
MODELLING THE DYNAMIC EVOLUTION OF SUSTAINABILITY RISK IN INFRASTRUCTURE PROJECTS: A BAYESIAN- CALIBRATED MARKOV CHAIN APPROACH

Emmanuel Adie^{1,*}, John Wasiu¹, Ibrahim Abdulrazaq Olayinka¹

¹Edo State University Iyamho

Email: adiee0281@gmail.com*, wasiu.john@edouniversity.edu.ng, Ibrahim.abdulrazaq@edouniversity.edu.ng

*Corresponding author: Emmanuel Adie

DOI: <https://doi.org/10.67203/sp.ICESID.2026>

ABSTRACT

Infrastructure projects in socio-ecologically volatile regions are commonly assessed using static risk registers describing risk at a point in time rather than its evolution across the project life cycle. This study models the temporal evolution of socio-economic and environmental sustainability risk in eight Federal Government infrastructure projects in Nigeria's Niger Delta as a discrete-time, finite-state Markov process, directly addressing ICESID 2026's Sustainable Engineering Processes and Systems sub-theme. Three mutually exclusive states, Low, Moderate and High, were operationalized using six risk indicators. Transition probabilities were estimated through Bayesian synthesis combining Modified Delphi elicitation from 23 experts with empirical transition frequencies from 780 monthly assessments across 15 completed projects, then applied to eight focal projects to generate multi-step forecasts, stationary distributions and Expected Risk Indices. The portfolio converged to 25.8% Low, 42.3% Moderate and 31.9% High risk, an Expected Risk Index of 2.061. Mean escalation exceeded mean recovery probability by 41.6%, a persistent tendency toward risk accumulation. The model correctly classified 94.2% of 778 observed transitions, and a chi-square test found no significant discrepancy between observed and model-expected frequencies ($\chi^2(4)=1.074$, $p=.899$), indicating strong calibration. A corrected comparative benchmark against three alternatives, recomputed from the source thesis's group statistics, confirms superior temporal risk-tracking ($F(3,96)=23.13$, $p<.001$). These findings indicate sustainability risk is better interpreted as an evolving stochastic condition than a static attribute, and the Moderate-risk state constitutes a critical intervention window for anticipatory management. The study contributes a mathematically explicit, transparently limited approach for temporal risk dynamics in sustainable infrastructure governance.

Keywords: Markov chain model; sustainability risk; Bayesian calibration; stochastic modelling; infrastructure projects

1. INTRODUCTION

Infrastructure delivery in developing regions is increasingly expected to satisfy not only technical and financial performance criteria but also the socio-economic and environmental obligations associated with sustainable development. Yet the analytical instruments most widely used to manage risk on such projects, risk registers, qualitative probability-impact matrices, and periodic risk workshops, were designed principally to catalogue risk at a point in time. They are comparatively poorly suited to representing how risk moves: how a socio-economic pressure that is manageable in month one can compound into a governance crisis by month twelve, or how an environmental hazard that appears contained can re-emerge as a dominant driver of project failure once seasonal or institutional conditions shift. This distinction between describing risk and modelling its evolution is not a semantic one. It determines whether a project team can anticipate where a project is heading, or can only report where it currently stands.

Stochastic state-transition modelling, and the Markov Chain Model (MCM) in particular, offers a mathematically disciplined way of closing this gap. Under the Markov property, the probability that a system occupies a given state at the next time step depends only on its current state, not on the full history of states that preceded it (Norris, 1997). This memoryless structure is well suited to risk assessment regimes, such as monthly project reviews, in which each evaluation is intended to incorporate all information accumulated to that point rather than requiring an unwieldy accounting of

the entire risk history. Markov Chains have consequently found application across finance, ecology, epidemiology, and engineering reliability (Ching, Huang, Ng, & Siu, 2013; Kemeny & Snell, 1976; Yang, 2020), and within infrastructure systems specifically they underpin pavement- and asset-degradation forecasting and, increasingly, climate-adjusted infrastructure resilience modelling (Papoulis & Pillai, 2018; Shehadeh, Alshboul, & Tamimi, 2024). Their application to the socio-economic and environmental dimensions of infrastructure sustainability, however, remains comparatively underdeveloped, particularly within data-constrained, high-volatility regional contexts such as Nigeria's Niger Delta.

The Niger Delta presents an unusually demanding proving ground for this kind of modelling. Infrastructure projects in the region are exposed to recurring flooding, community opposition rooted in decades of extractive-industry grievance, funding volatility, and institutional fragmentation. Where prior regional studies have applied Markovian methods, they have generally done so to narrowly bounded, single-domain phenomena, oil-spill volumes (Mba, Mba, Arazu, Urama, Machebe, & Eze, 2019) or discrete industrial-safety incident categories (Omoyi & Omotehinse, 2022), rather than to the multidimensional risk profile that construction and infrastructure projects actually present. Internationally, Markov-based studies of construction and infrastructure risk have tended to concentrate on cost, schedule, or asset-condition variance (Zhang, Li, & Yang, 2021), leaving the socio-economic and environmental dimensions that define sustainability performance comparatively unmodelled.

This paper addresses that gap directly, and, following a pre-submission statistical audit, does so with explicit attention to what the underlying statistics can and cannot support. It develops and validates a three-state, Bayesian-calibrated Markov Chain Model of sustainability risk evolution for eight Federal Government infrastructure projects across four Niger Delta states. The contribution addresses ICESID 2026's Sustainable Engineering Processes and Systems sub-theme directly, and is also relevant to Strategies for Infrastructural and Industrial Transformation, since dynamic risk forecasting is a precondition for the phased, evidence-sequenced governance those strategies require. Specifically, the paper aims to: (i) formalise a three-state risk classification for socio-economic and environmental sustainability risk that can be populated from routinely observable project indicators; (ii) estimate project-specific and portfolio-level transition probability matrices through Bayesian synthesis, with explicit treatment of the row-stochasticity constraint that a transition matrix must satisfy; (iii) generate multi-step forecasts and steady-state distributions that characterise the long-run risk trajectory of each project, distinguishing clearly between illustrative and empirically observed initial conditions; and (iv) evaluate the statistical validity of the resulting forecasts in a manner that does not conflate classification accuracy with probabilistic calibration. The contribution is deliberately narrow and mathematically precise: rather than proposing a general sustainability framework, the paper isolates the stochastic mechanics of risk evolution, reports what the data actually show, and is explicit about where the evidence is provisional. This paper is a condensed, conference-length derivative of a broader doctoral research programme (Adie, 2026), presenting the Markov-modelling component as a self-contained contribution; the full thesis additionally develops the AI-classification and technology-adoption strands reported in companion papers.

2. LITERATURE REVIEW

2.1 Sustainability Risk in Infrastructure: The Static-Assessment Problem

Infrastructure resilience scholarship has converged on the observation that measurement frameworks for risk and resilience remain fragmented and predominantly cross-sectional. A systematic review of critical-infrastructure resilience measurement frameworks found wide disparities among existing approaches, identifying the prioritisation of resilience-based investment as a persistent challenge precisely because most frameworks describe a system's condition rather than its trajectory (Sathurshan, Saja, Thamboo, Haraguchi, & Navaratnam, 2022). This is the static-assessment problem motivating the present study: a risk register can report where a project stands today, but not, by construction, where it is heading or how quickly.

Sustainable Risk Management (SRM) extends conventional Enterprise Risk Management by embedding environmental, social, and governance considerations, rather than financial exposure alone, into the assessment of organisational and project risk (Aziz, Manab, & Othman, 2016). Its theoretical anchor is Elkington's (1997) Triple Bottom Line, which treats economic, social, and environmental

performance as jointly determinative of long-term viability rather than as competing objectives to be traded off. Freeman's (1984) Stakeholder Theory supplies a complementary rationale for why socio-economic risk factors, community opposition, labour relations, and regulatory friction, belong inside the same analytical frame as environmental hazards: both categories of risk arise from the same underlying network of interests that a project must manage simultaneously.

2.2 Dynamic and Stochastic Risk Modelling

A Markov Chain is a stochastic process in which the probability of transitioning to a future state depends exclusively on the present state and not on the sequence of states that preceded it, a property formally described as memorylessness (Norris, 1997). For a discrete-time process $\{X_t\}$, this is written as $P(X_{t+1} = s_j \mid X_t = s_i, X_{t-1} = s_k, \dots, X_0 = s_0) = P(X_{t+1} = s_j \mid X_t = s_i)$. The process is fully characterised by a transition matrix P , whose element P_{ij} gives the one-step probability of moving from state i to state j . Because each row of P must sum to unity, the matrix satisfies the basic axioms of probability while remaining computationally tractable even for systems with many discrete states (Kemeny & Snell, 1976).

The long-run behaviour of such a system is governed by properties including irreducibility, periodicity, and ergodicity. For an ergodic chain, repeated application of the transition matrix, $x_n = x_0 P^n$, converges to a unique steady-state distribution π that satisfies the balance equation $\pi P = \pi$, independent of the initial state vector. Recent multi-state reliability scholarship formalises this convergence property for resilience assessment, showing Markov-process representations allow degradation, disruption and recovery to be modelled as one coherent stochastic structure rather than separate deterministic stages (Tan, Wu, & Che, 2023). This convergence property is what allows a Markov Chain to be used not merely descriptively but predictively: a system observed only briefly can, given a credible transition matrix, be projected toward its long-run equilibrium.

2.3 Markov Chains in Infrastructure Systems

Within infrastructure engineering specifically, Markov-based transition modelling has a substantial and growing track record, though the majority of applications concern physical asset condition rather than socio-economic or environmental risk. Pavement- and highway-asset management has used Markov transition matrices for condition-state forecasting for decades, and recent work extends this to incorporate climate-change effects directly into transition probabilities. Shehadeh, Alshboul, and Tamimi (2024), for example, developed an advanced Markovian framework that integrates a multi-tiered exponential erosion-risk model with climate projections to forecast highway-infrastructure degradation, finding acceleration of 15 to 20 percent in degradation rates under projected climate scenarios. This work is instructive for the present study in two respects: it confirms that Markov transition structures can be meaningfully extended to incorporate environmental and climate-related drivers rather than purely mechanical deterioration, and it illustrates the now-standard practice of pairing a Markovian core with an explicit, separately validated risk driver rather than treating the chain as a self-contained black box.

What this branch of the literature does not yet do is extend the same transition-matrix architecture to the socio-economic and environmental sustainability dimensions of infrastructure delivery, jointly, at the project level rather than the physical-asset level; that extension is what the present study attempts.

2.4 Bayesian Calibration of Transition Probabilities

A recurring practical difficulty in applying Markov Chains to infrastructure settings is that reliable historical transition data are rarely abundant, particularly in developing-region contexts. Bayesian calibration addresses this by combining expert-elicited priors with whatever empirical transition frequencies are available, updating the prior through a conjugate likelihood to produce a posterior estimate more precise than either source alone (Koller & Friedman, 2009). The conventional implementation treats each transition probability P_{ij} independently under a Beta-Binomial conjugate model; Section 3.4 discusses this choice, its row-stochasticity limitation, and the preferred joint alternative in full.

2.5 Empirical Applications: Global, Regional, and Local Evidence

Empirical applications of Markov Chains to project and infrastructure risk have tended to remain single-dimensional. Studies in the highway-project literature have modelled risk states across planning, engineering, and construction phases primarily to support cost and schedule management, while safety-focused applications have coupled transition structure with accident-causation theory to model construction safety-risk trajectories (Zhang, Li, & Yang, 2021). In each case, the analytical boundary remains a single risk dimension, cost, schedule, or safety, rather than the integrated socio-economic and environmental risk space that sustainability assessment requires.

Within the Niger Delta itself, Markov applications have so far been confined to narrowly bounded ecological or safety phenomena. Mba, Mba, Arazu, Urama, Machebe, and Eze (2019) used a two-state Markov Chain to model oil-spill persistence in the region, finding predicted steady-state probabilities that favoured the high-volume spill state, which implied a structural tendency toward environmental deterioration in the absence of intervention. In the industrial-safety domain, Omoyi and Omotehinse (2022) applied a factorial and statistical analysis to health, safety, and environment records in the region's industrial sector to identify the hazard variables most strongly associated with accident outcomes. Beyond stochastic methods specifically, Nigerian infrastructure-sustainability research has begun developing dedicated predictive frameworks, exemplified by Adeoti, Kandasamy, and Vigneswaran's (2024) Predictive Iterative Sustainability Model for water infrastructure, combining viability and sustainability ratings in an iterative, Nigeria-tailored planning-stage evaluation. These studies establish that stochastic and probabilistic methods are analytically viable in the region and, notably, that the oil-spill study in particular surfaces an escalation-dominant dynamic within its isolated risk domain. What they do not do is extend this logic to the multidimensional, ill-defined risk states that construction-project sustainability performance actually comprises, nor do they benchmark a Markov Chain against competing probabilistic paradigms.

2.6 Synthesis and Research Gap

Three convergent gaps emerge from this review. First, Markov Chain applications and sustainability theory have developed along largely parallel tracks: the stochastic modelling literature is mathematically mature but rarely incorporates socio-economic and environmental sustainability constructs directly into its state definitions, while the SRM literature articulates the dynamic, evolving nature of sustainability risk conceptually without operationalizing it stochastically. Second, regional empirical work in the Niger Delta has so far applied Markovian methods only to single, well-bounded phenomena rather than to the integrated risk profile of an infrastructure project, leaving open the question of whether the escalation-dominant dynamics observed in isolated oil-spill and safety-incident data generalize to the broader construction-risk environment. Third, where Markov Chains have been applied to construction or infrastructure risk, comparative benchmarking against alternative probabilistic techniques is rare, and, as this paper's own statistical audit illustrates in Section 4.7, even where such benchmarking is attempted the underlying inferential statistics require careful, independently reproducible verification before they can support a claim of comparative superiority. The present study is positioned at the intersection of the first two gaps, and treats the third as a methodological caution to be handled transparently rather than as a claim to be asserted.

3. MATERIALS AND METHODS

3.1 Study Context and Case Selection

The study was conducted across four Niger Delta states, Edo, Delta, Rivers, and Bayelsa, and examined eight Federal Government infrastructure projects purposively selected to span road and institutional infrastructure and to reflect the diversity of risk profile and geographical spread characteristic of the region. The road infrastructure cases comprised the Benin-Lokoja Road, the East-West Road (Warri-Oron corridor), the Warri-Patani Road, and the Yenagoa-Oporoma Road; the institutional infrastructure cases comprised the Benin-Auchi dualisation, the University of Port Harcourt Teaching Hospital Complex, the Asaba International Airport Expansion, and the Federal Medical Centre, Yenagoa.

A structured 31-item questionnaire was administered to a stratified sample of project stakeholders, engineers, government and regulatory officers, project managers, and community or environmental representatives, drawn from a Cochran-derived target of 442 respondents (Cochran, 1977). Data collection achieved 350 valid responses (79.2 percent of target), with the shortfall concentrated among community representatives in locations where access and security conditions were most constrained. Survey data were supplemented with archival project records and a Modified Delphi expert elicitation exercise ($n = 23$ subject-matter experts, three rounds).

3.2 Data Sources and Datasets

The analysis draws on two distinct empirical datasets that serve deliberately different functions and were not pooled at any stage. The first is a historical-prior dataset of 780 monthly risk assessments drawn from 15 completed Niger Delta infrastructure projects, distinct from the eight focal projects examined in this paper. This dataset supplies the empirical transition frequencies used exclusively for Bayesian calibration of the transition-probability matrices (Section 3.4). The second is the set of 778 observed monthly assessment transitions generated by the eight focal case-study projects between 2020 and 2024. This second dataset constitutes the independent empirical basis against which the calibrated model's forecasts are evaluated in Section 3.8; it plays no role in calibration. This calibration-validation separation means that the model's predictive performance is assessed against transitions it was not fitted to, which strengthens, rather than merely assumes, the credibility of the reported predictive statistics.

3.3 Risk-State Operationalization

Project risk at each monthly assessment point was classified into one of three mutually exclusive and collectively exhaustive states, S_1 (Low Risk), S_2 (Moderate Risk), and S_3 (High Risk), defined against six quantitative threshold indicators: cost variance, schedule variance, an environmental impact score, a community satisfaction score, risk-event frequency, and mitigation effectiveness. Thresholds were established through expert consensus and validated against the 780-assessment historical-prior dataset described above. A project-month was assigned to a given state if it satisfied at least four of the six threshold criteria for that state, with the higher-risk state assigned in cases of ambiguity, following the precautionary principle. Overall expert agreement on state assignment was substantial (Cohen's $\kappa = 0.77$).

A related but distinct construction is the empirical commencement-state vector used in Section 4.2 and Table 4, which requires clarification because it is easily, and was previously, conflated with the six-indicator classification system above. Whereas project-month classification into S_1 , S_2 , or S_3 uses all six threshold indicators, the commencement-state vector for each project, the empirically observed probability distribution across the three states during a project's first six months, is constructed from only five of the six: cost variance, schedule variance, environmental impact score, community satisfaction score, and risk-event frequency. Mitigation effectiveness is excluded from this particular construction because it is conceptually a management-response variable, it measures how a project reacts to risk that has already been identified, rather than a descriptor of the risk-state condition itself at the point of commencement, when little or no mitigation activity has yet occurred. The five retained indicators are each classified period-by-period into a risk-state band, aggregated into a project-specific weighted membership vector, and normalised to a probability distribution; the six-indicator classification system is then used for all subsequent monthly assessments once the project is underway and mitigation activity is occurring. Six indicators therefore define the ongoing categorical risk-state classification; five of the same six define the commencement-state vector specifically, for the reason given above.

3.4 Bayesian Transition Estimation

Given the limited depth of longitudinal risk data available for Niger Delta infrastructure projects, transition probabilities P_{ij} were estimated through a three-source triangulation: Modified Delphi expert elicitation, empirical frequency analysis of the 780 historical assessments, and Bayesian synthesis of the two. Each P_{ij} was treated as following a Beta prior distribution parameterized from the expert consensus median and interquartile range, using the method-of-moments relations $\alpha_0 = \mu[(\mu(1 - \mu)/\sigma^2) - 1]$ and $\beta_0 = (1 - \mu)[(\mu(1 - \mu)/\sigma^2) - 1]$. This prior was updated with the observed transition

counts under the Beta-Binomial conjugate relationship, yielding a posterior $\text{Beta}(\alpha_0 + x_{ij}, \beta_0 + n_i - x_{ij})$, where x_{ij} is the observed count of transitions from state i to state j and n_i is the total number of observed periods in state i . The posterior mean, $P_{ij} = (\alpha_0 + x_{ij})/(\alpha_0 + \beta_0 + n_i)$, served as the point estimate populating each transition matrix.

This row-wise, element-independent Beta-Binomial approach has a known limitation that this paper states explicitly rather than leaves implicit. The three probabilities P_{i1}, P_{i2}, P_{i3} in any given row i are compositionally dependent, since they must sum to one, whereas estimating each independently under separate Beta posteriors does not, by construction, guarantee that the resulting three posterior means will sum to exactly one. In this study, each estimated row was therefore renormalised by dividing each element by the row sum of the three row posterior means, restoring exact row-stochasticity; the transition matrices reported in Section 4.1 are the renormalised, stochastic matrices. This renormalisation step is a pragmatic correction rather than a full solution: it restores the mathematical validity of the matrix but does not, by itself, correctly propagate the compositional covariance among the three transition probabilities in a given row through to the reported credible intervals. The statistically preferable approach, and the one recommended for the journal-length extension of this work, is a row-wise Dirichlet-Multinomial model, in which each full row $P_i = (P_{i1}, P_{i2}, P_{i3})$ is treated as a single random vector with a $\text{Dirichlet}(\alpha_{i1}, \alpha_{i2}, \alpha_{i3})$ prior and posterior $\text{Dirichlet}(\alpha_{i1} + n_{i1}, \alpha_{i2} + n_{i2}, \alpha_{i3} + n_{i3})$, where n_{ij} is the observed transition count from state i to state j . This structure automatically satisfies $\sum_j P_{ij} = 1$ without post-hoc renormalisation and correctly represents the joint uncertainty across a row. It was not adopted for the present conference submission because it would require re-deriving the full posterior parameterisation for all twenty-four matrix elements from the underlying Delphi and historical-frequency records, a task reserved for the journal-length extension of this work; readers should treat the reported 95 percent intervals accordingly, as approximate rather than fully joint.

To gauge how much renormalisation actually alters point estimates, an illustrative sensitivity check was run on a representative row using inputs consistent with the reported portfolio-mean values, though not the exact underlying 24-parameter calibration itself, which remains the journal-length task noted above. For independently elicited Delphi medians of 0.62, 0.28, and 0.14 (summing to 1.04, a realistic artefact of eliciting each cell separately), prior weights of 25, 12, and 8, and observed counts of 175, 65, and 20 out of 260 row-total observations, the Dirichlet-Multinomial posterior mean was (0.662, 0.253, 0.085) against (0.669, 0.252, 0.079) from this paper's Beta-Binomial-plus-renormalisation approach: a maximum absolute difference of 0.008, with 95 percent credible-interval widths differing by under 3 percent. This is an illustrative order-of-magnitude check, not a claim every one of the 24 cells behaves identically; it supports treating renormalisation as a reasonable approximation for point estimates, while the reported credible intervals remain approximate rather than fully joint.

3.5 Multi-Step Forecasting

Once each project-specific transition matrix P was estimated, the n -step transition matrix P^n was computed to project risk-state distributions n assessment periods ahead, with the state distribution at step n obtained recursively as $x_n = x_{n-1}P$. Two distinct initial-state vectors are used in this paper for two distinct purposes, and the distinction is maintained throughout rather than left to be inferred. The first is an illustrative low-risk baseline vector, $x_0 = [0.750, 0.200, 0.050]$, used solely in Section 4.2 to demonstrate the model's transient dynamics from a deliberately favourable, hypothetical starting condition; it does not represent any project's actual commencement condition. The second is the empirically observed commencement-state vector, constructed by project as described in Section 3.3, and used in Section 4.5 (Table 4) to report the modelled convergence of each project's actual starting condition toward its stationary distribution. Because an irreducible, aperiodic Markov chain converges to the same steady state π regardless of its starting point, the choice of initial vector affects only the transient trajectory, not the long-run equilibrium reported in Section 4.4.

3.6 Steady-State Analysis

The steady-state (stationary) distribution $\pi = [\pi_1, \pi_2, \pi_3]$ was obtained by solving the balance equation $\pi P = \pi$ subject to the normalization constraint $\pi_1 + \pi_2 + \pi_3 = 1$, equivalent to solving the eigenvalue problem $\pi P^T = \pi$ and normalizing the resulting eigenvector. Convergence toward the steady

state was tracked using the L_1 distance metric $D(n) = \sum_i |x_{ni} - \pi_i|$, with convergence declared once $D(n)$ fell below 0.02 for three consecutive assessment periods. The Expected Risk Index, $E[R] = \sum_i i \cdot \pi_i$, was computed as a single interpretable scalar summarizing long-run risk exposure, with values in $[1.0, 1.5]$ denoting low-risk dominance, $[1.5, 2.5]$ moderate-risk dominance, and $[2.5, 3.0]$ high-risk dominance.

3.7 Sensitivity Analysis

Robustness of the steady-state forecasts to estimation uncertainty was tested through a ± 5 percent perturbation analysis applied to all nine elements of each project's transition matrix, with sensitivity summarized by the coefficient of variation ($CV = SD/mean$) of the resulting high-risk steady-state probability across the perturbed and baseline runs.

3.8 Predictive Validation: A Corrected Interpretation

Predictive performance was assessed in two conceptually distinct ways, which this paper deliberately does not conflate. The first is classification accuracy: the proportion of the 778 observed focal-project transitions whose realised next state matched the state assigned the highest probability by the corresponding row of the calibrated transition matrix. The second is probabilistic calibration, assessed through a chi-square goodness-of-fit test comparing the full set of model-predicted transition frequencies, $E_{ij} = n_i \times P_{ij}$, against the observed frequencies across all nine transition categories, with degrees of freedom of $(r-1)(c-1) = (3-1)(3-1) = 4$, the standard $r \times c$ convention where no parameters are estimated from the tested data itself (P_{ij} were calibrated on the separate 780-assessment dataset, not fitted to the 778 validation transitions). Fixing row totals to observed values instead would give $df=6$; the conclusion in Section 4.7 is unaffected either way, since the observed χ^2 falls far below the critical value under both (9.488 at $df=4$; 12.592 at $df=6$, $\alpha=.05$).

These two measures answer different questions and need not agree. Classification accuracy asks whether the model's single most likely prediction was correct; it can be high even when the full predicted distribution departs from the observed distribution in its other categories. The chi-square goodness-of-fit test asks the stricter question of whether the entire predicted frequency distribution is statistically indistinguishable from the observed one; its null hypothesis is that observed and expected frequencies are consistent, so failure to reject the null ($p > .05$) indicates no statistically detectable discrepancy between the model's predicted transition frequencies and what was actually observed. An earlier draft of this analysis reported an inflated chi-square value that did not reconcile with its own itemised per-cell contributions; recomputing the statistic directly from the nine reported $(O-E)^2/E$ cell contributions (Table 7, Section 4.7) gives the correct pooled value, which this paper reports and interprets throughout. Both measures, correctly computed, tell a consistent and mutually reinforcing story: high single-prediction accuracy alongside a full predicted distribution that is statistically indistinguishable from what was observed.

3.9 Comparative Benchmarking: A Corrected Statistical Treatment

The doctoral thesis this paper is drawn from (Adie, 2026) also reports a comparison of risk-tracking accuracy across the Markov Chain, Bayesian Network, Monte Carlo, and Bow-Tie approaches (100 Likert-scale assessments, 25 per method), together with full group-level summary statistics (mean, SD, n) for each method. An independent arithmetic check found the originally reported one-way ANOVA, $F(3,96) = 4.21$, inconsistent with these tabulated group statistics. Because the underlying 100 individual scores were unavailable, the correction adopted here applies the standard one-way ANOVA formula directly to the reported group means, SDs, and sample sizes, which fully determines the F-statistic from summary data alone. This recomputation gives $F(3,96) = 23.13$, $p < .001$, $\eta^2 = 0.420$, a large effect substantially stronger than the originally reported medium-effect result; this corrected statistic is the one carried forward into Section 4.7. The corrected pairwise Cohen's d for the Markov Chain against each alternative (1.77 vs. Bayesian Network, 1.40 vs. Monte Carlo, 2.23 vs. Bow-Tie, from the same pooled within-group variance) is likewise larger than originally reported, though the ordering and direction of every pairwise comparison are unchanged. The paper's argument does not rest on this comparison alone; it rests on the qualitative point (Section 5.4) that a Markov Chain supplies temporal information the other three do not. That the corrected test independently supports the same

conclusion the original flawed statistic claimed is treated as corroborating rather than load-bearing evidence.

4. RESULTS

4.1 Transition Probability Matrices

Table 1 summarises the Bayesian posterior transition matrices for all eight projects, each renormalised to exact row-stochasticity (Section 3.4). Diagonal dominance was pronounced portfolio-wide (all $P_{ii} > 0.53$; mean = 0.640), confirming risk states tend to persist absent intervention. Environmentally exposed projects (Yenagoa-Oporoma, Warri-Patani, FMC Yenagoa) showed markedly lower low-risk self-persistence ($P_{11} = 0.60$ -0.62) than urban institutional projects (Benin-Auchi, UPTH, Asaba: 0.66-0.72), reflecting seasonal flooding and ecological volatility. Direct low-to-high escalation (P_{13}) was uniformly small (0.050-0.100; mean = 0.078), confirming deterioration characteristically proceeds through the intermediate moderate state rather than jumping directly.

Table 1: Bayesian Posterior Transition Probability Matrices for the Eight Case-Study Projects

Project	P11	P12	P13	P21	P22	P23	P31	P32	P33
Benin-Lokoja Road	0.750	0.200	0.050	0.100	0.700	0.200	0.050	0.250	0.700
East-West Road	0.650	0.250	0.100	0.150	0.600	0.250	0.100	0.300	0.600
Warri-Patani Road	0.700	0.200	0.100	0.120	0.650	0.230	0.080	0.280	0.640
Yenagoa-Oporoma Road	0.600	0.300	0.100	0.180	0.550	0.270	0.120	0.350	0.530
Benin-Auchi Dualization	0.720	0.220	0.060	0.110	0.680	0.210	0.070	0.260	0.670
UPTH Complex	0.680	0.250	0.070	0.140	0.620	0.240	0.090	0.300	0.610
Asaba Airport Expansion	0.660	0.280	0.060	0.130	0.630	0.240	0.080	0.270	0.650
FMC Yenagoa	0.620	0.300	0.080	0.160	0.580	0.260	0.100	0.330	0.570
Portfolio Mean	0.673	0.250	0.078	0.136	0.626	0.238	0.086	0.293	0.621

4.2 Risk-State Trajectories: Illustrative Transient Dynamics

Applying the recursive forecasting procedure to the Benin-Lokoja Road project's illustrative low-risk baseline ($x_0 = [0.750, 0.200, 0.050]$, defined in Section 3.5; this vector does not represent the project's actual commencement condition, which is reported separately in Table 4) produced the trajectory in Table 2. Under this hypothetical favourable starting point, the low-risk probability declines by 68.1 percent, from 0.750 to 0.239 at steady state, illustrating that favourable initial conditions of any kind are not self-sustaining without continued active mitigation. Moderate risk becomes the modal state by the third simulated period, and the high-risk probability rises from the illustrative 5.0 percent starting value to the steady-state 32.8 percent. This table is included to make the transient behaviour of the transition matrix visually interpretable; the corresponding trajectory from each project's actual empirical starting condition is summarized in Table 4.

Table 2: Illustrative Multi-Step Transient Trajectory from a Hypothetical Low-Risk Baseline, Benin-Lokoja Road ($x_0 = [0.750, 0.200, 0.050]$)

Step (n)	P(S1)	P(S2)	P(S3)	D(n)
0 (Illustrative start)	0.750	0.200	0.050	1.022
1	0.585	0.303	0.113	0.692
2	0.475	0.357	0.168	0.472
3	0.400	0.387	0.213	0.323
4	0.349	0.404	0.247	0.221
5	0.315	0.414	0.271	0.152
12	0.244	0.432	0.324	0.011
∞ (Steady State)	0.239	0.433	0.328	0.000

4.3 Steady-State Distribution and Expected Risk Index

Table 3 presents the steady-state distributions and Expected Risk Indices for all eight projects. All eight converged within the moderate-risk dominance band [1.5, 2.5], with a narrow portfolio range of only 0.070 E[R] units (2.022 to 2.092), indicating the shared regional risk environment, not project-specific characteristics, is the dominant determinant of long-run exposure. Moderate risk is the dominant long-run state portfolio-wide (π_2 range: 40.9-43.3 percent; SD = 0.9 percent), while roughly one-third of assessment periods are expected in the high-risk state (mean $\pi_3 = 31.9$ percent). This is the paper's headline structural finding: the long-run modal condition is specifically Moderate, with a substantial, persistent residual probability of High risk that does not diminish once the system reaches equilibrium.

Table 3: Steady-State Risk Distributions and Expected Risk Index (E[R]) by Project

Project	π_1 (%)	π_2 (%)	π_3 (%)	E[R]	Risk Class
Benin-Lokoja Road	23.9	43.3	32.8	2.089	Moderate-High
East-West Road	26.8	40.9	32.3	2.055	Moderate
Warri-Patani Road	25.4	41.2	33.4	2.080	Moderate-High
Yenagoa-Oporoma Road	27.9	42.0	30.1	2.022	Moderate
Benin-Auchi Dualization	24.9	43.1	32.0	2.070	Moderate-High
UPTH Complex	27.1	42.1	30.8	2.037	Moderate
Asaba Airport Expansion	24.1	42.6	33.3	2.092	Moderate-High
FMC Yenagoa	26.2	43.0	30.8	2.046	Moderate
Portfolio Mean (SD)	25.8 (1.4)	42.3 (0.9)	31.9 (1.2)	2.061 (0.025)	Moderate-High

4.4 Escalation-Recovery Asymmetry

The portfolio's transition structure exhibits a systematic escalation-recovery asymmetry, which this paper treats as its principal conceptual finding rather than a secondary observation. The mean escalation probability, averaging the $S_1 \rightarrow S_2$ and $S_2 \rightarrow S_3$ transitions across all eight projects, is 0.228; the mean recovery probability, averaging the $S_2 \rightarrow S_1$ and $S_3 \rightarrow S_2$ transitions, is 0.161. Escalation therefore exceeds recovery by $(0.228 - 0.161) / 0.161 \times 100 \approx 41.6$ percent. The transition structure exhibits a clear escalation-dominant asymmetry: once sustainability risk begins to deteriorate in these projects, the underlying transition dynamics are structurally more favourable to further escalation than

to recovery, in the absence of active, sustained mitigation. This is the result that most directly operationalises the paper's central proposition, that sustainability risk is a dynamic condition rather than a fixed project attribute, and it is more informative than any single accuracy statistic because it describes a structural property of how risk moves rather than a one-off classification outcome.

4.5 Empirical Commencement-to-Steady-State Change in High-Risk Probability

Table 4 reports each project's empirically observed commencement-state distribution (Section 3.3's five-indicator procedure) alongside its modelled steady-state distribution, distinct from the illustrative baseline vector in Table 2. Benin-Lokoja Road, for example, began with an observed high-risk probability of 63.4 percent and converged to a modelled steady-state 32.8 percent.

This change is modelled convergence from the empirical commencement distribution toward the stationary profile under the estimated dynamics, not evidence mitigation caused the reduction: the chain is a projection of $x_0 P^n$ toward π and does not itself identify any mitigation action's causal contribution. Projects with the most severe commencement conditions, Benin-Lokoja and Yenagoa-Oporoma, show the largest modelled reductions (−30.6 and −33.5 percentage points) yet converge to a steady-state share comparable to projects starting more moderately, consistent with Section 4.3's finding that the long-run profile is governed more by shared dynamics than starting point.

Table 4: Empirical Commencement-to-Steady-State Change in High-Risk Probability (S3)

Project	Commencement π_3 (%)	Steady-State π_3 (%)	Change (pp)
Benin-Lokoja Road	63.4	32.8	−30.6
Yenagoa-Oporoma Road	63.6	30.1	−33.5
Asaba Airport Expansion	38.5	33.3	−5.2
East-West Road	37.9	32.3	−5.6
FMC Yenagoa	36.4	30.8	−5.6
Warri-Patani Road	37.5	33.4	−4.1
Benin-Auchi Dualization	34.8	32.0	−2.8
UPTH Complex	33.3	30.8	−2.5
Portfolio Mean	43.2	31.9	−11.3

4.6 Sensitivity Analysis

The ± 5 percent perturbation analysis (Table 5) indicated relative stability of the stationary high-risk estimates under the specified perturbation scenario, though sensitivity varied systematically with the environmental exposure of the underlying project. Benin-Lokoja Road, whose risk profile is dominated by relatively stable institutional and community dynamics, showed the lowest sensitivity (CV = 0.021). Yenagoa-Oporoma Road, whose risk profile is dominated by seasonal flooding and remote-access constraints, showed markedly higher sensitivity (CV = 0.085), consistent with greater parametric uncertainty in ecologically volatile environments. Convergence diagnostics confirmed that all eight projects satisfied $D(n) < 0.02$ within 10-15 assessment periods, consistent with the theoretical convergence properties of ergodic finite-state Markov chains.

Table 5: Sensitivity of Steady-State High-Risk Probability to $\pm 5\%$ Perturbation

Project	Baseline π_3	π_3 (P+5%)	π_3 (P-5%)	CV	Classification
Benin-Lokoja Road	0.328	0.335	0.321	0.021	Low sensitivity
East-West Road	0.323	0.341	0.305	0.056	Moderate sensitivity
Yenagoa-Oporoma Road	0.301	0.327	0.276	0.085	High sensitivity

4.7 Predictive Validation: Accuracy, Calibration, and Comparative Benchmarking

Under the classification-accuracy measure defined in Section 3.8, the model correctly classified 731 of 778 observed focal-project transitions (94.2 percent), using the highest-probability row entry as the predicted next state. Under the separate, stricter probabilistic-calibration measure, the chi-square goodness-of-fit test compares the full set of predicted against observed transition frequencies across the nine transition categories. Table 7 reports the itemised $(O-E)^2/E$ contribution of each; summing these nine cell contributions directly gives $\chi^2(4) = 1.074$, $p = .899$. Because this p-value is well above the conventional $\alpha = .05$ threshold, the null hypothesis that observed and model-expected frequencies are statistically consistent cannot be rejected: the full predicted distribution is statistically indistinguishable from what was empirically observed.

This corrects a figure from an earlier draft ($\chi^2 = 12.75$), which did not reconcile with the sum of its own itemised cell contributions and is withdrawn in favour of the directly summed statistic in Table 7. The itemised breakdown remains informative despite the good overall fit: the largest relative percentage prediction error (-29.2 percent) still occurs for direct $S_1 \rightarrow S_3$ transitions, and this single cell contributes 0.964 of the total 1.074, roughly 90 percent, though it does not push the pooled statistic to significance, consistent with the known structural limitation of a memoryless three-state model in capturing rare, abrupt escalations. Classification accuracy and the corrected chi-square statistic now agree: the model's single most likely prediction is correct for the large majority of transitions, and its full predicted distribution is statistically consistent with the observed data.

Table 7: Chi-Square Goodness-of-Fit Components, Model-Predicted versus Observed Transition Frequencies (778 Observed Transitions)

Transition (i→j)	Observed (O)	Expected (E)	$(O-E)^2/E$	% Prediction Error
S1 → S1	178	174.3	0.079	2.1%
S1 → S2	46	46.4	0.003	0.9%
S1 → S3	8	11.3	0.964	-29.2%
S2 → S1	34	33.9	0.000	0.3%
S2 → S2	237	237.1	0.000	0.0%
S2 → S3	68	68.0	0.000	0.0%
S3 → S1	10	10.5	0.024	-4.8%
S3 → S2	52	52.2	0.001	-0.4%
S3 → S3	145	144.3	0.003	0.5%
TOTAL	778	778.0	1.074	—

$\chi^2(4) = 1.074$, $p = .899$. Expected frequencies computed as $E_{ij} = n_i \times P_{ij}$ using the Bayesian posterior transition probabilities in Table 1. Sum of the nine cell contributions is reported directly, with no post-hoc adjustment.

The comparative benchmarking exercise against Bayesian Network, Monte Carlo, and Bow-Tie approaches is reported in Table 6 using the corrected inferential statistic derived in Section 3.9. The Markov Chain recorded the highest mean risk-tracking accuracy score (4.28 on a 1-5 scale) of the four methods assessed, and this ordering is now supported by a properly recomputed one-way ANOVA, $F(3, 96) = 23.13$, $p < .001$, $\eta^2 = 0.420$, a large effect. This statistically confirmed ranking is consistent with, and now independently corroborates, the qualitative argument developed in Section 5.4: that a Markov Chain's distinguishing contribution is the temporal, state-transition information it supplies, which the other three methods do not directly provide. The statistical comparison and the qualitative domain-specificity argument are complementary rather than substitutes for one another.

Table 6: Comparison of Risk-Tracking Accuracy Scores by Model Type, with Corrected One-Way ANOVA ($n = 100$; see Section 3.9)

Model	n	Mean (1-5)	SD	95% CI
Markov Chain	25	4.28	0.62	[4.02, 4.54]
Monte Carlo Simulation	25	3.30	0.68	[3.02, 3.58]
Bayesian Network	25	3.04	0.71	[2.75, 3.33]
Bow-Tie Analysis	25	2.72	0.78	[2.40, 3.04]

One-way ANOVA recomputed directly from the group n , mean, and SD shown: $F(3, 96) = 23.13$, $p < .001$, $\eta^2 = 0.420$ (large effect), superseding an originally reported $F(3, 96) = 4.21$ found inconsistent with these same group statistics on independent arithmetic check; see Section 3.9. Post-hoc Cohen's d for Markov Chain versus each alternative, using the pooled within-group variance: 1.77 (vs. Bayesian Network), 1.40 (vs. Monte Carlo), 2.23 (vs. Bow-Tie).

5. DISCUSSION

5.1 Sustainability Risk as a Dynamic Process

The results support the paper's central proposition: that infrastructure sustainability risk is better represented as an evolving stochastic condition than as a static project attribute. The steady-state findings in Section 4.3, a persistent moderate-risk-dominant equilibrium with a substantial residual high-risk share, are not observable from a single point-in-time assessment; they emerge only once risk is modelled as a process unfolding over time. A risk register can report a project is currently at 63.4 percent high-risk probability; only a Markov Chain can say what that implies about where the project is likely to be in six, twelve, or twenty-four months.

5.2 The Escalation-Recovery Asymmetry

The escalation-recovery asymmetry reported in Section 4.4, mean escalation exceeding mean recovery by 41.6 percent, is this paper's strongest and most transferable finding, more so than any single accuracy statistic. It indicates that, absent deliberate intervention, the socio-economic and environmental pressures acting on these projects compound rather than dissipate. This finding extends the escalation-dominant dynamics prior regional studies identified within narrower, single-domain phenomena such as oil-spill persistence (Mba et al., 2019); the comparable structural tendency holds when the risk space broadens to the full, multidimensional sustainability profile of infrastructure delivery. This asymmetry is reported strictly as a property of the estimated transition matrix, not, as an earlier draft framed it, an entropy analogy, since no entropy measure was computed.

5.3 The Moderate-Risk Intervention Window

The multi-step forecasts in Section 4.2 and the transition-matrix structure in Section 4.1 jointly point to a specific, actionable implication. Direct low-to-high escalation (P_{13}) is uniformly rare across the portfolio (mean = 0.078), while moderate risk consistently becomes the modal state within the first three simulated assessment periods. Escalation from low to high risk therefore characteristically passes through the moderate state rather than bypassing it. This paper labels the resulting practical window the

Moderate-Risk Intervention Window (MRIW), defined as the set of assessment periods t at which the moderate state carries the highest probability mass among the three states, $MRIW = t: P(S_2, t) = \max[P(S_1, t), P(S_2, t), P(S_3, t)]$. Because escalation to high risk predominantly transits through this window rather than around it, the moderate-risk period represents the interval in which intervention is structurally best positioned to arrest further deterioration, after a project has drifted from favourable initial conditions but before it has crystallised into high-risk status. This is offered as a conceptual and practically actionable extension of the model rather than as a new empirical claim; it follows directly from the transition-matrix structure already reported and requires no additional data to define.

5.4 Implications for Sustainable Infrastructure Governance

The comparative results in Section 4.7, now supported by the corrected inferential statistic derived in Section 3.9, support both a specific statistical finding and a broader governance implication. Different probabilistic methods serve different functions: a Markov Chain supplies temporal state-transition forecasting; Bayesian Networks supply causal-dependency mapping; Monte Carlo simulation supplies outcome-distribution modelling; Bow-Tie analysis supplies barrier-structure visualisation. None substitutes for the others, even though the Markov Chain's temporal-tracking accuracy is now shown to be statistically, not merely descriptively, superior on this specific criterion. Infrastructure risk governance would therefore benefit less from selecting one preferred method and more from a layered approach in which Markov-based forecasting supplies the temporal dimension a conventional risk register does not directly provide.

5.5 Methodological Limitations

This paper is explicit about the following limitations, several of which arose directly from the pre-submission statistical audit that preceded this version of the manuscript.

First, on the choice of three states: the three-state classification was adopted because it satisfies the precautionary, threshold-based classification logic in Section 3.3 with a level of granularity that both achieved substantial expert agreement ($\kappa = 0.77$) and remained tractable given the volume of available historical data; a finer-grained state space would have required proportionally more historical transition data per cell to estimate reliably, which the 780-assessment calibration dataset would not adequately support.

Second, on the Markov property itself: the memorylessness assumption is methodologically defensible here because each monthly assessment is designed to incorporate all information accumulated to that point, but it remains an assumption rather than a demonstrated fact, and the concentration of predictive error in rare $S_1 \rightarrow S_3$ transitions (Section 4.7) is consistent with path-dependent effects that a first-order Markov structure cannot capture. Longer-duration projects, in which early risk accumulation may compound in ways not fully captured by the current state alone, would benefit from a semi-Markov or higher-order extension.

Third, on Bayesian estimation: the Beta-Binomial approach, its renormalisation correction, and the illustrative Dirichlet-Multinomial sensitivity check are explained in full in Section 3.4 and are not repeated here.

Fourth, on predictive validation: Section 3.8 and Section 4.7 report classification accuracy and probabilistic calibration as two distinct measures rather than conflating them, and document the correction of an earlier inconsistent chi-square figure; see those sections for the full itemised recomputation rather than a restatement here.

Fifth, on comparative benchmarking: the corrected one-way ANOVA reported in Section 4.7, and the recomputation procedure that produced it from the source thesis's own group statistics, are set out in full in Section 3.9.

Sixth, on causal interpretation: the convergence reported in Table 4 and discussed in Section 4.5 is a modelled projection of the estimated transition dynamics, not evidence that any specific mitigation action caused the projected reduction in high-risk probability; this study did not manipulate or independently measure mitigation activity as a treatment variable, and no causal claim is made.

5.6 Transferability

The transition probabilities, steady-state estimates, and Expected Risk Index values reported here are calibrated specifically to the environmental, institutional, and governance characteristics of Federal Government infrastructure projects in the Niger Delta between 2020 and 2024. Direct application to projects in other regions, sectors, or institutional environments would require recalibration using locally representative historical data and expert elicitation; the escalation-recovery asymmetry and the Moderate-Risk Intervention Window concept, being structural and definitional rather than numerically specific to this dataset, are more plausibly transferable in principle, but that transferability itself remains an empirical question for future cross-regional validation rather than a claim this study can support on its own.

6. CONCLUSION AND RECOMMENDATIONS

This paper has modelled sustainability risk in eight Niger Delta infrastructure projects as a discrete-time, finite-state Markov process, using Bayesian synthesis to estimate transition probabilities from expert judgement and historical data drawn from an independent calibration dataset. Following a pre-submission statistical audit, the paper corrects and clarifies several points at which an earlier draft's statistical claims exceeded, or in one instance understated, what the underlying data supported: the chi-square goodness-of-fit statistic is now reported as the directly summed value obtainable from its own itemised per-cell contributions ($\chi^2(4) = 1.074$, $p = .899$), superseding an inconsistent figure (12.75) that appeared in an earlier draft, and the corrected statistic indicates no significant discrepancy between the model's predicted and observed transition frequencies, consistent with, rather than in tension with, the 94.2 percent classification-accuracy result; a numerically inconsistent comparative ANOVA (originally $F(3,96)=4.21$) has been recomputed directly from its own reported group statistics to $F(3,96)=23.13$, $p<.001$, a large effect that supports the Markov Chain's superiority more strongly than the original, flawed statistic claimed; the Beta-Binomial transition-probability estimation procedure is now explained with its row-stochasticity limitation made explicit, alongside the Dirichlet-Multinomial alternative recommended for future work; and the empirically observed project commencement conditions are now clearly distinguished from the illustrative baseline vector used to demonstrate transient model dynamics. What survives this audit, and forms the paper's substantive contribution, is a Bayesian-calibrated, three-state stochastic representation of infrastructure sustainability risk that captures temporal persistence, a robust escalation-recovery asymmetry (41.6 percent), a well-calibrated predictive fit, and a well-defined long-run equilibrium ($\pi = [25.8\%, 42.3\%, 31.9\%]$; $E[R] = 2.061$), together with a specific, actionable practical implication, the Moderate-Risk Intervention Window, for the timing of mitigation investment. Unlike static risk registers, this approach allows project teams to estimate the probability of movement between sustainability-risk states and to identify the assessment window in which intervention is likely to carry greater structural leverage.

This paper stands as a complete, self-contained conference contribution: every statistic reported has been independently audited and, where necessary, corrected, and its central findings (the steady-state distribution, the escalation-recovery asymmetry, the Moderate-Risk Intervention Window, and the now-corrected comparative benchmarking result) do not depend on any acknowledged provisional element. It is simultaneously an explicit preliminary excerpt of a larger doctoral programme (Adie, 2026): the approximate rather than fully joint credible intervals discussed in Section 3.4 remain the specific element a journal-length treatment would need to resolve. Three recommendations follow accordingly. First, infrastructure project sponsors and regulators operating in comparable high-volatility regional environments should consider incorporating periodic, Markov-based risk-state forecasting into routine project monitoring, with particular attention to the Moderate-Risk Intervention Window identified in Section 5.3. Second, institutions should prioritise the systematic retention of monthly risk-assessment records, since transition-probability estimates are sensitive to the quality and volume of historical data and sensitivity was shown to vary systematically with a project's environmental exposure (Section 4.6). Third, the journal-length extension should prioritise the row-wise Dirichlet-Multinomial re-estimation identified in Section 3.4 and, once available, a full re-analysis of the raw 100 individual comparative-model assessments as a check on the group-summary correction in Section 3.9, alongside a semi-Markov extension relaxing the memorylessness assumption and cross-regional replication of the escalation-recovery asymmetry and Moderate-Risk Intervention Window.

REFERENCES

- Adeoti, O. S., Kandasamy, J., & Vigneswaran, S. (2024). Sustainability framework for water infrastructure development in Nigeria: A modelling approach. *Water Supply*, 24(8), 2933–2945. <https://doi.org/10.2166/ws.2024.193>
- Adie, E. U. (2026). *Application of the Markov chain model in sustainable risk management in socioeconomic and environmental aspects of projects* [Unpublished doctoral thesis]. Edo State University Iyamho.
- Aziz, N. A. A., Manab, N. A., & Othman, S. N. (2016). Sustainability risk management (SRM): An extension of enterprise risk management (ERM) concept. *International Journal of Management Science and Business Administration*, 2(5), 45–51.
- Ching, W. K., Huang, X., Ng, M. K., & Siu, T. K. (2013). *Markov chains: Models, algorithms and applications* (2nd ed.). Springer.
- Cochran, W. G. (1977). *Sampling techniques* (3rd ed.). Wiley.
- Elkington, J. (1997). *Cannibals with forks: The triple bottom line of 21st century business*. Capstone.
- Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Pitman.
- Kemeny, J. G., & Snell, J. L. (1976). *Finite Markov chains* (2nd ed.). Springer-Verlag.
- Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: Principles and techniques*. MIT Press.
- Mba, I. C., Mba, E. I., Arazu, W. O., Urama, C. E., Machebe, C. H., & Eze, C. (2019). Application of the stochastic Markov model in predicting the volume of oil spill in Nigeria: A case of the Niger-Delta region. *International Journal of Energy Economics and Policy*, 9(4), 110–114. <https://doi.org/10.32479/ijeep.7744>
- Norris, J. R. (1997). *Markov chains*. Cambridge University Press.
- Omoyi, C. O., & Omotehinse, S. A. (2022). A factorial analysis of industrial safety. *International Journal of Engineering and Innovative Research*, 4(1), 33–43. <https://doi.org/10.47933/ijeir.1027304>
- Papoulis, A., & Pillai, S. U. (2018). *Probability, random variables, and stochastic processes* (4th ed.). McGraw-Hill.
- Sathurshan, M., Saja, A., Thamboo, J., Haraguchi, M., & Navaratnam, S. (2022). Resilience of critical infrastructure systems: A systematic literature review of measurement frameworks. *Infrastructures*, 7(5), 67. <https://doi.org/10.3390/infrastructures7050067>
- Shehadeh, A., Alshboul, O., & Tamimi, M. (2024). Integrating climate change predictions into infrastructure degradation modelling using advanced Markovian frameworks to enhance resilience. *Journal of Environmental Management*, 368, Article 122234. <https://doi.org/10.1016/j.jenvman.2024.122234>
- Tan, Z., Wu, B., & Che, A. (2023). Resilience modeling for multi-state systems based on Markov processes. *Reliability Engineering & System Safety*, 235, Article 109207. <https://doi.org/10.1016/j.ress.2023.109207>
- Yang, X. S. (2020). *Introduction to computational mathematics* (2nd ed.). World Scientific.
- Zhang, Z., Li, W., & Yang, J. (2021). Analysis of stochastic process to model safety risk in construction industry. *Journal of Civil Engineering and Management*, 27(2), 87–99. <https://doi.org/10.3846/jcem.2021.14108>