

Agent-Based Simulation of Offshore Well Decommissioning Campaigns: Learning Curves, Weather Coupling, and the Economics of Rigless Plug and Abandonment

Ilan Gleiser¹, Erica Crandon¹

¹AlphaIQ

August 2026

Abstract

Offshore well decommissioning is the petroleum industry's next structural cost wave, yet permanent plug and abandonment (P&A) campaigns are still planned predominantly on static schedules that cannot represent the two couplings that dominate realised outcomes: weather risk compounds with campaign duration, and crew learning compounds with activity batching. We present *welldecom*, a pure-Python agent-based model in which wells, vessels, an operator, and a regulator interact under a discrete-event engine with six-hour metocean resolution. Task durations follow PERT distributions; crew proficiency follows Crawford geometric learning curves tracked per crew and activity family; downhole risk events are Bernoulli trials with triangular duration penalties; barrier placement follows the 30 CFR § 250.1715 rulebook with plug integrity testing as explicit success/failure nodes; and a seasonal AR(1) significant-wave-height process gates transit, dynamic positioning, and heavy lifts. The model passes a twelve-check verification and validation protocol, including exact reproduction of analytic limits, bit-exact cross-language replication against an independent TypeScript implementation, and comparative benchmark error of 3.2% (MAPE) against the specification's published strategic table on matched cost scope. Across 10,000-iteration Monte Carlo ensembles of a ten-well subsea campaign, a batched rigless programme using perforate-wash-cement technology delivers a 69% reduction in P50 well-work cost and a 45% reduction in P50 duration relative to a sequential rig-based section-milling campaign, with a 74% reduction in fleet emissions. We further show that seed-only ensembles understate honest contingency: stacking plausible input-parameter ranges on execution uncertainty widens the P10–P90 cost band by a factor of 2.1, and we argue that uncertainty bands must be labelled with the scope of what they vary.

Keywords: *agent-based modelling, well decommissioning, plug and abandonment, learning curves, Monte Carlo simulation, metocean risk, rigless intervention, PWC*

1. Introduction

More than a hundred thousand offshore wells worldwide will require permanent plug and abandonment over the coming three decades, and in mature basins the associated liability now rivals new-development capital budgets [10]. Unlike development drilling, decommissioning generates no revenue against which overruns can be absorbed: every day a spread waits on weather, every failed barrier verification, and every unplanned downhole event lands directly on a provision that regulators, auditors, and investors scrutinise. Yet the dominant planning instruments for P&A campaigns remain deterministic Gantt schedules and factored cost lines — tools that treat durations as known, weather as an allowance, and crews as interchangeable.

Two couplings, each individually well documented, interact to make static planning systematically wrong. First, weather risk compounds with duration: a campaign that runs longer samples more storms, and because waiting time is bought at the vessel spread rate, the marginal cost of schedule slip is convex. A static plan that books a fixed weather allowance cannot represent this feedback, and it penalises short campaigns and flatters long ones. Second, crew learning compounds with batching: repetitive execution of an identical operation produces geometric duration decay [1, 2], but the compounding only accrues when a crew works an unbroken sequence of the same activity. A campaign strategy decision — sequential well-by-well execution versus batched activity passes — therefore changes not just logistics but the shape of the learning curve each crew rides.

This paper presents *welldecom*, an agent-based model (ABM) of offshore well decommissioning campaigns in which these couplings are structural rather than assumed. Wells, vessels, an operator, and a regulator are explicit agents; the campaign unfolds through a discrete-event engine at six-hour metocean resolution; and cost, schedule, emissions, and regulatory compliance emerge from the interaction of stochastic task durations, discrete risk events, barrier-verification outcomes, and a seasonal wave-height process. The agent formulation follows the general argument for ABM in systems dominated by interaction effects [11]: the quantity of interest — outturn cost of a campaign — is not a sum of independent line items but the product of feedbacks that only a simulation can carry.

Our contributions are fourfold. *(i)* A complete, reproducible stochastic specification of a P&A campaign: PERT task durations, Crawford learning per crew and activity family, Bernoulli–triangular risk coupling, rule-based barrier placement under 30 CFR § 250.1715 [5] with integrity testing as success/failure nodes, and an AR(1) seasonal metocean process gating weather-sensitive operations. *(ii)* A verification and validation protocol with twelve programmatic checks, including exact analytic limits, bit-exact cross-language replication, and a matched-scope comparative benchmark achieving 3.2% MAPE. *(iii)* A 10,000-iteration Monte Carlo comparison of rig-based sequential milling against batched rigless perforate-wash-cement (PWC) execution, decomposing a 69% P50 well-work cost reduction into its drivers. *(iv)* A methodological finding on uncertainty reporting: seed-only ensembles price execution luck but not input error, and the honest contingency band — $2.1\times$ wider in our base case — stacks both.

2. Model Architecture

The model comprises four agent classes interacting in a shared environment. The environment holds the lease geometry (a well cluster approximately 120 nautical miles offshore), the metocean state, and the site-clearance verification rules. Time advances on a continuous day clock per vessel, with the weather sampled at six-hour resolution; an operational task makes progress only while the sea state is inside the acting vessel's limit for that task class, and otherwise the vessel accrues waiting-on-weather (WOW) time at or near its full spread rate.

2.1 Well agents

Each well carries a casing profile with explicit string tops and shoes (preventing double-counted footage across sidetracks), a water depth, a measured depth that scales task durations, flags for open hole below casing, open perforated intervals, casing stubs, and poor primary cement, and the NSTA mechanical-state machine ACTIVE/SUSPENDED → AB1 → AB2 (or AB2-Derogated) → AB3 [6]. State transitions are earned, not scheduled: a well reaches AB1 only when its reservoir barriers are placed and verified, and AB3 only when conductor removal and site clearance complete. A long-term gas-leakage probability falls with each independently verified barrier, and poor primary cement doubles the barrier-verification failure rate.

2.2 Vessel agents

Vessels are the execution units: light well-intervention vessels (LWIV), semisubmersible rigs, and workover units, each carrying a daily spread rate, transit speed, crane and deck limits, fuel-burn rates by activity state, and per-task sea-state ceilings — one limit for sailing, a tighter one for dynamic positioning, and the tightest for heavy lifts. Vessels occupy the states IDLE, MOBILIZING, TRANSITING, OPERATING, WOW, and NPT, and their fuel burn accumulates into a CO₂e account that the regulator prices.

2.3 Operator and regulator agents

The operator chooses the campaign strategy (sequential single-spread execution versus batched activity passes), the Phase-2 technology (section milling on a heavy rig versus single-trip perforate-wash-cement from a light vessel), whether to run Phase-0 pre-abandonment diagnostics (which reduce downhole risk-event probability from 0.30 to 0.05 per event class), and routes vessels to the nearest well with outstanding work. The regulator enforces a one-year post-CoP plugging deadline with per-day penalties, audits fleet emissions against a carbon price, and holds the methane-monitoring windows. Barrier placement itself is not an operator choice: it is compiled per well from the 30 CFR § 250.1715 if-then rulebook (open hole → straddle the deepest shoe by ±100 ft; open perforations → squeeze or a 200-ft plug; casing stub → straddle the stub), plus an environmental surface plug and, for poor-cement wells, an annular remediation barrier [5]. Every plug faces an integrity test — 15,000 lb of pipe weight or 1,000 psi held with less than 10% decay in fifteen minutes — modelled as a success/failure node whose failures append remedial cement work to the critical path.




Figure 1: Campaign progress ensembles (200 realisations): wells at AB3 versus time for the sequential rig campaign (left) and the batched rigless campaign (right), with sample metocean realisations against vessel operating limits and the two Crawford learning curves beneath. The batched campaign's plateau-then-sprint completion signature is a structural property of phased execution.

3. Stochastic Specification

3.1 Task durations

Each work-breakdown phase on each well draws its baseline duration from a PERT distribution parameterised by optimistic, most-likely, and pessimistic values (a , m , b), following standard three-point estimation practice [3, 4, 12]:

$$\mu_{\text{PERT}} = \frac{a + 4m + b}{6}, \quad \sigma_{\text{PERT}} = \frac{b - a}{6} \quad (1)$$

sampled as a scaled Beta with the exact shape parameters

$$\alpha = \frac{4m + b - 5a}{b - a}, \quad \beta = \frac{5b - a - 4m}{b - a}, \quad D \sim a + (b - a) \text{Beta}(\alpha, \beta) \quad (2)$$

Durations scale with measured depth, and Beta variates are generated by the Marsaglia–Tsang method [9] so that the TypeScript and Python implementations consume identical random streams (Section 5).

3.2 Crew learning

The duration of the N -th repetition of an activity by a given crew follows Crawford's geometric learning curve [1, 2]:

$$T_N = T_1 \cdot N^b, \quad b = \frac{\ln L}{\ln 2} \quad (3)$$

where L is the learning rate ($L = 0.95$ for conservative section milling; $L = 0.85$ for a standardised PWC tool set). Critically, the repetition counter is tracked *per crew and per activity family*: proficiency gained by the LWIV crew setting reservoir barriers does not transfer to the milling crew, and a sequential campaign interrupts every activity's sequence at each phase change. This bookkeeping choice is what makes campaign strategy and learning interact — batching is not merely a logistics preference but a decision about which learning curves the campaign gets to ride, an effect long observed in repetitive offshore operations [8].

3.3 Discrete risk events and cost coupling

Downhole and surface risk events (stuck tubing, lost circulation, gas migration during waiting-on-cement, swarf blockages, stuck mills, cut misruns) are Bernoulli trials evaluated at task start, each with a triangular duration penalty and a fixed tooling cost:

$$X_{n,p,i} \sim \text{Bernoulli}(p_i), \quad \Delta D_i \sim \text{Triangular}(a_i, m_i, b_i), \quad \Delta C_i = C_{\text{tool},i} + \Delta D_i \cdot R_{\text{spread}} \quad (4)$$

Phase-0 diagnostics scale the probabilities of downhole event classes by 0.05/0.30. Campaign cost is then fully coupled to realised durations:

$$C_{\text{total}} = C_{\text{mob}} + \sum_{n=1}^{N_w} \left[\sum_{p=1}^P D_{n,p} R_{\text{spread}(p)} + C_{\text{mat},n} + \sum_{i=1}^M X_{n,p,i} \Delta C_{n,p,i} \right] + C_{\text{CO}_2} \quad (5)$$

with costs booked across an eleven-category work-breakdown ledger and AACE-class contingency bands. The emissions term prices fleet fuel burn at a carbon price (default \$85/tCO₂e). We distinguish throughout between *full-scope* cost (all eleven WBS lines, including post-CoP running costs, project management, site clearance, recycling, and monitoring) and *well-work* cost (vessel spread time including mobilisation and WOW, materials, and risk costs — WBS categories 3, 4, and 7), because published strategic comparisons are typically quoted on the latter scope.

3.4 Metocean process

Significant wave height follows a seasonal AR(1) process in log space at six-hour resolution:

$$\ln H_s(t) = s(t) + \rho[\ln H_s(t-1) - s(t-1)] + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_w^2) \quad (6)$$

with the seasonal mean

$$s(t) = \ln \bar{H}_s + \frac{A}{\bar{H}_s} \cos\left(\frac{2\pi(\text{doy}(t) - 15)}{365.25}\right) \quad (7)$$

peaking in winter (day-of-year ≈ 15) and correlated wind drawn alongside. The autocorrelation $\rho = 0.90$ produces realistic multi-day weather windows and multi-day WOW spells — the phenomenon a fixed weather allowance cannot represent. Heavy lifts additionally require a continuous six-hour clean window to begin. One implementation detail matters for experimental design: the metocean field is drawn as a single leading block from the random stream *before* any task-level draw, so two campaign strategies compared on the same seed face bit-identical weather. Without this separation, strategy-dependent task draws would silently shift the weather sequence and same-seed A/B comparisons would not be weather-controlled — a subtle failure mode we detected and corrected in a derived implementation of this model.

4. Experimental Design

The reference experiment is a ten-well subsea campaign on a fixed synthetic lease (well geometry held constant across iterations so that ensemble dispersion reflects execution risk, not geology). Two strategies are compared. Scenario A executes sequentially from a semisubmersible at a \$350k/day spread rate with Phase-2 section milling (baseline 8–20 days per well) and a 95% learning rate. Scenario B batches phases across the campaign from an LWIV at \$180k/day with single-trip PWC (baseline 2–5 days) and an 85% learning rate: the light vessel sets reservoir barriers across all wells, executes the Phase-2 remediations, then returns for conductor removal and site clearance. Both scenarios run Phase-0 diagnostics and face the same spring campaign start (day-of-year 90).

Each scenario is simulated for 10,000 independent Monte Carlo iterations from spawned random substreams of one recorded master seed, making every figure in this paper exactly reproducible. Comparative acceptance ensembles (250 iterations each) probe seasonality (summer versus winter starts), the value of Phase-0 diagnostics, storm-scale sensitivity, and a learning-rate sweep.

Table 1: Phase table of the reference campaign: PERT duration parameters (days, first well, before depth scaling and learning), governing weather limit class, and leading risk events.

Phase	PERT (a, m, b)	Vessel	Limit	Leading risks (p)
0 — Diagnostics	(1, 2, 4)	LWIV	DP	logging tool failure (0.04)
1 — Reservoir isolation	(4, 6, 10)	LWIV	DP	stuck tubing (0.08); lost circulation (0.06)
2 — PWC remediation	(2, 3, 5)	LWIV	DP	stuck assembly (0.03); re-squeeze (0.05)
2 — Section milling	(8, 10, 20)	Rig	DP	swarf NPT (0.15); stuck mill (0.05)
3 — Conductor removal	(1.5, 2.5, 5)	LWIV	Lift	cut misrun (0.07); dropped object (0.02)

Site clearance	(0.75, 1.5, 3)	LWIV	DP	—
----------------	----------------	------	----	---

5. Verification and Validation

The model is graded against a twelve-check protocol: five verification checks that the code obeys its specified mathematics exactly, six behavioural acceptance checks on independent comparative ensembles, and one benchmark check against the specification's published strategic table. All twelve pass.

5.1 Verification

Under a unit learning rate the Monte Carlo mean campaign duration converges to the analytic sum of phase PERT means with 0.004% relative error. At the extreme learning rate $L = 0.5$ the implementation reproduces the analytic decay $T_n = T_1 \cdot N^{-1}$ to machine precision. With all risk probabilities set to zero, sampled durations fall exactly inside the $[\Sigma a, \Sigma b]$ PERT bounds. The full engine produces finite outputs with all wells reaching AB3 and ordered percentiles, and repeated ensembles from the same master seed are bit-identical. Beyond these, the model exists twice — in Python and in an independent TypeScript port sharing only the seeded Mulberry32 generator — and the two implementations agree bit-exactly on all 48 compared output fields across seeds and scenarios.

5.2 Behavioural acceptance

On 250-run comparative ensembles: percentiles are ordered; the rigless batched strategy dominates on both cost and duration; a winter campaign start waits 53.0 days on weather at P50 against 25.1 for a summer start on identical scope; Phase-0 diagnostics cut P50 non-productive time from 11.5 to 6.3 days; milling carries more NPT than PWC (7.8 versus 6.3 days at P50); and P50 duration is monotonic in the learning rate across $L \in [0.80, 1.00]$, a sweep spanning 142 to 210 days.

5.3 Benchmark validation

Absent audited outturn records for the synthetic lease, the model is benchmarked against its specification's published ten-well strategic comparison on scale-free comparative metrics (percentage reductions, Scenario A $\rightarrow$ B), using the MAPE criterion of Equation 8 with a 5% gate:

$$\text{MAPE} = \frac{100\%}{K} \sum_{k=1}^K \left| \frac{A_k - F_k}{A_k} \right| \quad (8)$$

Scope matters: the published table is vessel-spread-centric, so the comparison is made on well-work cost. On that matched scope the model reproduces the published P50 cost reduction to 0.6 percentage points (69.4% modelled versus 68.8% published) and achieves 3.2% MAPE across the four comparative metrics, inside the gate. The full-scope comparison (6.7% MAPE) is reported alongside for transparency; its residual error concentrates in the duration metric, where the model's

weather–duration feedback extends the rig campaign in a way a static benchmark cannot represent. Against real outturn data the identical procedure applies unchanged.

Validation score: 12/12	
PASS	V1 Mathematical conservation (L=1) MC mean 72.081 d vs analytic 72.083 d (rel. err. 0.0038%, tol 1.0%)
PASS	V2 Extreme learning ($L=0.5 \Rightarrow N^{-1}$) max $ N^b - 1/N $ over $N = 1...30$: 0.00e+00
PASS	V3 Null-risk PERT bounds observed [52.84, 85.38] d within [43.50, 117.00] d
PASS	V4 Full-engine smoke finite=True, all wells AB3=True, P10 \$17.5M ≤ P50 \$18.8M ≤ P90 \$21.5M
PASS	V5 Bit-exact reproducibility identical results from identical master seed (15×2 runs)
PASS	A1 Percentile ordering P10 ≤ P50 ≤ P90 for cost and duration, both scenarios
PASS	A2 Technology direction (B cheaper than A) P50 cost −69.4%, duration −45.1%
PASS	A3 Winter WOW > summer WOW (seasonality) winter P50 53.0 d vs summer 25.1 d
PASS	A4 Phase-0 diagnostics reduce NPT NPT P50 6.3 d with diagnostics vs 11.5 d without
PASS	A5 Milling NPT > PWC NPT P50 NPT 7.8 d milling vs 6.3 d PWC
PASS	A6 Learning monotonicity (duration ↑ with L) P50 duration non-decreasing across $L = 0.80 \dots 1.00$
PASS	H1 Benchmark MAPE < 5% (spec §6 table, matched scope) MAPE 3.2% on WBS 3+4+7 well-work scope (full scope: 6.7%)

Figure 2: The twelve-check validation scorecard: five verification checks (V), six behavioural acceptance checks (A), and the matched-scope benchmark (H). All twelve pass.

6. Results

6.1 Strategic comparison

Table 2 reports the 10,000-iteration ensembles. The batched rigless PWC campaign delivers P50 well-work cost of \$31.3M against \$102.5M for the sequential rig campaign — a 69.4% reduction — and completes in 155 days against 282 (−45.1%). The distributions barely overlap: Scenario B's P90 full-scope cost sits \$73M below Scenario A's P10. Fleet emissions fall 73.7%, from 16,760 to 4,416 tCO₂e at P50, because the semisubmersible both burns more fuel per day and remains on location far longer.

Table 2: Ten-well subsea campaign, 10,000 Monte Carlo iterations per scenario. Well-work scope is WBS 3+4+7 (vessel spread including mobilisation and WOW, materials, risk); full scope adds running costs, project management, clearance, recycling, monitoring, and carbon.

Metric (P10 / P50 / P90)	A: Rig + Milling	B: LWIV + PWC	Δ (P50)	Reduction
Well-work cost (\$M)	93.7 / 102.5 / 113.2	29.1 / 31.3 / 34.1	71.2	−69.4%
Full-scope cost (\$M)	124.5 / 135.5 / 148.9	45.1 / 48.2 / 51.9	87.3	−64.4%
Campaign duration (days)	257 / 282 / 313	144 / 155 / 169	127	−45.1%
Waiting on weather (days)	24.8 / 43.3 / 68.2	10.8 / 19.8 / 31.5	23.5	−54.3%
Non-productive time (days)	2.6 / 7.8 / 16.4	2.2 / 6.3 / 11.9	1.5	−19.2%
Fleet emissions (tCO ₂ e)	16,011 / 16,760 / 17,562	4,202 / 4,416 / 4,648	12,344	−73.7%




Figure 3: Full-scope campaign cost across 10,000 iterations per scenario: probability density (left) and cumulative S-curves with P10/P50/P90 markers (right). The distributions are essentially disjoint.

The advantage decomposes into four reinforcing drivers: the spread-rate differential (\$350k versus \$180k per day); the Phase-2 baseline (a 10-day milling median versus 3-day PWC); the hotter learning curve on the standardised PWC tool set compounding across the batch; and weather exposure, which is not additive but multiplicative with duration — the rig campaign is on location 82% longer and samples storms accordingly, accruing 43.3 P50 WOW days against 19.8. This last driver is precisely the coupling static planning omits, and it is why the model's duration advantage for the rigless campaign (45.1%) exceeds the static benchmark's (41.8%).

6.2 Sensitivity and comparative dynamics

Figure 4 summarises the comparative ensembles. Campaign timing is a first-order lever: an identical work scope started in winter roughly doubles P50 weather downtime. Phase-0 diagnostics nearly halve NPT, paying for the logging run on the first avoided stuck string. The learning-rate sweep is monotonic and steep — one-at-a-time elasticities put campaign duration at 2.1 with respect to the learning rate, 0.7–0.8 for WOW with respect to storm scale, and 0.85 for well-work cost with respect to campaign size. A structural finding merits regulatory attention: under batching,

no well reaches final abandonment for months, then all complete within weeks (Figure 1). An interim progress audit reading zero completed wells would misread an on-schedule phased campaign as stalled; the completion-curve shape should be communicated to compliance stakeholders in advance.




Figure 4: Comparative ensembles (P50 dot, P10–P90 bar): seasonality, storm scale, and Phase-0 diagnostics (left); P50 campaign duration versus crew learning rate with P10–P90 band (right).

7. What an Ensemble Prices: Seeds versus Inputs

A Monte Carlo band answers a narrower question than it appears to. Varying only the random seed with all input parameters fixed prices execution uncertainty — weather, downhole events, integrity-test outcomes — conditional on the inputs being correct. It says nothing about whether the day rates, learning rates, or risk probabilities are themselves right, and a tight seed-only band around a mis-parameterised model is precision without accuracy.

We therefore label every reported band with its scope. On the reference campaign, the seed-only P10–P90 well-work cost band spans \$6.4M. Stacking plausible input ranges — spread rates $\pm 15\%$ (lognormal, reflecting charter-market movement), learning rate ± 0.04 absolute, carbon price uniform on \$50–130/tCO₂e, and metocean severity $\pm 20\%$ — widens the band to \$13.4M, a factor of 2.1. The stacked band is the appropriate contingency basis; the seed-only band, presented alone in a budget, reads as false precision. Two further reporting disciplines follow from the same principle: every displayed result carries the seed (or recorded seed base) that reproduces it, and tail percentiles used for contingency are taken from the 10,000-iteration ensemble, where the bootstrap 90% confidence interval on the P50 cost estimate is 17 $\times$ narrower than at 50 iterations.

Table 3: Uncertainty scope on the reference rigless campaign (10-well, full-scope cost). The band must state what it varies.

Ensemble scope	P10 (\$M)	P50 (\$M)	P90 (\$M)	Band (\$M)
Seeds only (execution luck)	45.5	48.2	51.9	6.4
Seeds + input ranges	41.6	47.7	55.0	13.4

8. Discussion and Limitations

The headline economics — rigless batched PWC at roughly one-third the cost and half the duration of rig-based sequential milling — are consistent in direction and magnitude with the probabilistic rigless-P&A literature [8] and with published strategic comparisons, but the model's contribution is the decomposition: it separates the spread-rate effect, the technology baseline, the learning-curve compounding, and the weather–duration feedback, and it shows that the last of these systematically penalises long campaigns beyond what static estimates carry. For operators this argues that campaign strategy and start-date selection are not scheduling details but first-order economic decisions; for regulators it explains why phased campaigns produce completion curves that look stalled when they are not.

Limitations are deliberate and disclosed. The lease geometry, metocean statistics, day rates, and risk tables are stylised and specification-derived, not calibrated to a specific basin; validation therefore establishes model mechanics and comparative behaviour, not predictive accuracy for a named asset. Deployment requires real well schematics, hindcast metocean series, broker day rates, and re-validation against completed campaign outturns through the same MAPE gate. The model also excludes multi-vessel parallelism within a pass, port logistics, and supply-chain constraints on tool availability — extensions the agent architecture accommodates naturally. Nothing in this paper constitutes investment, provisioning, or regulatory advice.

9. Conclusion

We have presented a fully specified, twice-implemented, and twelve-check-validated agent-based model of offshore well decommissioning campaigns. Its central claim is methodological as much as economic: the costs that dominate P&A outturns arise from couplings — weather with duration, learning with batching, integrity failures with cement quality — and those couplings are exactly what deterministic planning removes. In the reference campaign the couplings are worth two-thirds of the cost and nearly half the schedule. A model that carries them, states what its uncertainty bands vary, and reproduces from a recorded seed is, we argue, the minimum standard for defensible decommissioning provisioning.

References

- [1] Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, 3(4), 122–128.
- [2] Crawford, J. R. (1944). *Learning Curve, Ship Curve, Ratios, Related Data*. Lockheed Aircraft Corporation.
- [3] Clark, C. E. (1962). The PERT model for the distribution of an activity time. *Operations Research*, 10(3), 405–406.
- [4] Malcolm, D. G., Roseboom, J. H., Clark, C. E., & Fazar, W. (1959). Application of a technique for research and development program evaluation. *Operations Research*, 7(5), 646–669.
- [5] Bureau of Safety and Environmental Enforcement. 30 CFR Part 250, Subpart Q — Decommissioning Activities, § 250.1715: How must I permanently plug a well?

- [6] North Sea Transition Authority (2024). Well Decommissioning Guidelines and well abandonment phase definitions (AB1–AB3).
- [7] Kaiser, M. J. (2019). Decommissioning Forecasting and Operating Cost Estimation. Gulf Professional Publishing.
- [8] Moeinikia, F., Ford, E. P., Lohne, H. P., Arild, Ø., Majoumerd, M. M., & Fjelde, K. K. (2014). A probabilistic methodology to evaluate the cost efficiency of rigless technology for subsea multiwell abandonment. SPE-169203-MS, SPE Bergen One Day Seminar.
- [9] Marsaglia, G., & Tsang, W. W. (2000). A simple method for generating gamma variables. *ACM Transactions on Mathematical Software*, 26(3), 363–372.
- [10] Offshore Energies UK (2023). Decommissioning Insight Report.
- [11] Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl. 3), 7280–7287.
- [12] Vose, D. (2008). *Risk Analysis: A Quantitative Guide* (3rd ed.). Wiley.