lost revenue at the central price. The whole decision consequently lives in the production channel, which a direct-cost comparison would miss entirely. Although the catamaran is the cheapest vessel to charter, the SES yields the lowest total cost at all three prices, and the price sweep of Figure 42 makes the margins explicit: the SES beats the reference vessel for any price above 30.7 USD/MWh and beats the catamaran above 9.9 USD/MWh, both far below any plausible offshore revenue.

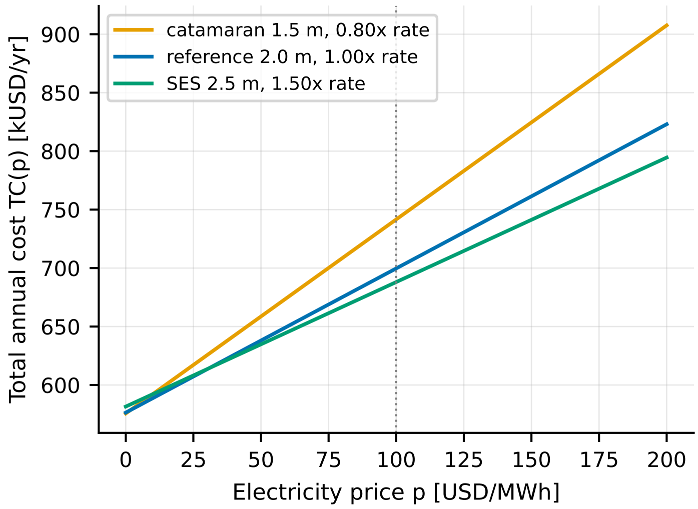

Figure 42: Candidate total cost as a function of the electricity price. The crossings define break-even prices of 9.9 and 30.7 USD/MWh.

The break-even frontier of Figure 43 generalises the comparison from three candidates to the whole capability–rate plane, and at the central price it shows that a 2.5 m vessel would remain competitive up to 7627 USD/day, 2.38 times the reference rate, whereas the 1.5 m catamaran could not match the total cost of the reference vessel even at a day rate of zero. Therefore, as the higher-capability vessel recovers its premium many times over through avoided lost production, the SES-class candidate is the preferred charter under the stated conventions, and the conclusion is robust across the whole price range considered. This is consistent with the observation of Sperstad et al. [44] that fleet rankings hinge on the access limit, and it quantifies exactly why.

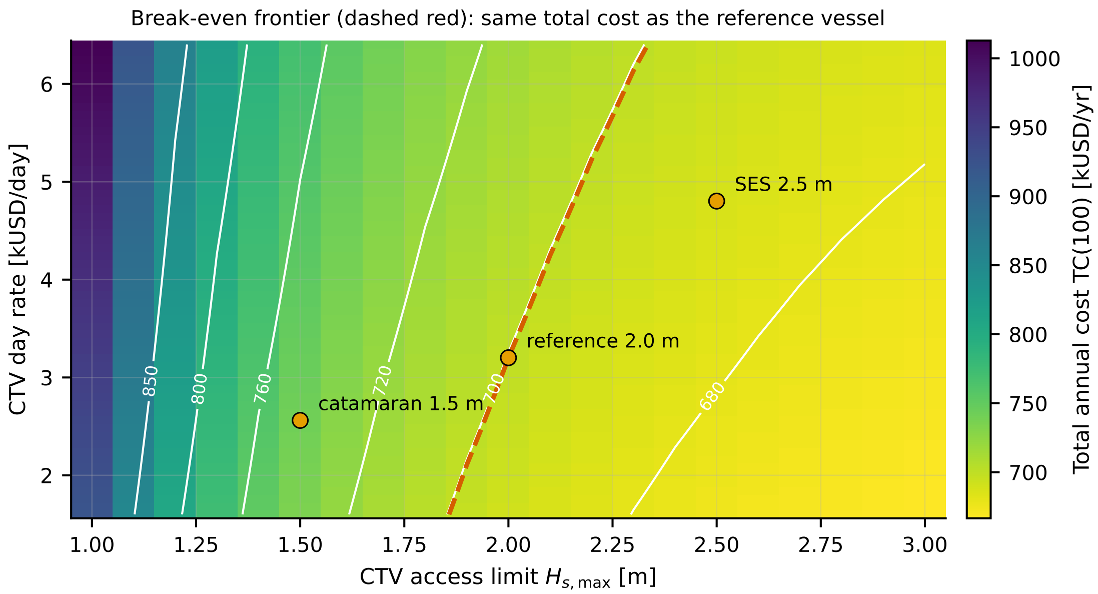

Figure 43: Total annual cost at 100 USD/MWh over the CTV capability–rate plane. The dashed red frontier marks cost parity with the reference vessel. The markers are the three candidates.

Takeaway. CTV selection is decided in the production channel and not the charter invoice. The three candidates differ by 6 kUSD per year of direct cost but by 59.5 kUSD per year of lost revenue. The higher-capability vessel wins at any electricity price above 30.7 USD/MWh, remains competitive to 2.38 times the reference day rate, and the low-capability catamaran cannot reach cost parity even chartered for free. Access capability buys availability where day rate buys almost nothing, so procurement comparisons made on direct cost alone select the wrong vessel and motivate the use of the transient model.

5.3.4 The lifting envelope: HLV wind limit

Major replacements require crane operations whose limiting factor is typically the wind speed at hub height, and a considerable engineering literature exists on widening that envelope through single-blade installation tooling, tugger-line damping and related devices [112, 118]. The super-elastic cost response found in Subsection 5.3.1 ($\varepsilon_C = -1.21$) suggests the value of such technology is large, and the model turns the suggestion into a number. Figure 44 sweeps the HLV lifting wind limit from 8 to 14 m/s around the 10 m/s baseline.

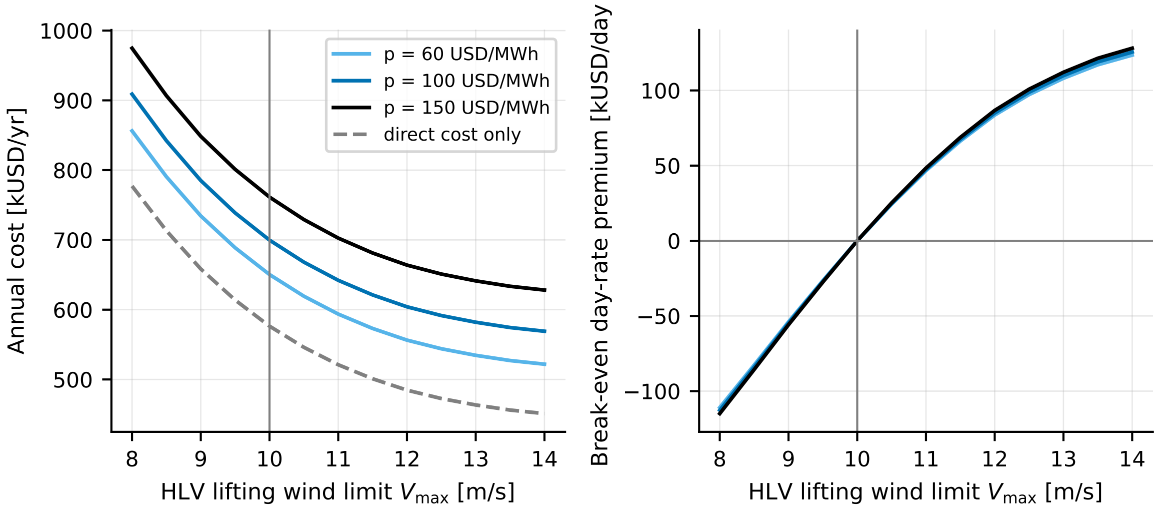

Figure 44: Value of the HLV lifting envelope. Left: total annual cost against the wind limit at three prices, with the direct-cost-only curve dashed. Right: break-even day-rate premium relative to the 400 kUSD/day baseline vessel.

Raising the limit from 10 to 12 m/s lowers the direct cost from 576.4 to 484.7 kUSD/yr and the total cost at the central price by 95 565 USD/yr per turbine, a 13.7 % reduction, while availability gains a modest 0.062 pp. Expressed as a procurement rule, a vessel or lifting system certified to 12 m/s is worth a day-rate premium of up to 84 987 USD/day, 21 % above the 400 kUSD/day baseline, before the advantage is consumed, and the capitalised value of the upgrade is 1.10 MUSD per turbine over the 20-year life. The curve also shows clearly diminishing returns beyond roughly 13 m/s, where the remaining weather-stretch on the lifting spans is small, so the economic case supports incremental envelope improvements near the current limit far more strongly than pursuit of an unrestricted envelope. It is worth noting that, unlike the CTV study, this value sits almost entirely in the direct-cost channel, since it is the billed weather-stretched span of the 400 kUSD/day vessel that shrinks.

Takeaway. The lifting envelope is where technology buys cost directly, as widening the HLV wind limit from 10 to 12 m/s saves 95,565 USD per turbine-year at the central price, justifies a day-rate premium of up to 21%, and capitalises to 1.10 MUSD per turbine, with clearly diminishing returns beyond 13 m/s. Unlike the CTV case the value sits in the direct-cost channel, so single-blade tooling and damping systems that extend lifting limits pay for themselves in billed vessel hours, and incremental envelope gains near the current limit dominate the pursuit of unrestricted capability.

5.3.5 Mobilisation time and charter strategy

The mobilisation paradox of Subsection 5.3.1 showed that the direct cost rewards longer mobilisation, which is plainly the wrong signal, and the literature confirms both that mobilisation lead time significantly affects availability [119] and that charter strategy, from spot chartering with long lead times to framework or long-term arrangements with short ones, is a genuine strategic choice with material cost implications [120]. Figure 45 resolves the paradox by pricing the production channel.

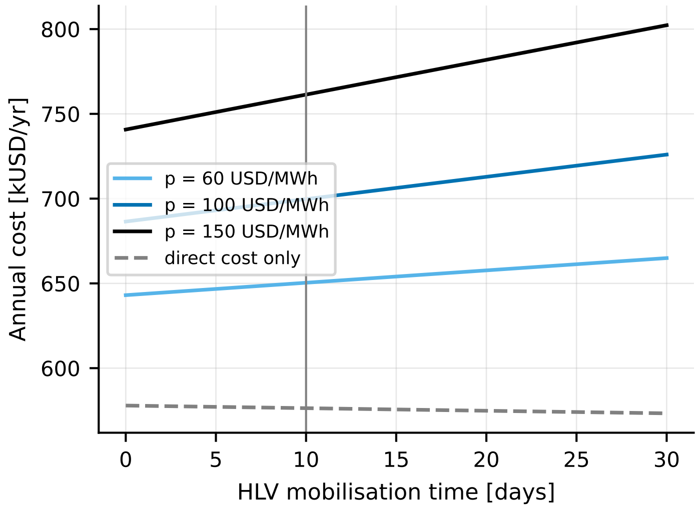

Figure 45: Total annual cost against HLV mobilisation time at three prices. The dashed direct-cost curve decreases with mobilisation time, the total cost increases at any positive price.

Across mobilisation times from zero to 30 days the direct cost falls from 577.9 to 573.3 kUSD/yr while the total cost at the central price rises monotonically from 686.5 to 726.0 kUSD/yr, therefore any positive electricity price restores the correct preference for shorter mobilisation. The local willingness to pay around the 10-day baseline is 1320 USD/yr per turbine per mobilisation day at 100 USD/MWh, scaling approximately in proportion to the price, which with 0.0971 HLV interventions per turbine-year corresponds to roughly 13.6 kUSD per saved mobilisation day per actual intervention. Whether a framework charter that guarantees, say, a ten-day-shorter response is worthwhile then reduces to comparing its annual fixed premium per turbine against 13200 USD/yr, a comparison the operator can perform directly with Figure 45. For a single turbine the number is modest, and this is one of the places where the single-turbine scope of the reference case must be respected: an HLV arrangement is shared across a farm, so the farm-level value scales with the number of turbines while the premium does not, and the strategic conclusion is likely to be considerably stronger at farm scale (Subsection 5.3.7).

Takeaway. Direct cost rewards longer mobilisation, which is the wrong signal, and any positive electricity price reverses it. A shorter response is worth about 1,320 USD per turbine-year per mobilisation day at the central price, roughly 13.6 kUSD per saved day per actual heavy intervention. Per turbine the number is modest, but an HLV arrangement is shared across a farm while its premium is not, so the framework-charter case scales with turbine count and the per-turbine figure is a lower bound on the farm-level value.

5.3.6 Reliability investment: where a design margin pays

The final study asks the reliability question in procurement form: how much additional capital expenditure does a 10 % reduction of a single failure rate mode justify? This is precisely the direction of current industry interest, where the required reliability of next-generation turbines is an open design target [121] and drive-train configuration choices are argued through their O&M consequences [51]. For each of the 22 modes the model is solved with the rate of the mode multiplied by 0.90, the total-cost saving is computed at the three prices, and the saving is capitalised with the 20-year annuity. Because the per-mode response is nearly linear for small reductions, the resulting figures scale approximately linearly in the reduction percentage.

The concentration is extreme (Figure 46). A 10 % rate reduction of the power converter replacement mode alone is worth 43121 USD/yr, that is a justified capex of 494592 USD per turbine, while the same relative improvement applied to every one of the 22 modes simultaneously is worth 68588 USD/yr, so the converter carries 63% of the total reliability value by itself. The next candidates follow at an order of magnitude less: the direct-drive generator major replacement at 82885 USD, the main shaft replacement at 51649 USD and the blades major replacement at 30729 USD of justified capex.

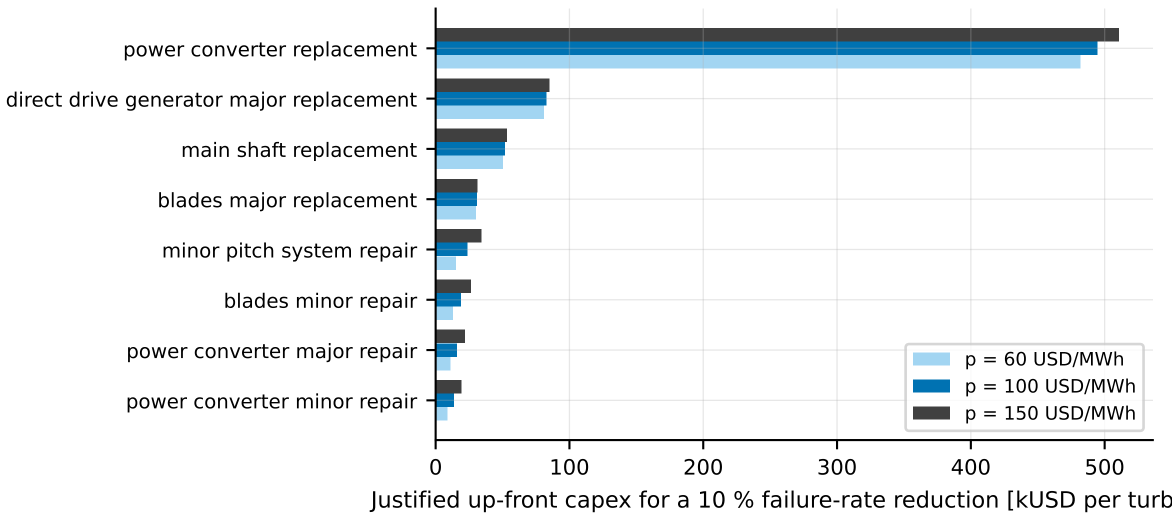

Figure 46: Justified up-front capital expenditure per turbine for a 10 % failure-rate reduction of each mode (top eight shown), capitalised at 6 % over 20 years.

The design guidance is therefore unambiguous: reliability margin, redundancy or condition monitoring spend should be targeted at the power converter before anywhere else, and a manufacturer able to demonstrate even a modest converter reliability improvement can defend a per-turbine price premium approaching half a million dollars under the stated conventions. Although these figures inherit the failure-rate uncertainty quantified in Subsection 5.3.2, the ranking itself is driven by the cost structure of the modes rather than by the uncertain rates.

These figures rest on stated economic conventions, and their dependence on them is worth making explicit. The valuation uses fixed electricity prices, a 6% real discount rate over the twenty-year annuity, and no market coupling or curtailment, so the absolute levels move with the market while the structure does not. A higher electricity price scales the lost-production component approximately linearly, which is why the three-price bars widen together, a higher HLV day rate raises the value of every avoided replacement in proportion to its billed span, and a higher discount rate compresses all capitalised figures through the annuity factor, roughly fourteen per cent for two points of discount rate. Under any of these movements the ranking is preserved, because it is carried by the cost structure of the modes, the converter combining a meaningful rate with an expensive, weather-stretched intervention, and only the absolute size of the justified premium shifts. The one condition that could reorder the list is a change in the intervention definitions themselves, vessel class or repair content per mode, which is a design question the same machinery prices in milliseconds.

Takeaway. Reliability value is extremely concentrated, with a ten per cent rate reduction of the power converter replacement alone justifying 494,592 USD of capital expenditure per turbine, 63% of the value of improving all 22 modes at once, with the next candidates an order of magnitude behind. Design margin, redundancy, and condition-monitoring spend should target the converter first, and a manufacturer able to demonstrate the improvement can defend a per-turbine premium approaching half a million dollars under the stated conventions, a ranking robust to the input uncertainty because it is carried by the cost structure of the modes.

5.3.7 Limitations and scope of the exploitation studies

The limitations of this section fall into two classes, and separating them matters because they bound different things. The first class belongs to the analytical framework itself. The reference case is a single turbine served by dedicated vessels, so effects that are approximately linear in turbine count, the CTV selection and the reliability ranking, extrapolate naturally to farm scale, whereas shared resources do not: an HLV charter serves the whole farm, vessel queueing couples turbines, and the mobilisation values of Subsection 5.3.5 are per-turbine lower bounds on a farm-level value, with a farm formulation under vessel sharing beyond the scope of this thesis (A-9). The studies are steady-state, so questions with intrinsic calendar structure, the optimal start month of a maintenance campaign [55] or the seasonality of charter rates [116], belong to the transient TM-S formulation and are not addressed here. And the marginal interruptibility treatment of A-3 leaves the coupling between spans and the concurrent weather state at transition instants unresolved, bounded empirically by the summit residuals of Chapter 4.

The second class belongs to the case and its inputs rather than to the framework. All results inherit the $\beta = 1$ build, and by the invariance theorem the availability-side conclusions carry unchanged to the mixed-shape arms, although the timing distribution of costs within the lifetime would require the transient or sampler members. The maintenance strategy is held fixed across every sweep and the economic layer uses fixed prices with no market coupling, so the extreme points of Subsection 5.3.1 are stress tests of a frozen strategy rather than forecasts for an operator who would re-optimise long before availability fell towards 64%. And the input-distribution choices of Subsection 5.3.2 shape the propagated spreads in the ways stated there. On generalisability beyond this case, the transfer argument of Section 4.8 applies unchanged: the mechanism conclusions travel with the mathematics, the numbers are tied to this site, machine, and input baseline, and the method reruns on any configuration the pipeline accepts.

Within those bounds, the section has demonstrated the intended payload of RQ4. The certified analytical model prices vessel capability, lifting envelopes, mobilisation lead time and component reliability in seconds, it propagates realistic joint input uncertainty with exact machinery whose worst-case error is measured in fractions of a percentage point, and it exposes structure, from the super-elastic lifting envelope to the non-monotonic energy gap and the dominant role of failure-rate dependence, that would sit below the noise floor of the discrete-event reference at any practical seed budget.

5.4 Synthesis: the take-home messages

The chapter closes by stating the main four messages its analysis amounts to, each carried by results above and each feeding one verbatim answer in Chapter 6.

The memoryless simplification is adequate exactly as far as the means reach, and time is where it ends. In steady state the exponential family is protected by the invariance theorem, and every long-run conclusion of this chapter holds for it identically. In time it first misprices, overstating the cold-start settling by sixty per cent and the boundary credit by 46% even at $\beta = 1$, and then goes structurally silent, since at $\beta = 2$ a twenty-year life is one long, validated transient of replacement waves that the exponential description cannot represent at any accuracy. Calendar-indexed commitments on ageing fleets require the semi-Markov transient as a matter of representability, not precision (RQ2).

Interruptibility is a declaration-worthy assumption, two-regime and priced. On short repairs the execution choice is a second-order convention, at the reference site 0.167 pp and roughly 11,300 USD per turbine-year. On heavy interventions it decides whether the model describes the system at all, with stretch factors leaving physical plausibility and the largest job infeasible under whole-window execution, unrescuable by wave-limit relaxation because the wind limit caps the window supply. An O&M model should state its execution treatment as explicitly as its failure rates (RQ3).

The analytical route converts fidelity into reach. With millisecond solves and exact derivatives, the chapter ran studies that are structurally unavailable to the reference at any practical budget: exact elasticities reproducing their closed forms to 4×10^{-8}, forty thousand joint uncertainty samples in under a second against an equivalent of order 8×10^{6} lifetime simulations, and effects far below the ensemble noise floor, the super-elastic lifting envelope, the sevenfold dependence-driven risk shift, the non-monotonic 0.70 pp energy gap. The procurement studies returned decisions with break-evens attached, and the machinery reruns under altered assumptions as a routine robustness check (RQ4).

The production-based estimand carries money, and only the semi-Markov family carries it. The $A - A_E$ gap prices lost revenue correctly where the time-based figure understates it by construction, it peaks in exactly the mid-severity regime where real sites operate, and it is available along the whole trajectory rather than as an afterthought correction, which the vessel-selection study converts directly into a reversed procurement decision (RQ2, RQ4).

Together with the fidelity established in Chapter 4, these are the results the thesis stands on, and Chapter 6 answers the research questions in their terms.

6 Conclusions

This thesis set out to close a specific gap: no analytical offshore wind availability model handled non-exponential durations, interruptible repairs, or time-resolved solutions, and none had been validated against a modern discrete-event reference at mechanism level under aligned conventions. The chapters built the four-model family that occupies that gap, certified it against an audited reference, and exercised it on the engineering questions the speed makes reachable. This chapter answers the four research questions verbatim, states the contributions in two registers, discusses the limitations and makes suggestions for future work, and closes with the messages the thesis asks the reader to retain.

6.1 Answers to the research questions

RQ1 (fidelity). Can an analytical availability model, built on strictly aligned inputs, reproduce the availability and cost estimates of a certified DES reference within the statistical resolution of the reference itself?

Yes, but only when the models are aligned deliberately rather than assumed to be comparable. Chapter 4 defines the conditions required for that alignment. The reference model is audited at both source and event level, the patch series is verified by signature and separates defect correction from the declared change in clock convention, the input baseline is frozen and checked for parity, and the validation campaign is preregistered so that its predictions and acceptance criteria are fixed before the results are generated. With these controls in place, the steady-state semi-Markov model reproduces the verification twin closely. For the first 200-seed batch, the availability difference corresponds to $z = +0.20$, while comparison with the pooled 1000-seed population gives $z = -1.02$. The replication batch also satisfies the sealed criteria using seeds that had not yet been generated when the campaign was registered. The remaining discrepancy can be traced to one identified source, namely the sampled count of a rare failure mode, rather than to any persistent structural difference between the models. The mechanism-accretion ladder reaches the same conclusion incrementally. Agreement is established first through closed-form identities and then through the weather operator, which is validated event by event to the resolution imposed by hourly rounding. The transient formulation is checked at both ends of the shape-parameter range. It recovers the steady-state result at $\beta = 1$ and matches the cold-start ageing batch x19 at $\beta = 2$ (Subsection 4.6.3).

RQ2 (exponential adequacy boundary). For which output parameters and operating regimes does the memoryless simplification remain adequate, and where does the semi-Markov description become necessary?

The simplification is adequate for long-run quantities determined by mean durations, but a semi-Markov treatment is required once timing or energy production becomes relevant. The appropriate conclusion is therefore conditional rather than absolute. As shown by the invariance theorem in Section 3.7, steady-state availability, state occupancies, task frequencies, and costs depend on the duration distributions only through their means. A mean-matched exponential model therefore reproduces these long-run quantities exactly, regardless of the shape parameter. This result follows analytically from the proof in Section A.3; it is not merely an empirical observation.

The limits of the simplification appear in two areas, both quantified in Chapter 5. The first concerns transient behaviour. At $\beta = 1$, the exponential and semi-Markov models have the same steady state, but they do not approach it in the same way. The exponential model overestimates the cold-start settling time by 60% and the boundary credit by 46%, with both discrepancies explained by its fixed second cycle moment. For $\beta \neq 1$, the exponential trajectory remains unchanged even though the underlying system behaviour varies. At $\beta = 2$, it underestimates year-one availability by 1.47 percentage points. The validated replacement waves also persist throughout the twenty-year analysis horizon, which means that the system does not reach steady state within a typical design life. Any quantity tied to calendar time, such as first-year availability, warranty-period energy production, or the timing of expenditure over a debt period, therefore requires the transient semi-Markov formulation because the exponential model cannot represent the relevant dynamics.

The second limitation concerns the quantity being estimated. Production-based availability cannot be represented by the exponential model because its memoryless sojourn times contain no information about when downtime occurs in the calendar year, as expressed in Equation 10. In the reference case, production-based availability is 0.1660 percentage points lower than time-based availability. This difference is sensitive to weather conditions and reaches 0.70 percentage points in the mid-severity range most relevant to operating sites. Wherever lost energy is assigned a monetary value, this distinction has a direct economic effect.

RQ3 (interruptibility). How large is the pause-and-resume effect of weather-interrupted repairs, and does a semi-Markov model carrying interruption-shaped durations capture it?

The importance of interruptibility depends strongly on task duration. Its effect is relatively small for short repairs but becomes decisive for heavy interventions. The key quantity is the relationship between the active duration required by a task and the distribution of uninterrupted accessible periods in the corresponding weather mask. When the task duration is shorter than a typical accessible period, banked-progress and whole-window execution give similar results. For the sixteen CTV failure modes at the reference site, the whole-window assumption reduces availability by 0.167 percentage points and increases annual cost by approximately 11,300 USD per turbine. Although modest under the reference conditions, this difference grows by a factor of 230 as the wave limits become more restrictive across the severity sweep.

For heavy interventions, the whole-window assumption no longer provides a physically credible representation. The resulting stretch factors range from 25 to 143, compared with values below two when completed work is retained between weather interruptions. For the longest task, contiguous execution is not merely conservative but infeasible, since the twenty-year weather record contains no accessible period long enough to complete it. Relaxing the wave-height limit does not remove this problem because the wind-speed constraint still limits uninterrupted accessibility to approximately two weeks.

The semi-Markov formulation captures this behaviour directly because its duration distributions are derived from the recorded accessibility masks and include interruptions by construction. These span distributions were validated event by event across 1,701 weather interruptions in the mechanism-accretion ladder and against the unprocessed reference outputs at the final stage. The implication for O&M modelling is therefore conditional. Whole-window methods remain appropriate for genuinely uninterruptible tasks of moderate duration, but they can overstate heavy-intervention durations by orders of magnitude when progress can be retained. The assumed execution rule should consequently be reported as explicitly as the failure rates, since for the longest maintenance tasks it can have a greater effect on the results than any other modelling assumption.

RQ4 (computational reach and model applicability). What can the analytical route resolve and afford that the simulation route cannot, and what engineering studies does that reach enable?

The analytical approach reveals effects that simulation cannot resolve economically, and Chapter 5 demonstrates this directly. A certified steady-state solution requires about 13 milliseconds, whereas the corresponding reference analysis relies on batches of full-lifetime simulations to obtain converged estimates. At a matched tolerance, this represents a computational speedup of approximately ninety times. The closed-form formulation also provides exact derivatives, and the six registered elasticities matched their analytical values to within 4×10^{-8}.

This computational advantage was used in three main ways. First, it made it possible to resolve effects that would lie below the sampling variability of any practical simulation ensemble. These included the super-elastic cost response of -1.21 to the lifting envelope, the effect of failure-rate dependence in increasing the availability spread by a factor of 2.7 and multiplying the probability of sub-97% availability by seven under identical marginals, and the non-monotonic behaviour of the energy gap. Second, it enabled large-scale uncertainty propagation. Forty thousand joint input samples were evaluated in under one second, while an equivalent simulation study would require on the order of 8×10^6 lifetime simulations. The results also showed that the delta method understated the required cost contingency by roughly one third.

Finally, the analytical model supported procurement decisions with explicit break-even values. A higher-capability transfer vessel became economically justified above 30.7 USD/MWh. Expanding the lifting envelope supported a 21% day-rate premium and a capitalised value of 1.10 MUSD per turbine. Reducing mobilisation lead time was valued at 13.6 kUSD per saved day for each heavy intervention, while a 10% improvement in power-converter reliability justified approximately 0.5 MUSD of additional capital expenditure and accounted for 63% of the total reliability value in the system. Because the full analysis can be repeated within seconds, sensitivity to alternative assumptions can be examined routinely at a level of detail that would be impractical in a simulation-based workflow.

6.2 Contributions

The contributions of this thesis fall into two groups. The first concerns the scientific methods and findings established by the work, while the second concerns their practical value for offshore wind O&M decision-making.

Scientific contributions. S1, a validated model hierarchy. The thesis develops six solvers arranged in a two-by-three structure. Their internal relationships, including analytical equivalence, nesting of the exponential case, convergence from transient to steady state, and agreement between sampling methods, are tested explicitly rather than assumed. This allows the computationally cheaper models to inherit credibility from the more detailed ones, while any remaining differences are traced to clearly defined modelling conventions. **S2, a controlled characterisation of the reference model.** The work establishes a frozen and signature-verified patch series for a widely used open-source simulator. The build sequence separates defect correction from the deliberate change in clock convention, making it possible to audit exactly which version was compared at each stage. The paired deferral census in Chapter 4 also quantifies the effect of using the published tool without modification, providing a result that is directly relevant to future studies using the same reference. **S3, a formal result on the estimand.** Production-based availability is defined as a quantity that cannot be represented by the exponential model family, as shown in Equation 10. This distinction is carried consistently through both the steady-state and transient formulations. The supporting results, including the invariance theorem and the exactness of the weather operator, are proved in Appendix A.

Engineering contributions. E1, the importance of declaring execution assumptions. The treatment of interruptibility is shown to operate in two regimes. Its effect is small for short repairs but becomes dominant for heavy interventions. This demonstrates that execution assumptions should be reported as explicitly as failure rates, because they can materially change the results. **E2, a practical decision-support instrument.** The thesis shows how a validated analytical model, exact surrogates and dense parameter sweeps can turn availability analysis into direct procurement support. The framework is used to evaluate vessel selection, lifting-envelope value, charter strategy and component reliability, with each result expressed through a clear recommendation and an associated break-even value. Because each analysis can be completed in seconds, the approach also makes repeated sensitivity and robustness studies practical.

Overall, the dashed ellipse of the literature review, analytical cost with simulator mechanisms, reachable in principle and occupied by nobody, is for this system class no longer empty. The models are built, the bridge to the reference is validated mechanism by mechanism, and the reach it buys is demonstrated on the decisions the field actually faces. That is the central contribution of this thesis, in the terms of Section 5.4, measured throughout.

6.3 Limitations and further work

The limitations of the present work also define the most direct avenues for further research. These can be separated into extensions of the modelling framework and improvements to the case-study parametrisation. At the framework level, the reference case considers a single turbine served by dedicated vessels. Extending the model to include vessel sharing, maintenance queues and farm-level charter strategies is therefore a natural next step beyond assumption A-9. These effects are particularly relevant to the mobilisation costs identified in Chapter 5, whose economic significance is likely to increase at wind-farm scale. The exploitation studies are also based on steady-state analysis, whereas seasonal charter rates and the timing of maintenance campaigns are inherently calendar-dependent. Such questions could be addressed using the transient formulation already developed in this thesis.

A further limitation arises from the marginal treatment of interruptibility in A-3, which neglects the dependence between task spans and the weather state prevailing during execution. Should this residual dependence prove important outside the bounds quantified here, it could be represented through a weather-state-augmented kernel. A more detailed task representation would also provide a useful intermediate case between the two execution assumptions examined in this work. In practice, many interventions combine interruptible labour with shorter phases that require uninterrupted weather windows. A phase-structured description could represent this more faithfully. The bounds reported in Chapter 5 suggest that the resulting corrections would be limited, but the formulation would better reflect real maintenance procedures.

At the case-study level, the parametrisation inherits the $\beta = 1$ configuration and the default WOMBAT inputs as part of the deliberate model-alignment strategy. Two developments in the available data would allow the framework to be extended more substantially. First, reliable estimates of renewal time-to-failure shape parameters remain scarce. Under the estimand criteria established in Chapter 2, only one defensible value was identified in the literature. As reliability datasets mature and begin to provide mode-specific renewal estimates, the mixed-shape formulations developed here can incorporate them directly. Second, the model family could be extended beyond corrective maintenance to preventive and condition-based strategies. The operational-age kernel already contains the state information on which such policies depend, providing a suitable basis for these extensions.

References

[1] Lazard, "Levelized cost of energy+ (lcoe+): Version 18.0," tech. rep., Lazard, June 2025.

[2] K. Rohrig, V. Berkhout, D. Callies, M. Durstewitz, S. Faulstich, B. Hahn, M. Jung, L. Pauscher, A. Seibel, M. Shan, M. Siefert, J. Steffen, M. Collmann, S. Czichon, M. Dörenkämper, J. Gottschall, B. Lange, A. Ruhle, F. Sayer, B. Stoevesandt, and J. Wenske, "Powering the 21st century by wind energy—options, facts, figures," *Applied Physics Reviews*, vol. 6, p. 031303, 08 2019.

[3] T. Stehly, P. Beiter, P. Duffy, and M. Shields, "Cost of wind energy review: 2024 edition," Tech. Rep. NREL/TP-5000-91775, National Renewable Energy Laboratory (NREL), Golden, CO, USA, Dec. 2024.

[4] S. Faulstich, B. Hahn, and P. J. Tavner, "Wind turbine downtime and its importance for offshore deployment," *Wind Energy*, vol. 14, no. 3, pp. 327–337, 2011.

[5] R. Hammond and A. Cooperman, "Windfarm operations and maintenance cost-benefit analysis tool (WOMBAT)," Tech. Rep. NREL/TP-5000-83712, National Renewable Energy Laboratory, Golden, CO, 2022.

[6] J. Feuchtwang and D. Infield *Wind Energy*, vol. 16, no. 7, pp. 1049–1066, 2013.

[7] J. Carroll, A. McDonald, and D. McMillan, "Failure rate, repair time and unscheduled O&M cost analysis of offshore wind turbines," *Wind Energy*, vol. 19, no. 6, pp. 1107–1119, 2016.

[8] I. Dinwoodie, O.-E. V. Endrerud, M. Hofmann, R. Martin, and I. B. Sperstad, "Reference cases for verification of operation and maintenance simulation models for offshore wind farms," *Wind Engineering*, vol. 39, no. 1, pp. 1–14, 2015.

[9] IPCC, "Climate change 2021: The physical science basis. summary for policymakers," working group i contribution to the sixth assessment report, Intergovernmental Panel on Climate Change (IPCC), 2021.

[10] UNFCCC, "Paris agreement." United Nations Framework Convention on Climate Change, 2015.

[11] Council of the European Union, "How the EU is greening energy." Web page, n.d.

[12] K. Khan, A. Khurshid, J. Cifuentes-Faura, and D. Xianjun, "Does renewable energy development enhance energy security?," *Utilities Policy*, vol. 87, p. 101725, 2024.

[13] International Energy Agency (IEA), "Energy security," June 2025.

[14] European Commission, "Repowereu: A plan to rapidly reduce dependence on russian fossil fuels and fast-forward the green transition," May 2022.

[15] International Renewable Energy Agency (IRENA), "Renewable capacity statistics 2025," tech. rep., IRENA, Mar. 2025.

[16] International Energy Agency (IEA), "Wind." IEA Energy System web page, n.d.

[17] Global Wind Energy Council (GWEC), "Offshore wind installed capacity reaches 83 gw as new report finds 2024 a record year for construction and auctions." GWEC news release, June 2025.

[18] U.S. Department of Energy, "Wind turbines: Bigger, better." U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy (EERE), article, n.d.

[19] WindEurope, "Wind energy in europe: 2024 statistics and the outlook for 2025–2030," tech. rep., WindEurope, Brussels, Belgium, Feb. 2025. Published in February 2025.

[20] J. Paquette, M. Williams, R. Clarke, M. Devin, S. Sheng, C. Constant, C. Clark, J. Fields, V. Gevorgian, M. Hall, J. Jonkman, J. Keller, A. Robertson, L. Sethuraman, and J. van Dam, "An operations and maintenance roadmap for u.s. offshore wind: Enabling a cost-effective and sustainable u.s. offshore wind energy industry through innovative operations and maintenance," tech. rep., U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy (EERE), Wind Energy Technologies Office, May 2024.

[21] A. McCoy, W. Musial, R. Hammond, D. Mulas Hernando, P. Duffy, P. Beiter, P. Perez, R. Baranowski, G. Reber, and P. Spitsen, "Offshore wind market report: 2024 edition," Tech. Rep. NREL/PR-5000-90897, National Renewable Energy Laboratory (NREL), Golden, CO, Aug. 2024.

[22] International Energy Agency (IEA), "Offshore wind outlook 2019," world energy outlook special report, International Energy Agency, Paris, Nov. 2019. Licence: CC BY 4.0. PDF: https://iea.blob.core.windows.net/assets/495ab264-4ddf-4b68-b9c0-514295ff40a7/Offshore-Wind-Outlook-2019.pdf.

[23] United Nations Environment Programme (UNEP), "Ocean, seas and coasts." Web page, Feb. 2026. Last updated 24 Feb 2026.

[24] World Forum Offshore Wind (WFO), "Global offshore wind report 2024," tech. rep., World Forum Offshore Wind e.V., Apr. 2025. Published 30 Apr 2025.

[25] J. McMorland, C. Flannigan, J. Carroll, M. Collu, D. McMillan, W. Leithead, and A. Coraddu, "A review of operations and maintenance modelling with considerations for novel wind turbine concepts," *Renewable and Sustainable Energy Reviews*, vol. 165, p. 112581, 2022.

[26] Accenture and Offshore Renewable Energy (ORE) Catapult, "Wind industry digital landscape review," tech. rep., Accenture and ORE Catapult, Oct. 2020.

[27] Offshore Renewable Energy (ORE) Catapult, "Offshore wind operations & maintenance: A £9 billion per year opportunity by 2030 for the UK to seize," tech. rep., ORE Catapult, June 2021. Published 1 June 2021.

[28] GL Garrad Hassan, "A guide to uk offshore wind operations and maintenance," tech. rep., Scottish Enterprise and The Crown Estate, May 2013.

[29] I. Dinwoodie, *Modelling the Operation and Maintenance of Offshore Wind Farms*. PhD thesis, University of Strathclyde, Glasgow, United Kingdom, 2014.

[30] H. Seyr and M. Muskulus, "Decision support models for operations and maintenance for offshore wind farms: A review," *Applied Sciences*, vol. 9, no. 2, p. 278, 2019.

[31] N. Avanessova, *Developing a Holistic Operation and Maintenance Simulation Tool for Emerging Offshore Wind Projects*. PhD thesis, University of Strathclyde, Glasgow, United Kingdom, 2025. PhD Thesis.

[32] BVG Associates, "Guide to a floating offshore wind farm." `https://guidetofloatingoffshorewind.com`. Published with ORE Catapult and Crown Estate Scotland.

[33] M. J. Kaiser and B. F. Snyder, "Modeling offshore wind installation vessel day-rates in the United States," *Maritime Economics & Logistics*, 2012.

[34] O. Donnelly, F. Anderson, and J. Carroll, "Operation and maintenance cost comparison between 15 MW direct-drive and medium-speed offshore wind turbines," *Wind Energy Science*, vol. 9, pp. 1345–1362, 2024. DOI: 10.5194/wes-9-1345-2024. Vessel inputs therein compile wave limits from Dalgic et al. and McMorland et al. and charter rates from ORE Catapult (2019, 2020).

[35] Y. Dalgic, I. Lazakis, and O. Turan, "Vessel charter rate estimation for offshore wind O&M activities," in *Developments in Maritime Transportation and Exploitation of Sea Resources*, 2013.

[36] O. Donnelly and J. Carroll, "Analysing the cost impact of failure rates for the next generation of offshore wind turbines," *Wind Energy*, 2024.

[37] DNV, "80% of insurance claims in offshore wind are related to subsea cable failures: How can the industry manage these risks?." DNV web article, n.d.

[38] M. Ford, "Unravelling subsea cable failure in offshore wind." International Marine Contractors Association (IMCA) commentary, Nov. 2023.

[39] M. Li, X. Jiang, H. Polinder, and R. R. Negenborn, "A review of maintenance strategy optimization for wind energy," in *Developments in Renewable Energies Offshore*, pp. 469–480, CRC Press / Balkema (Taylor & Francis Group), 2020.

[40] M. Li, M. Wang, J. Kang, L. Sun, and P. Jin, "An opportunistic maintenance strategy for offshore wind turbine system considering optimal maintenance intervals of subsystems," *Ocean Engineering*, vol. 216, p. 108067, 2020.

[41] J. J. Nielsen and J. D. Sørensen, "On risk-based operation and maintenance of offshore wind turbine components," *Reliability Engineering & System Safety*, vol. 96, no. 1, pp. 218–229, 2011.

[42] M. I. H. Tusar and B. R. Sarker, "Maintenance cost minimization models for offshore wind farms: A systematic and critical review," *International Journal of Energy Research*, vol. 46, pp. 3739–3765, 2022.

[43] R. Dawid, D. McMillan, M. Revie, and I. Lazakis, "Review of markov models for maintenance optimization in the context of offshore wind," in *Annual Conference of the PHM Society*, vol. 7, 2015.

[44] I. B. Sperstad, M. Stålhane, I. Dinwoodie, O.-E. V. Endrerud, R. Martin, and E. Warner, "Testing the robustness of optimal access vessel fleet selection for operation and maintenance of offshore wind farms," *Ocean Engineering*, vol. 145, pp. 334–343, 2017.

[45] J. S. Nielsen and J. D. Sørensen, "Methods for risk-based planning of o&m of wind turbines," *Energies*, vol. 7, no. 10, pp. 6645–6664, 2014.

[46] M. Centeno-Telleria, J. I. Aizpurua, and M. Penalba, "Computationally efficient analytical O&M model for strategic decision-making in offshore renewable energy systems," *Energy*, vol. 285, p. 129374, 2023.

[47] M. Stålhane, E. E. Halvorsen-Weare, L. M. Nonås, and G. Pantuso, "Optimizing vessel fleet size and mix to support maintenance operations at offshore wind farms," *European Journal of Operational Research*, vol. 276, no. 2, pp. 495–509, 2019.

[48] C. A. Irawan, D. Ouelhadj, D. Jones, M. Stålhane, and I. B. Sperstad, "Optimisation of maintenance routing and scheduling for offshore wind farms," *European Journal of Operational Research*, vol. 256, no. 1, pp. 76–89, 2017.

[49] T. Stehly, P. Duffy, and D. Mulas Hernando, "2022 cost of wind energy review," Tech. Rep. NREL/TP-5000-88335, National Renewable Energy Laboratory, Golden, CO, 2023.

[50] Z. Ren, A. S. Verma, Y. Li, J. J. E. Teuwen, and Z. Jiang, "Offshore wind turbine operations and maintenance: A state-of-the-art review," *Renewable and Sustainable Energy Reviews*, vol. 144, p. 110886, 2021.

[51] J. Carroll, A. McDonald, I. Dinwoodie, D. McMillan, M. Revie, and I. Lazakis, "Availability, operation and maintenance costs of offshore wind turbines with different drive train configurations," *Wind Energy*, vol. 20, no. 2, pp. 361–378, 2017.

[52] K. Anastasiou and C. Tsekos, "Persistence statistics of marine environmental parameters from Markov theory, Part 1: analysis in discrete time," *Applied Ocean Research*, vol. 18, no. 4, pp. 187–199, 1996.

[53] International Electrotechnical Commission, "IEC 61400-26-1:2019: Wind energy generation systems – Part 26-1: Availability for wind energy generation systems," international standard, International Electrotechnical Commission, Geneva, 2019. Edition 1.0.

[54] Y. Dalgic, I. Lazakis, I. Dinwoodie, D. McMillan, and M. Revie, "Advanced logistics planning for offshore wind farm operation and maintenance activities," *Ocean Engineering*, vol. 101, pp. 211–226, 2015.

[55] I. B. Sperstad, F. D. McAuliffe, M. Kolstad, and S. Sjømark, "Investigating key decision problems to optimize the operation and maintenance strategy of offshore wind farms," *Energy Procedia*, vol. 94, pp. 261–268, 2016.

[56] M. Hofmann, "A review of decision support models for offshore wind farms with an emphasis on operation and maintenance strategies," *Wind Engineering*, vol. 35, no. 1, pp. 1–15, 2011.

[57] I. Dinwoodie, D. McMillan, M. Revie, I. Lazakis, and Y. Dalgic, "Development of a combined operational and strategic decision support model for offshore wind," *Energy Procedia*, vol. 35, pp. 157–166, 2013.

[58] M. Hofmann and I. B. Sperstad, "NOWIcob – a tool for reducing the maintenance costs of offshore wind farms," in *Energy Procedia (DeepWind 2013)*, vol. 35, pp. 177–186, 2013.

[59] F. P. Santos, Â. P. Teixeira, and C. Guedes Soares, "Modelling and simulation of the operation and maintenance of offshore wind turbines," *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, vol. 229, no. 5, pp. 385–393, 2015.

[60] B. Le and J. Andrews, "Modelling wind turbine degradation and maintenance," *Wind Energy*, vol. 19, no. 4, pp. 571–591, 2016.

[61] M. Scheu, D. Matha, M. Hofmann, and M. Muskulus, "Maintenance strategies for large offshore wind farms," *Energy Procedia*, vol. 24, pp. 281–288, 2012.

[62] M. Shafiee, "Maintenance logistics organization for offshore wind energy: Current progress and future perspectives," *Renewable Energy*, vol. 77, pp. 182–193, 2015.

[63] L.-J. Huang, S. Mancini, D. M. Hernando, and R. Hammond, "Verification-based assessment of modelling assumptions in discrete-event simulation of operation and maintenance for floating offshore wind." Wind Energy Science Discussions, preprint wes-2025-279, in review, 2026. DOI: 10.5194/wes-2025-279.

[64] J. McMorland, M. Collu, D. McMillan, and J. Carroll, "Operation and maintenance for floating wind turbines: A review," *Renewable and Sustainable Energy Reviews*, vol. 163, p. 112499, 2022.

[65] E. Byon, L. Ntaimo, and Y. Ding, "Optimal maintenance strategies for wind turbine systems under stochastic weather conditions," *IEEE Transactions on Reliability*, vol. 59, no. 2, pp. 393–404, 2010.

[66] S. M. Ross, *Stochastic Processes*. New York: John Wiley & Sons, 2nd ed., 1996.

[67] K. S. Trivedi, *Probability and Statistics with Reliability, Queuing and Computer Science Applications*. New York: John Wiley & Sons, 2nd ed., 2002.

[68] R. E. Barlow and F. Proschan, *Statistical Theory of Reliability and Life Testing: Probability Models*. New York: Holt, Rinehart and Winston, 1975.

[69] R. A. Howard, *Dynamic Probabilistic Systems, Volume II: Semi-Markov and Decision Processes*. New York: John Wiley & Sons, 1971.

[70] C. Moler and C. Van Loan, "Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later," *SIAM Review*, vol. 45, no. 1, pp. 3–49, 2003.

[71] F. Besnard, K. Fischer, and L. B. Tjernberg, "A model for the optimization of the maintenance support organization for offshore wind farms," *IEEE Transactions on Sustainable Energy*, vol. 4, no. 2, pp. 443–450, 2013.

[72] L.-L. Huang, Y. Fu, Y. Mi, J.-L. Cao, and P. Wang, "A Markov-chain-based availability model of offshore wind turbine considering accessibility problems," *IEEE Transactions on Sustainable Energy*, vol. 8, no. 4, pp. 1592–1600, 2017.

[73] G. Abeynayake, T. Van Acker, D. Van Hertem, and J. Liang, "Analytical model for availability assessment of large-scale offshore wind farms including their collector system," *IEEE Transactions on Sustainable Energy*, vol. 12, no. 3, pp. 1974–1983, 2021.

[74] R. Sun, G. Abeynayake, J. Liang, and K. Wang, "Reliability and economic evaluation of offshore wind power DC collection systems," *Energies*, vol. 14, no. 10, p. 2922, 2021.

[75] J. Zhang, S. Chowdhury, J. Zhang, W. Tong, and A. Messac, "Optimal selection of time windows for preventive maintenance of offshore wind farms subject to wake losses," *Wind Energy*, vol. 26, no. 11, pp. 1103–1122, 2023.

[76] M. Centeno-Telleria, H. Yue, J. Carroll, J. I. Aizpurua, and M. Penalba, "O&m-aware techno-economic assessment for floating offshore wind farms: A geospatial evaluation off the north sea and the iberian peninsula," *Applied Energy*, vol. 371, p. 123684, 2024.

[77] M. Centeno Telleria, *PhD Thesis: Offshore wind O&M modelling*. PhD thesis, n.d., n.d.

[78] R. Pyke, "Markov renewal processes: Definitions and preliminary properties," *The Annals of Mathematical Statistics*, vol. 32, no. 4, pp. 1231–1242, 1961.

[79] W. L. Smith, "Regenerative stochastic processes," *Proceedings of the Royal Society of London. Series A*, vol. 232, no. 1188, pp. 6–31, 1955.

[80] E. Çinlar, "Markov renewal theory," *Advances in Applied Probability*, vol. 1, no. 2, pp. 123–187, 1969.

[81] N. Limnios and G. Oprişan, *Semi-Markov Processes and Reliability.* Boston: Birkhäuser, 2001.

[82] F. Grabski, *Semi-Markov Processes: Applications in System Reliability and Maintenance.* Amsterdam: Elsevier, 2015.

[83] M. Rausand and A. Høyland, *System Reliability Theory: Models, Statistical Methods, and Applications.* Hoboken, NJ: Wiley-Interscience, 2 ed., 2004.

[84] G. Corradi, J. Janssen, and R. Manca, "Numerical treatment of homogeneous semi-Markov processes in transient case – a straightforward approach," *Methodology and Computing in Applied Probability*, vol. 6, no. 2, pp. 233–246, 2004.

[85] M. d. C. Moura and E. L. Droguett, "Mathematical formulation and numerical treatment based on transition frequency densities and quadrature methods for non-homogeneous semi-Markov processes," *Reliability Engineering & System Safety*, vol. 94, no. 2, pp. 342–349, 2009.

[86] G. D'Amico, F. Petroni, and F. Prattico, "First and second order semi-Markov chains for wind speed modeling," *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 5, pp. 1194–1201, 2013.

[87] R. Dawid, D. McMillan, and M. Revie, "Time series semi-Markov decision process with variable costs for maintenance planning," in *Proceedings of the European Safety and Reliability Conference (ESREL)*, (Glasgow, UK), 2016.

[88] T. M. Welte and K. Wang, "Models for lifetime estimation: an overview with focus on applications to wind turbines," *Advances in Manufacturing*, vol. 2, no. 1, pp. 79–87, 2014.

[89] B. Hagen, I. Simonsen, M. Hofmann, and M. Muskulus, "A multivariate Markov weather model for O&M simulation of offshore wind parks," *Energy Procedia*, vol. 35, pp. 137–147, 2013.

[90] M. Martini, R. Guanche, I. J. Losada, and C. Vidal, "Accessibility assessment for operation and maintenance of offshore wind farms in the North Sea," *Wind Energy*, vol. 20, no. 4, pp. 637–656, 2017.

[91] I. B. Sperstad, E. E. Halvorsen-Weare, M. Hofmann, L. M. Nonås, M. Stålhane, and M. Wu, "A comparison of single- and multi-parameter wave criteria for accessing wind turbines in strategic maintenance and logistics models for offshore wind farms," *Energy Procedia*, vol. 53, pp. 221–230, 2014.

[92] T. Gintautas and J. D. Sørensen, "Improved methodology of weather window prediction for offshore operations based on probabilities of operation failure," *Journal of Marine Science and Engineering*, vol. 5, no. 2, p. 20, 2017.

[93] V. Monbet, P. Ailliot, and M. Prevosto, "Survey of stochastic models for wind and sea state time series," *Probabilistic Engineering Mechanics*, vol. 22, no. 2, pp. 113–126, 2007.

[94] O. Dahmani, S. Bourguet, M. Machmoum, P. Guerin, and P. Rhein, "Reliability analysis of the collection system of an offshore wind farm," in *Proceedings of the Ninth International Conference on Ecological Vehicles and Renewable Energies (EVER)*, pp. 1–6, IEEE, 2014.

[95] P. J. Tavner, J. Xiang, and F. Spinato, "Reliability analysis for wind turbines," *Wind Energy*, vol. 10, no. 1, pp. 1–18, 2007.

[96] F. Spinato, P. J. Tavner, G. J. W. van Bussel, and E. Koutoulakos, "Reliability of wind turbine subassemblies," *IET Renewable Power Generation*, vol. 3, no. 4, pp. 387–401, 2009.

[97] S. Faulstich, V. Berkhout, J. Mayer, and D. Siebenlist, "Modelling the failure behaviour of wind turbines," *Journal of Physics: Conference Series*, vol. 749, p. 012019, 2016.

[98] F. Harzendorf, R. Schelenz, and G. Jacobs, "Reducing cost uncertainty in the drivetrain design decision with a focus on the operational phase," *Wind Energy Science*, vol. 6, no. 2, pp. 571–584, 2021.

[99] J. Walgern, F. Anderson, A. Kolios, and K. Fischer, "Field-data-based wind turbine reliability modelling: quantifying effects of operating age, design and technological development," *Wind Energy*, vol. 29, no. 5, p. e70108, 2026. DOI: 10.1002/we.70108.

[100] H. Ascher and H. Feingold, *Repairable Systems Reliability: Modeling, Inference, Misconceptions and Their Causes*. New York: Marcel Dekker, 1984.

[101] B. H. Lindqvist, "On the statistical modeling and analysis of repairable systems," *Statistical Science*, vol. 21, no. 4, pp. 532–551, 2006.

[102] M. Kijima, "Some results for repairable systems with general repair," *Journal of Applied Probability*, vol. 26, no. 1, pp. 89–102, 1989.

[103] V. Slimacek and B. H. Lindqvist, "Reliability of wind turbines modeled by a Poisson process with covariates, unobserved heterogeneity and seasonality," *Wind Energy*, vol. 19, no. 11, pp. 1991–2002, 2016.

[104] M. Shrivastava and A. S. Baghel, "Failure analysis and optimization of maintenance for the main shaft of 1500 kW wind turbine," *International Journal of Engineering Research and Advanced Technology*, vol. 4, no. 5, pp. 25–29, 2018. Field-fitted renewal time-to-failure shape, drivetrain (main shaft), $\beta = 1.43$.

[105] M. Asgarpour and J. D. Sørensen, "Bayesian based prognostic model for predictive maintenance of offshore wind farms," *International Journal of Prognostics and Health Management*, vol. 9, no. 1, p. 010, 2018.

[106] G. Smart, A. Smith, E. Warner, I. B. Sperstad, B. Prinsen, and R. Lacal-Arántegui, "IEA wind task 26: Offshore wind farm baseline documentation," Tech. Rep. NREL/TP-6A20-66262, National Renewable Energy Laboratory, Golden, CO, 2016.

[107] R. G. Sargent, "Verification and validation of simulation models," *Journal of Simulation*, vol. 7, no. 1, pp. 12–24, 2013.

[108] O.-E. V. Endrerud, J. P. Liyanage, and N. Keseric, "Marine logistics decision support for operation and maintenance of offshore wind parks with a multi method simulation model," in *Proceedings of the 2014 Winter Simulation Conference*, pp. 1712–1722, IEEE, 2014.

[109] D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *The Journal of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.

[110] T. Bedford and B. H. Lindqvist, "The identifiability problem for repairable systems subject to competing risks," *Advances in Applied Probability*, vol. 36, no. 3, pp. 774–790, 2004.

[111] DNV, "DNV-ST-N001: Marine operations and marine warranty," standard, DNV AS, 2023. Edition December 2023.

[112] Z. Jiang, "Installation of offshore wind turbines: A technical review," *Renewable and Sustainable Energy Reviews*, vol. 139, p. 110576, 2021.

[113] L. Mishnaevsky, Jr., "Repair of wind turbine blades: Review of methods and related computational mechanics problems," *Renewable Energy*, vol. 140, pp. 828–839, 2019.

[114] R. Martin, I. Lazakis, S. Barbouchi, and L. Johanning, "Sensitivity analysis of offshore wind farm operation and maintenance cost and availability," *Renewable Energy*, vol. 85, pp. 1226–1236, 2016.

[115] A. Ioannou, A. Angus, and F. Brennan, "Informing parametric risk control policies for operational uncertainties of offshore wind energy assets," *Ocean Engineering*, vol. 177, pp. 1–11, 2019.

[116] Y. Dalgic, I. Lazakis, and O. Turan, "Vessel charter rate estimation for offshore wind O&M activities," in *Developments in Maritime Transportation and Exploitation of Sea Resources (IMAM 2013)*, (A Coruña, Spain), pp. 899–907, CRC Press, 2013.

[117] Y. Dalgic, I. Lazakis, and O. Turan, “Investigation of optimum crew transfer vessel fleet for offshore wind farm maintenance operations,” *Wind Engineering*, vol. 39, no. 1, pp. 31–52, 2015.

[118] W. Guachamin-Acero, Z. Jiang, and L. Li, “Numerical study of a concept for major repair and replacement of offshore wind turbine blades,” *Wind Energy*, vol. 23, no. 8, pp. 1673–1692, 2020.

[119] Y. Hadjoudj and R. K. Pandit, “Improving O&M decision tools for offshore wind farm vessel routing by incorporating weather uncertainty,” *IET Renewable Power Generation*, vol. 17, no. 6, pp. 1488–1499, 2023.

[120] Y. Dalgic, I. Lazakis, O. Turan, and S. Judah, “Investigation of optimum jack-up vessel chartering strategy for offshore wind farm O&M activities,” *Ocean Engineering*, vol. 95, pp. 106–115, 2015.

[121] O. Donnelly, J. Carroll, and M. Howland, “Analysing the cost impact of failure rates for the next generation of offshore wind turbines,” *Wind Energy*, vol. 27, no. 7, pp. 695–710, 2024.

A Mathematical background and derivations

This appendix supports Appendix 3: parts A.1, A.2 and A.4 state the standard results the chapter uses, with the textbook treatments of [66, 67, 81, 82, 69, 68] as the general references, and parts A.3 and A.5 prove the two results that are specific to this thesis.

A.1 Continuous-time Markov chain machinery

A continuous-time Markov chain on a finite state space is specified by its generator $\mathbf{Q}$, whose off-diagonal entries are the jump rates and whose diagonals are fixed by row sums. The state probabilities obey the Kolmogorov forward equations, whose solution from any initial condition is the matrix exponential $\boldsymbol{p}(t) = \boldsymbol{p}(0)\, e^{\mathbf{Q}t}$ quoted in Subsection 3.4.1 [66, 67].

The derivation the chapter promises is that of the closed form Equation 12. The stationary distribution solves $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ with unit sum, and on the topology of Equation 2 with the rates of Equation 11 the balance equations decouple: the cut around W_i gives $\pi_{\mathrm{UP}}\, \lambda_i^h = \pi_{\mathrm{W}_i}\, \alpha_i$, and the cut around R_i gives $\pi_{\mathrm{W}_i}\, \alpha_i = \pi_{\mathrm{R}_i}\, \mu_i$. Therefore

$$\pi_{\mathrm{W}_i} \;=\; \pi_{\mathrm{UP}}\, \lambda_i^h\, \tau_i^{\mathrm{W}}, \qquad \pi_{\mathrm{R}_i} \;=\; \pi_{\mathrm{UP}}\, \lambda_i^h\, \tau_i^{\mathrm{R}}, \tag{20}$$

with $\tau_i^{\mathrm{W}} = 1/\alpha_i$ and $\tau_i^{\mathrm{R}} = 1/\mu_i$. Normalising and substituting $\tau_{\mathrm{UP}} = 1/\sum_j \lambda_j^h$ and $q_i = \lambda_i^h / \sum_j \lambda_j^h$ yields exactly Equation 12, mean up-time over mean cycle, and the task frequency follows as $\omega_i = \pi_{\mathrm{UP}} \lambda_i^h$.

A.2 Renewal and Markov renewal theory

The renewal–reward theorem states that for independent, identically distributed cycles of length X carrying rewards Y with finite means, the long-run reward rate equals the ratio of means,

$$\lim_{t\to\infty} \frac{1}{t} \sum_{k:\, S_k \le t} Y_k \;=\; \frac{\mathbb{E}[Y]}{\mathbb{E}[X]} \qquad \text{almost surely,} \tag{21}$$

by the strong law applied to reward sums and cycle counts [66]. Every long-run occupancy in this thesis is an instance of Equation 21 with Y the time spent in the state of interest during a cycle.

A Markov renewal process is specified by the semi-Markov kernel $Q_{jk}(t)$, the probability that the next state is k and the holding time is at most t given the current state j. With $H_j(t) = \sum_k Q_{jk}(t)$, the transition functions of the associated semi-Markov process satisfy the Markov renewal equations

$$P_{jk}(t) \;=\; \delta_{jk}\,\bigl(1 - H_j(t)\bigr) \;+\; \sum_{l} \int_0^t \mathrm{d}Q_{jl}(u)\; P_{lk}(t-u), \tag{22}$$

a Volterra system whose solution on finite spaces exists and is unique [80, 81]. The discretisation used by the time-marching model replaces the integral by a convolution sum over kernel increments on a uniform grid, which marches forward in time since each value uses only earlier ones. The scheme is first-order accurate for kernels of bounded variation, preserves substochasticity when the increments are non-negative, which Section 3.8 asserts as a running invariant, and its error is controlled by grid refinement [84, 85].

A.3 The invariance theorem

Let Z be an irreducible semi-Markov process on a finite state space with embedded jump chain $\{J_m\}$, stationary jump-chain distribution $\boldsymbol{\nu}$, and holding-time laws with finite means m_j in each state j. Then the limiting occupancy of state j is

$$P_j \;=\; \frac{\nu_j\, m_j}{\sum_k \nu_k\, m_k}, \tag{23}$$

and therefore depends on the holding-time laws only through their means.

Proof. Consider the successive visits to any fixed reference state as regeneration points, which exist by irreducibility and finiteness. Within one regeneration cycle, the expected number of visits to state j is proportional to ν_j by the theory of the embedded chain, and each visit contributes an expected sojourn m_j, independent of the visit count by the semi-Markov property. By the Wald identity the expected time spent in j per cycle is proportional to $\nu_j m_j$, and the expected cycle length is proportional to $\sum_k \nu_k m_k$. The renewal–reward theorem Equation 21 then gives the occupancy as the ratio Equation 23. No property of the holding-time laws beyond the first moment enters at any step.

Applied to the cycle structure of Equation 16, the theorem yields the steady-state solver directly, and read in reverse it is the certification limit used in Appendix 2 and Section 3.7, as steady-state agreement can never certify distributional structure, so shape effects must be sought in transients and distributions.

A.4 Competing Weibull risks under partial reset

With mode i carrying the hazard and survival of Equation 13 and the modes conditionally independent given the age vector, each mode survives a further operational time u with probability $R_j(a_j + u)/R_j(a_j)$, the asset survives while every mode survives, and attributing the failure to the mode whose hazard fires gives the sub-density $h_i(a_i + u) \prod_j R_j(a_j +$

$u)/R_j(a_j)$. Together these are Equation 14, the classical competing-risks construction with the ages entering only through the conditioning [68]. The repair regime, failed mode restored to age zero, survivors keeping their accumulated operational age, is the identifiable partial-repair setting of [110] within the virtual-age families of [102], and because the clocks pause during downtime, age lives on the operational clock and the kernel is solvable independently of the weather operators.

The binned kernel replaces Equation 14 by its quadrature on M bins per mode and solves the resulting finite embedded chain by power iteration. Convergence in the grid limit follows from the continuity of the Weibull hazard and survival. On any truncated age domain the binned kernel converges to the continuous one as $M \to \infty$, the truncation error is controlled by extending the grid to several multiples of the largest mean life where the residual survival mass is negligible, and the stationary distributions converge accordingly. Discretisation is therefore the only error of the construction, consistent with the grid study reported in Section 3.6.

A.5 The span operator as a renewal–reward functional

The claim of Subsection 3.3.3 is that the uniformly weighted start-hour mean Equation 6 equals the exact long-run mean restoration span of a stationary failure stream over the recorded climate, with no stochastic error.

The argument can be divided into three steps:

- Firstly, given the record and its cyclic tiling, $\mathrm{span}(s; L)$ is a deterministic function of the start hour s, so the only randomness in a restoration span is the randomness of its start hour.
- Secondly, consider the sequence of restoration starts generated by the stationary failure stream that the model posits (A-2, A-8). Under stationarity the start hours are uniformly distributed over the record phase, and the assumption is well founded since the failure clocks advance on the operational clock and carry no explicit phase information, the record is tiled cyclically, and the one channel through which phase memory could persist (which is the tendency of repairs to end in accessible weather) decays over up-times that are long compared with the correlation scale of the mask. The uniform weighting ρ is therefore the phase law of the stationary stream, and any departure from it would require phase information that the model does not carry.
- Thirdly, apply the renewal–reward theorem Equation 21 with the cycle reward $Y = \mathrm{span}(S; L)$ where S is the start hour of the restoration of the cycle: the long-run mean restoration span equals $\mathbb{E}[\mathrm{span}(S; L)]$ under the limiting start-hour law, which by the second step is the uniform average Equation 6. Since the span values themselves are deterministic given the record, the resulting mean carries no sampling error whatsoever, which is the sense of the claim made in the chapter.

The same argument identifies the start-hour ensemble Equation 7 as the exact marginal law of the restoration span under the stationary stream, the object assumption A-3 names and the time-marching SMP consumes. What A-3 discards is only the joint information between the span and the concurrent weather state at the transition instants, and Chapter 5 measures the cost of that marginal treatment.

B Code access and download

The final version of the models developed for this thesis, the WOMBAT patches, and the reference input dataset are all available for download in *https://github.com/blisken1303/EWEM-MSc-Thesis/*.

C Declaration on the use of AI

Throughout the duration of this thesis, the generative artificial intelligence chatbots *ChatGPT* and *Claude* were used as support tools by the author to aid in the development of the thesis and the preparation of this document. Specifically, they were used in the formatting of equations and the creation of figures for Overleaf, the editing and rephrasing of certain paragraphs and sentences to improve readability and incorporate a more academic tone and voice, the formatting of references into BibTeX, and in the debugging of the code written by the author.