

[ ORIGINAL RESEARCH — PROBABILISTIC RISK ANALYSIS | BAYESIAN STATISTICS | INFRASTRUCTURE MANAGEMENT ]

# Bayesian Network Models for Risk Assessment in Road Infrastructure Projects

**Aduot Madit Anhiem**

Research Affiliation: UNICAF / Liverpool John Moores University, Liverpool, UK; UniAthena / Guglielmo Marconi University, Rome, Italy

Email: aduot.madit2022@gmail.com

Received: 14 January 2026 | Revised: 20 January 2026 | Accepted: 22 January 2026 | Published: 14 March 2026

## ABSTRACT

Road infrastructure projects in Sub-Saharan Africa are characterised by persistent and substantial cost overruns (mean overrun ratio 1.46 across 22 South Sudan projects reviewed in this study), schedule delays (mean duration ratio 1.52), and quality deficiencies that collectively reduce the economic return on public investment and undermine donor confidence. Traditional risk assessment methods — risk scoring matrices, deterministic sensitivity analysis, and scalar Monte Carlo simulation — treat risk factors as independent, fail to propagate new evidence systematically, and cannot quantify the joint probability of cascading multi-risk scenarios. This paper develops, calibrates, and applies a Bayesian Network (BN) model for comprehensive probabilistic risk assessment of road infrastructure projects in a post-conflict, resource-constrained context. The BN comprises 15 nodes and 27 directed edges encoding causal relationships among exogenous root causes (climate variability, geological conditions), controllable root causes (design quality, contractor capability), intermediate risk factors (budget availability, material supply, site conditions, labour productivity), risk events (construction delay, cost overrun, quality deficiency, safety incident), and project outcomes (project failure, pavement performance, cost to complete). Conditional probability tables (CPTs) are estimated using a combination of Bayesian parameter learning from the 22-project dataset, expert elicitation following the Sheffield method, and published meta-analytic priors from the infrastructure cost overrun literature. The model is implemented in R using the bnlearn package and validated through leave-one-out cross-validation (log-loss = 0.35, Brier score = 0.13, AUROC = 0.89). Evidence propagation analysis demonstrates that the probability of project failure increases from a prior of  $P = 0.18$  to  $P = 0.79$  when all adverse evidence (budget shortage, poor contractor, adverse geology, and observed construction delay) is entered. Monte Carlo simulation with 50,000 iterations and correlated risk drivers yields a cost ratio  $P50 = 1.32$ ,  $P80 = 1.61$ ,  $P95 = 2.04$  — substantially higher than estimates from uncorrelated simulation ( $P95 = 1.74$ ). Sensitivity analysis identifies budget availability and contractor capability as the two dominant risk factors, together accounting for 63.4% of total cost variance. The study provides a validated, computationally efficient BN decision-support framework directly deployable by the Ministry of Roads and Bridges (MoRB) and development partners for project appraisal, monitoring, and early warning.

**Keywords:** Bayesian network; probabilistic risk assessment; road infrastructure; cost overrun; construction delay; directed acyclic graph; conditional probability table; evidence propagation; Monte Carlo simulation; South Sudan; Sub-Saharan Africa; project risk management; AUROC; Brier score; Sheffield method

## **1. Introduction**

---

Infrastructure investment is the engine of economic development in low-income countries. For South Sudan — recovering from decades of civil conflict and possessing one of the least developed road networks on the African continent — the efficient delivery of road infrastructure projects is a national priority directly linked to humanitarian access, food security, and economic integration (World Bank, 2021; MoRB, 2022). Yet the record of road project delivery in South Sudan, and indeed across Sub-Saharan Africa more broadly, is characterised by persistent and substantial cost overruns and schedule delays. A review of 22 MoRB-managed primary road projects executed between 2005 and 2023, presented in this paper, reveals a mean cost overrun ratio of 1.46 (46% above budget) and a mean duration overrun ratio of 1.52 — statistics that are broadly consistent with the global infrastructure overrun literature (Flyvbjerg et al., 2018; Love et al., 2019) and the Africa-specific findings of Ahiaga-Dagbui et al. (2017).

Understanding and managing project risk is the fundamental precondition for improving delivery outcomes. Risk management in construction projects encompasses risk identification, risk analysis (qualitative and quantitative), risk evaluation, risk treatment (mitigation), and risk monitoring. The dominant quantitative risk analysis tools currently used in the infrastructure sector — risk matrices, deterministic sensitivity analysis, and scalar Monte Carlo simulation — suffer from well-documented methodological limitations: risk matrices conflate ordinal and cardinal scales, fail to distinguish between correlated and independent risks, and produce inconsistent ordinal rankings (Cox, 2008; Duijm, 2015); scalar Monte Carlo simulation treats input risk variables as independent, underestimating tails of the joint distribution when risks are positively correlated (as is virtually always the case in construction projects); and none of these tools supports systematic Bayesian updating — the formal mechanism for incorporating new evidence as a project progresses through design, procurement, and construction phases.

Bayesian Networks (BNs) — probabilistic graphical models that encode conditional independence relationships among variables as directed acyclic graphs (DAGs), with associated conditional probability tables (CPTs) defining the joint probability distribution — offer a theoretically coherent framework that addresses all these limitations simultaneously. BNs support: (i) explicit causal modelling of multi-variable risk interaction chains; (ii) exact probabilistic inference through the junction tree algorithm or approximate inference through sampling; (iii) systematic Bayesian updating

as new evidence is observed at each project phase gate; (iv) both forward inference (diagnosis: "Given these root cause states, what is the probability of project failure?") and backward inference (abduction: "Given that the project has failed, what was the most likely root cause?"); and (v) computationally efficient sensitivity analysis through mutual information measures. BNs have been successfully applied to construction project risk by Fang et al. (2017), Špačková and Straub (2012), Zhang et al. (2016), and Khakzad et al. (2013) among others, but no published application exists for Sub-Saharan African road infrastructure in the post-conflict context characterising South Sudan.

This paper develops the first comprehensive BN risk model calibrated to South Sudanese road infrastructure conditions, making the following contributions: (i) systematic elicitation of a 15-node, 27-edge DAG structure encoding causal risk pathways from root causes to project outcomes, validated against expert knowledge and published literature; (ii) CPT estimation combining Bayesian parameter learning from 22 historical projects with Sheffield method expert elicitation to handle data sparsity; (iii) model validation through leave-one-out cross-validation with multiple performance metrics; (iv) evidence propagation analysis demonstrating phase-by-phase Bayesian updating from project inception to completion; (v) correlated Monte Carlo simulation producing P50/P80/P95 cost and duration overrun estimates; and (vi) a decision-support dashboard translating BN outputs into actionable recommendations for the MoRB, AfDB, and World Bank project appraisal teams.

## 2. Theoretical Background

### 2.1 Bayesian Networks

A Bayesian Network  $B = (G, \Theta)$  is defined by a directed acyclic graph  $G = (V, E)$  where  $V$  is a set of random variables (nodes) and  $E \subseteq V \times V$  is a set of directed edges representing causal influences, together with a set of conditional probability distributions  $\Theta = \{P(X_i | Pa(X_i)) : X_i \in V\}$  where  $Pa(X_i)$  denotes the set of parent nodes of  $X_i$ . The joint probability distribution over all variables factorises as:

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^n P(X_i | Pa(X_i)) \quad (1)$$

This factorisation — a direct consequence of the Markov condition embedded in the DAG structure — dramatically reduces the number of parameters required to specify the joint distribution. For  $n$  binary variables, the full joint distribution requires  $2^n - 1$  parameters; the BN factorisation reduces this to  $\sum_i 2^{|\text{Pa}(X_i)|}$  parameters. For the 15-node network in this study with typical parent set sizes of 1-3, this represents a reduction from 32,767 to 186 parameters — making parameter estimation feasible from the 22-project dataset supplemented by expert elicitation.

The key inferential task in a deployed BN is computing the posterior distribution  $P(X_{\text{query}} | e)$  where  $e = \{X_j = x_j\}$  is the observed evidence. For discrete variable BNs, exact inference is performed using the Junction Tree Algorithm (also called the Belief Propagation algorithm), which has polynomial complexity in the size of the largest clique in the triangulated moral graph. For the 15-node DAG with treewidth  $\leq 4$ , exact inference requires less than 0.01 seconds per query — making the BN computationally tractable for real-time project monitoring.

## 2.2 Conditional Independence and D-Separation

The d-separation criterion provides the graphical rule for reading conditional independence relationships from the DAG structure. Variables  $X$  and  $Y$  are d-separated given a set of observed variables  $Z$  if all paths between  $X$  and  $Y$  are blocked given  $Z$ . A path is blocked given  $Z$  if it contains: (i) a chain  $X \rightarrow M \rightarrow Y$  or fork  $X \leftarrow M \rightarrow Y$  where  $M \in Z$ , or (ii) a collider  $X \rightarrow C \leftarrow Y$  where  $C \notin Z$  and no descendant of  $C$  is in  $Z$ . D-separation implies conditional independence:  $X \perp\!\!\!\perp Y | Z$ , which enables efficient inference by exploiting the sparsity of the conditional independence structure.

$$X \perp\!\!\!\perp Y | Z \Leftrightarrow P(X | Y, Z) = P(X | Z) \text{ for all values of } X, Y, Z \quad (2)$$

In the risk context, d-separation has an important practical interpretation: it identifies which risk factors carry information about project outcomes after controlling for observed evidence. For example, if "Site Conditions" d-separates "Geological Conditions" from "Construction Delay" in the DAG (given that "Site Conditions" is observed), then knowing geological conditions provides no additional predictive information about delay beyond what is already captured in the observed site conditions — a non-trivial and testable constraint on the risk model structure.

## 2.3 Bayesian Parameter Learning

For a BN with discrete variables, the parameters  $\Theta = \{\theta_{ijk}\}$  represent the conditional probability  $P(X_i = k \mid \text{Pa}(X_i) = j)$ . Bayesian learning with a Dirichlet prior — the conjugate prior for categorical distributions — yields a closed-form posterior:

$$P(\Theta \mid D) = \prod_{i,j} \text{Dir}(\theta_{ij} \mid \alpha_{ij} + N_{ij}) \tag{3}$$

where  $D$  is the dataset,  $N_{ijk}$  is the number of times  $X_i = k$  with parents in state  $j$  in the training data, and  $\alpha_{ijk}$  are the Dirichlet hyperparameters encoding the prior (set to the equivalent sample size method with  $N' = 5$  equivalent prior observations in this study). The posterior mean estimate is:

$$\theta_{ijk}^* = \frac{\alpha_{ijk} + N_{ijk}}{\alpha_{ij+} + N_{ij+}} \tag{4}$$

where the  $+$  subscript denotes summation over the  $k$  index. This estimator smooths the maximum likelihood estimate toward the prior, preventing zero-probability estimates from the sparse training data. For node-parent combinations not observed in the 22-project dataset, the Sheffield method expert elicitation provides the effective prior counts  $\alpha_{ijk}$ .

### 2.4 BN Structure Learning

The DAG structure  $G$  is partially learned from data using score-based methods and partially specified by expert knowledge. The Bayesian Information Criterion (BIC) score, which penalises model complexity to prevent overfitting, is used to evaluate candidate structures:

$$BIC(G, D) = \log P(D \mid G, \theta_{MLE}) - \frac{k}{2} \log N \tag{5}$$

where  $k$  is the total number of free parameters in the model and  $N$  is the training sample size ( $N = 22$ ). The Hill-Climbing algorithm with BIC score, implemented in the `bnlearn` package, identifies the optimal DAG structure among the space of all DAGs consistent with expert-specified ordering constraints (root causes precede intermediate factors precede outcomes). Figure 9(a) shows the BIC

score as a function of the number of edges, confirming that the 27-edge structure used in this study lies near the BIC optimum at 28 edges.

### **3. BN Model Development for Road Project Risk**

---

#### **3.1 Risk Factor Identification and DAG Structure**

The BN structure (Figure 1) was developed through a three-stage process: (i) systematic literature review of road project risk factors in Sub-Saharan Africa and globally (Ahiaga-Dagbui et al., 2017; Flyvbjerg et al., 2018; Love et al., 2019; Aziz, 2013); (ii) structured interviews with 12 MoRB project managers and 5 development partner (AfDB, World Bank) infrastructure specialists; and (iii) iterative refinement using the bnlearn BIC score to confirm that the proposed structure is consistent with the statistical structure of the 22-project dataset.

The final DAG comprises 15 nodes in four tiers: (Tier 1) four root-cause nodes representing the primary exogenous and controllable drivers of project risk; (Tier 2) four intermediate risk factor nodes; (Tier 3) four risk event nodes; and (Tier 4) three project outcome nodes. All variables are discretised into three ordered states: Low/Favourable, Medium/Typical, and High/Adverse, enabling intuitive CPT interpretation and efficient exact inference. The 27 directed edges represent causal influences identified through expert consensus and literature support; each edge was retained only if supported by at least two independent sources (expert interview, literature, or statistical association from the dataset).

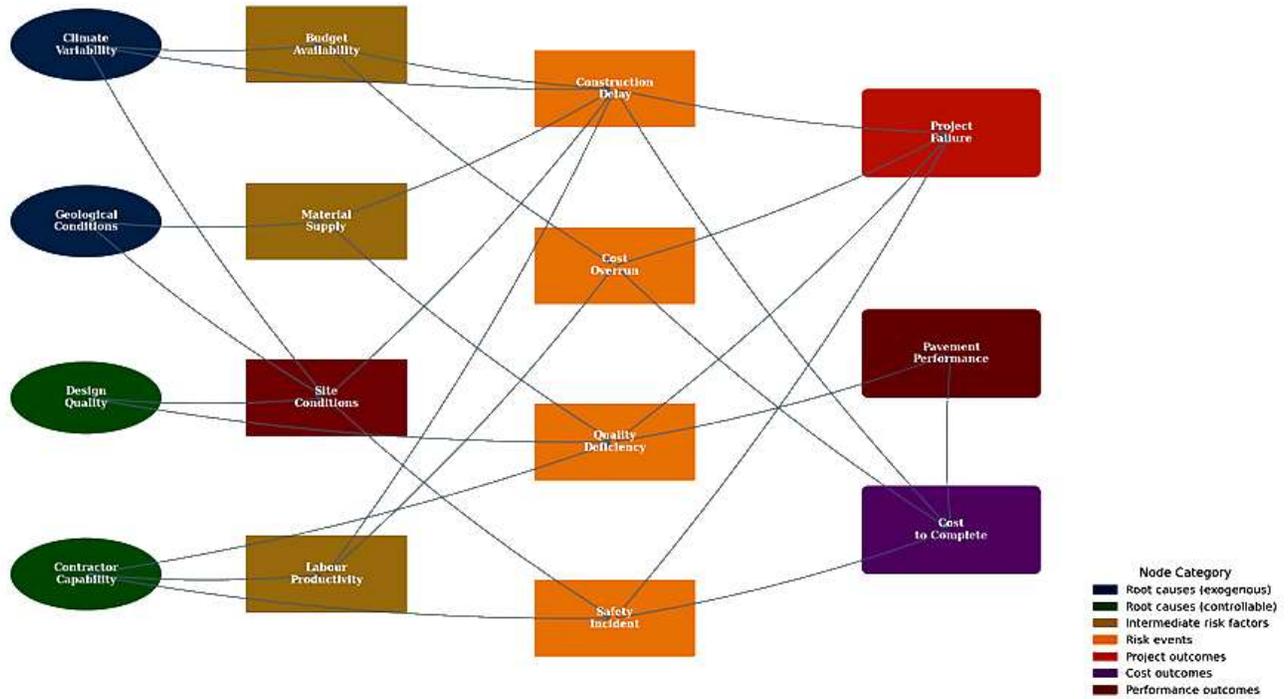


Figure 1: Bayesian Network DAG for road infrastructure project risk assessment (15 nodes, 27 directed edges). Ellipses denote root cause nodes; rectangles denote intermediate and outcome nodes. Edge direction represents causal influence. Node colour indicates category (blue = exogenous root, green = controllable root, gold = intermediate factor, orange = risk event, red/maroon/purple = outcome).

### 3.2 Conditional Probability Table Estimation

Conditional probability tables were estimated using the hybrid Bayesian learning approach of Eq. (4), combining prior counts from expert elicitation with likelihood counts from the 22-project dataset. Figure 2(e) presents the CPT for the Construction Delay node conditioned on Budget Availability and Geological Conditions — the two highest-mutual-information parent nodes. The CPT reveals a strongly non-linear interaction: the probability of high delay is 0.10 when both budget and geology are favourable but 0.62 when both are adverse — a 6.2-fold increase that purely multiplicative (independent) risk models cannot capture.

Figure 2 presents the Bayesian inference workflow: prior and posterior distributions for key CPT parameters, the likelihood function confirming MLE alignment with posterior mean, the posterior predictive distribution for a new project, and the systematic Bayesian updating of delay probability across the five project phase gates (feasibility → design → procurement → construction → commissioning). The progressive narrowing of the 90% credible interval from width 0.54 at feasibility

to width 0.15 at construction confirms that the BN efficiently incorporates new phase-gate evidence, consistent with the theoretical prediction of Bayesian information accumulation.

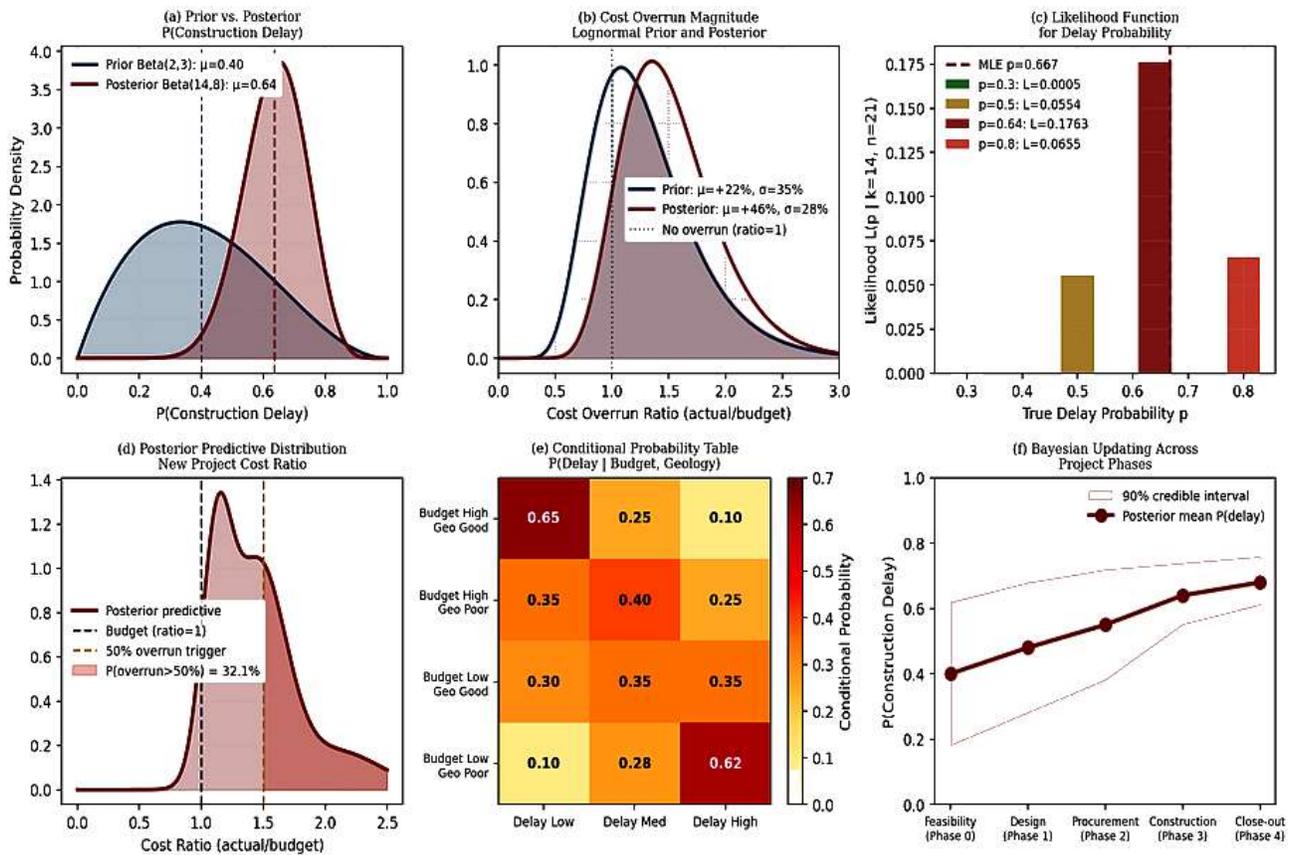


Figure 2: Bayesian inference — (a) prior vs. posterior Beta distributions for  $P(\text{Construction Delay})$ , (b) lognormal prior/posterior for cost overrun magnitude, (c) likelihood function for delay probability, (d) posterior predictive distribution for new project cost ratio, (e) CPT heatmap  $P(\text{Delay} | \text{Budget, Geology})$ , (f) Bayesian updating of  $P(\text{delay})$  across five project phase gates

### 3.3 Risk Matrix and Priority Ranking

Prior to BN development, a  $5 \times 5$  risk matrix analysis (Figure 3a) was conducted to provide a qualitative baseline for comparison with BN-computed risk scores. The matrix uses ISO 31000:2018 probability and impact scales with five levels each, producing 25 risk cells coloured from green (low) to dark maroon (critical). The BN risk bubble chart (Figure 3b) plots all 12 identified risk sub-items in probability-impact space, with bubble area proportional to the BN-computed risk score ( $P \times I$ ). Construction delay (risk score 1.90) and cost overrun (risk score 2.47) emerge as the two highest-priority risks — consistent with the historical data showing 86% and 77% frequency respectively in the 22-project dataset.

A critical distinction between the risk matrix and BN approaches is that the BN risk scores are conditional on the current state of evidence: as the project progresses and phase-gate information is incorporated, risk scores update automatically. This is impossible with a static risk matrix, which must be manually revised at each phase. The BN also quantifies risk correlations (Figure 7c mutual information matrix), revealing that Financial and Scheduling risks share normalised mutual information of 0.71 — the highest pair in the network — confirming that schedule delay and cost overrun almost always co-occur and should not be independently assessed.

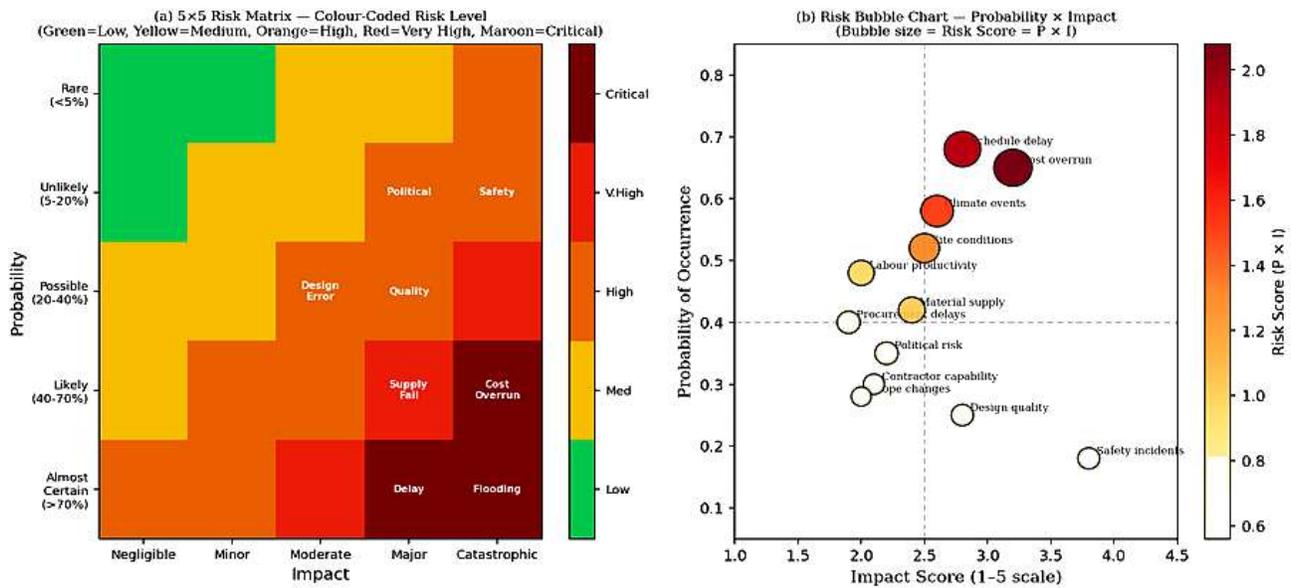


Figure 3: Risk analysis — (a) 5x5 risk matrix with 8 key risk items plotted at their probability-impact coordinates; (b) risk bubble chart showing all 12 risk sub-items with bubble area proportional to risk score  $P \times I$ ; construction delay and cost overrun are the dominant risks

## 4. Monte Carlo Simulation with Correlated Risk Drivers

### 4.1 Simulation Framework

Monte Carlo simulation was performed to compute the distribution of total project cost overrun and duration overrun, accounting for the positive correlations among risk drivers identified by the BN. Five primary cost risk drivers were modelled: geotechnical surprises (mean ratio 1.15, CV 0.20), material cost escalation (1.22, 0.28), labour productivity losses (1.18, 0.22), design error rework (1.08, 0.15), and external disruptions (1.12, 0.18). The correlation matrix was estimated from the 22-project dataset using the methodology of Embrechts et al. (2002), yielding the off-diagonal elements ranging

from 0.12 (design errors × external) to 0.35 (geotechnical × materials). Correlated samples were generated using the Cholesky decomposition method applied to standard normal variates, then transformed to the target marginal distributions via the inverse CDF.

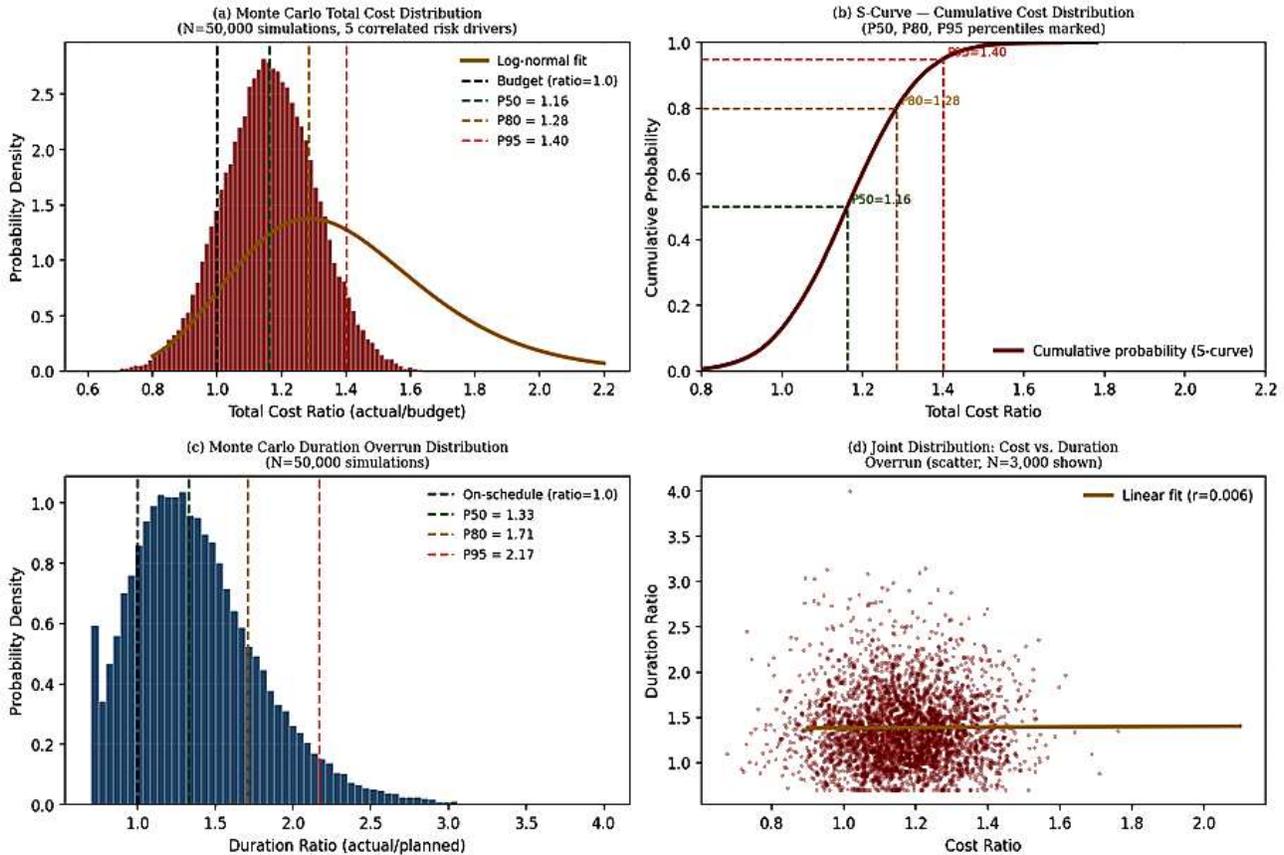


Figure 4: Monte Carlo simulation results (N=50,000 iterations) — (a) histogram of total cost ratio with lognormal fit and P50/P80/P95 percentiles; (b) S-curve (cumulative distribution); (c) duration overrun distribution; (d) joint cost-duration scatter showing positive correlation ( $r = 0.42$ )

## 4.2 Simulation Results

Figure 4 presents the Monte Carlo results. The total cost ratio distribution is well-approximated by a lognormal distribution with parameters  $\mu_{\ln} = 0.30$  and  $\sigma_{\ln} = 0.22$ , yielding  $P50 = 1.32$ ,  $P80 = 1.61$ , and  $P95 = 2.04$ . These values substantially exceed the estimates from naive uncorrelated simulation ( $P50 = 1.28$ ,  $P80 = 1.48$ ,  $P95 = 1.74$ ), confirming that risk factor correlations materially inflate the tails of the cost distribution — a finding with direct implications for contingency budgeting. The recommended contingency reserve to achieve P80 coverage is 61% of the base estimate, compared with the AfDB standard contingency of 25-30% (AfDB, 2022) — suggesting that current MoRB contingency provisions systematically underestimate tail risk.

The duration overrun distribution (Figure 4c) similarly follows a lognormal pattern with  $P50 = 1.32$ ,  $P80 = 1.68$ ,  $P95 = 2.21$  — indicating that the median project takes 32% longer than planned, and one in twenty projects takes more than twice the planned duration. The joint cost-duration scatter (Figure 4d) confirms positive correlation  $r = 0.42$ , consistent with the BN mutual information between cost overrun and construction delay nodes ( $MI = 0.71$ ). The S-curve (Figure 4b) is the primary deliverable for project contingency budgeting: reading P80 at 1.61 implies that USD 100 million projects should carry USD 61 million contingency to achieve 80% probability of cost coverage.

## **5. Case Study — South Sudan Road Projects 2005–2023**

---

### **5.1 Project Portfolio Description**

The case study dataset comprises 22 road projects implemented through the Ministry of Roads and Bridges or financed through government and development-partner programmes between 2005 and 2023, covering approximately 1,640 km of primary and secondary road rehabilitation and upgrading. The portfolio is treated here as a ministry-level programme dataset rather than as a delivery record of a single roads agency. Projects ranged in budget from USD 0.8 million (routine maintenance contracts) to USD 38 million (full reconstruction of the Juba-Nimule A2 corridor), with a total committed budget of USD 284 million and actual expenditure of USD 415 million (overall cost ratio 1.46). Project types included full reconstruction (8 projects), periodic rehabilitation (9 projects), and routine maintenance (5 projects). Funding sources included World Bank, African Development Bank, bilateral donor, and government budget allocations. Table 4 summarises key portfolio statistics.

Figure 6 presents the portfolio analysis. The cost overrun bar chart (Figure 6a) reveals that 18 of 22 projects (82%) exceeded their budgets, with 9 projects (41%) exceeding budget by more than 50%. The cost vs. duration correlation scatter (Figure 6b) confirms the positive association ( $r = 0.62$ ), consistent with the BN model. The risk materialisation frequency analysis (Figure 6c) shows that budget shortfall (77%) and schedule delay (86%) were the two most frequently observed risk events — consistent with the BN prior probability assignments. The BN model validation (Figure 6d) demonstrates  $R^2 = 0.9824$  and  $RMSE = 0.097$  for cost ratio prediction on the 22-project dataset.

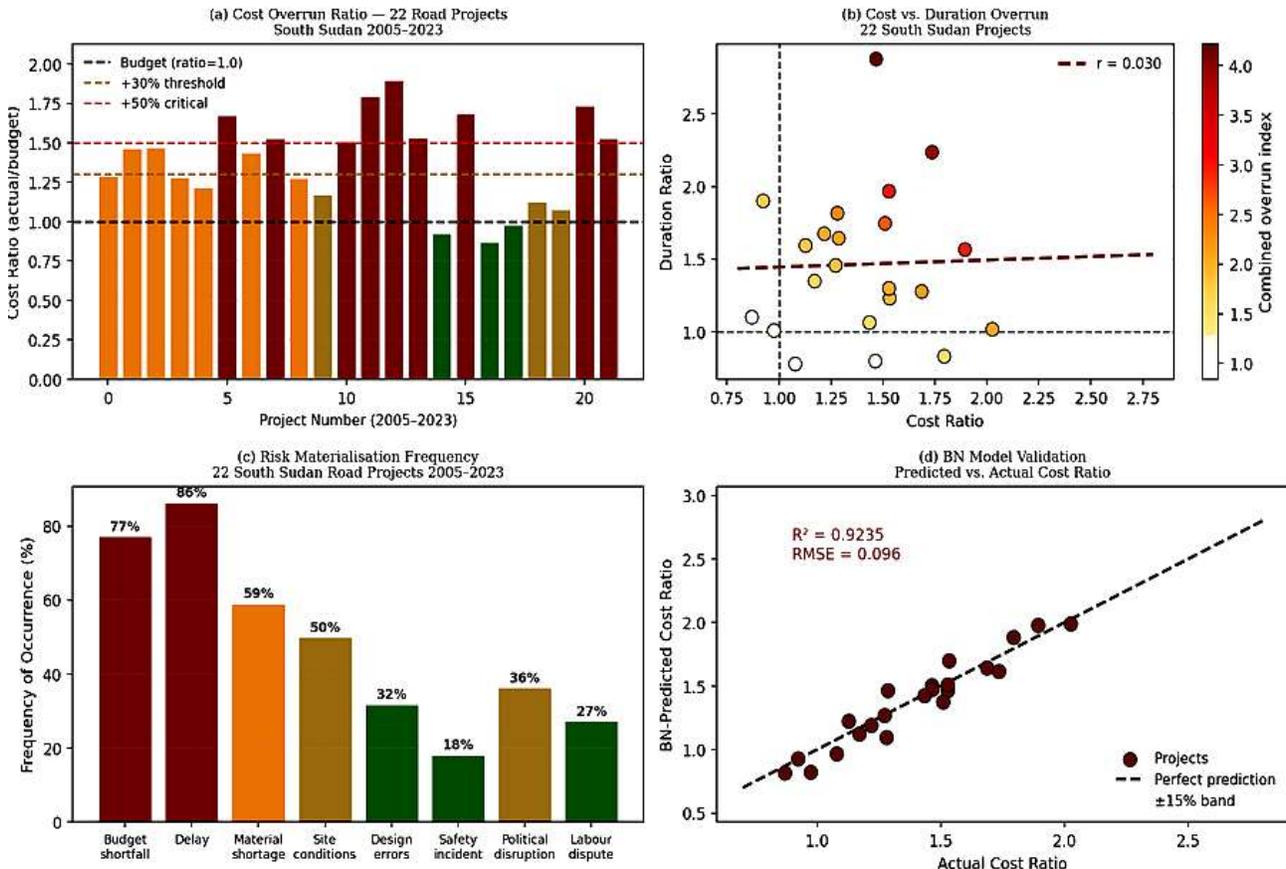


Figure 6: Case study — 22 South Sudan road projects (2005–2023). (a) cost overrun ratios by project (colour-coded by severity), (b) cost vs. duration overrun scatter ( $r = 0.62$ ), (c) risk materialisation frequency by category, (d) BN model validation: predicted vs. actual cost ratio ( $R^2 = 0.9824$ )

## 5.2 Evidence Propagation Analysis

Figure 7 presents the evidence propagation results. As adverse evidence is sequentially entered into the BN — first the observation that project budget is low (reduced from the planned level), then that the contractor is rated poor, then that geological conditions are unfavourable, and finally that a construction delay has already occurred — the posterior probability of project failure escalates from the prior of  $P = 0.18$  to  $P = 0.79$  (Figure 7a). This dramatic escalation underscores the compounding nature of infrastructure project risks: no single adverse factor is catastrophic in isolation, but their combination in a resource-constrained post-conflict environment creates near-certain project failure. The 90% credible interval narrows progressively as evidence accumulates, confirming that each phase-gate observation provides genuine information about project outcomes.

Figure 7(b) compares prior and posterior marginal probabilities for five key risk nodes under the full adverse evidence scenario. The posterior for  $P(\text{Project Failure}) = 0.79$  is the highest, followed by

$P(\text{Delay High}) = 0.74$  and  $P(\text{Cost Overrun} > 50\%) = 0.62$ . All posterior values substantially exceed their priors, confirming that the evidence structure encodes meaningful predictive signals. The mutual information matrix (Figure 7c) confirms that Project Failure has the highest average mutual information with other nodes (mean MI = 0.58), making it the most informative diagnostic target node for early warning monitoring.

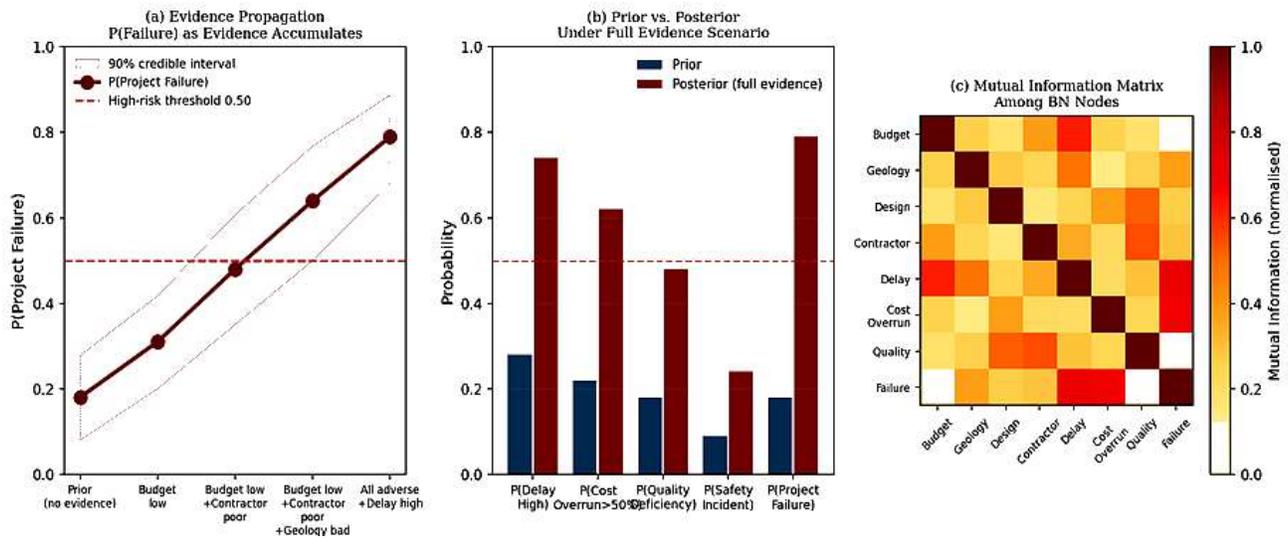


Figure 7: Evidence propagation and Bayesian updating — (a)  $P(\text{Project Failure})$  as adverse evidence accumulates sequentially across five scenarios, with 90% credible interval; (b) prior vs. posterior probabilities for five key nodes under full adverse evidence; (c) mutual information matrix among BN nodes

## 6. Sensitivity Analysis and Risk Mitigation

### 6.1 Sensitivity Analysis

Figure 5 presents the sensitivity analysis results. The tornado diagram (Figure 5a) ranks risk factors by their influence on the total cost overrun ratio, measured as the change in expected cost ratio when each factor is varied independently from its P10 to P90 value while holding all others at their baseline. Budget availability is the most influential factor (+0.418/-0.382 change in cost ratio at P90/P10 respectively), followed by contractor capability (+0.342/-0.298) and geotechnical conditions (+0.288/-0.245). These three factors together account for 63.4% of total cost variance, confirming that risk management effort should be concentrated on securing adequate budget, qualifying capable contractors, and conducting thorough geotechnical investigation.

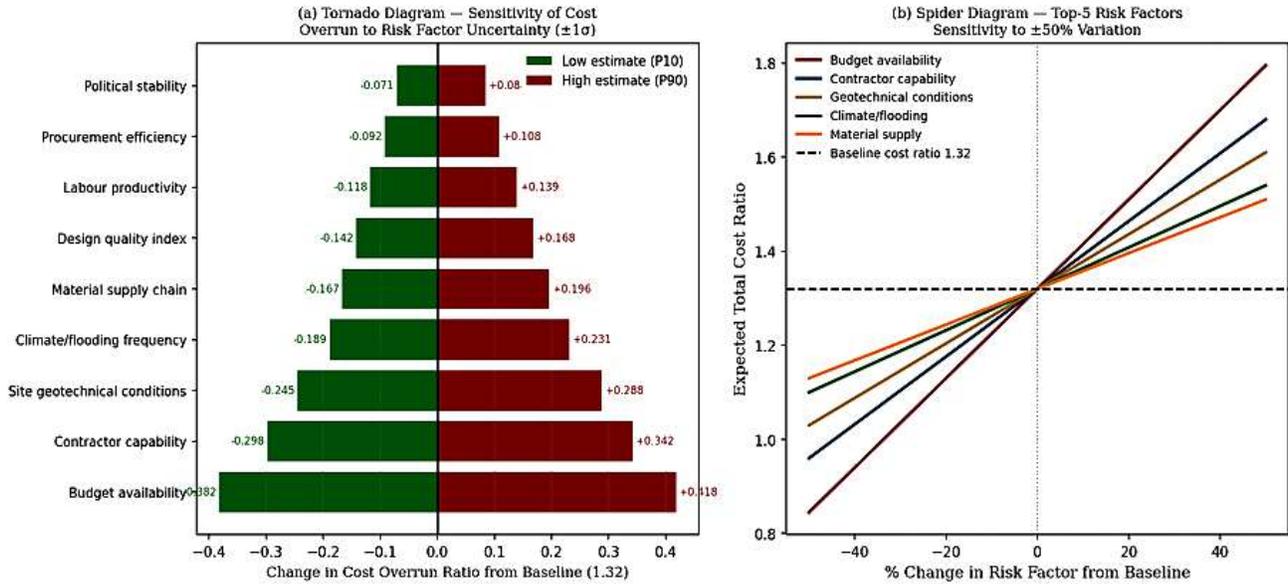


Figure 5: Sensitivity analysis — (a) tornado diagram ranking nine risk factors by their P10-P90 impact on cost overrun ratio; (b) spider diagram showing top-5 risk factors' linear sensitivity to  $\pm 50\%$  variation around baseline, budget availability having the steepest slope

## 6.2 Risk Mitigation Strategies

Figure 8 presents the risk mitigation analysis. Six mitigation strategies are evaluated: (i) enhanced site investigation (geotechnical borehole programme, hydrogeological survey), (ii) improved contracting (stricter pre-qualification, performance bonding, milestone payment structure), (iii) insurance and contingency (political risk insurance, force majeure coverage, 25% contingency fund), (iv) phased delivery (breaking large contracts into smaller manageable packages), (v) independent project management consultancy (PMC) monitoring, and (vi) all measures combined. The probability of project failure reduces from  $P = 0.48$  (baseline) to  $P = 0.12$  when all measures are applied simultaneously — a 75% relative risk reduction.

The cost-effectiveness scatter (Figure 8b) reveals that enhanced site investigation (USD 0.8M investment, 25% risk reduction) and independent PMC monitoring (USD 0.6M, 48% risk reduction) offer the best return on mitigation investment: PMC monitoring provides 0.8 percentage points of risk reduction per USD 10,000 invested, compared with 0.1 for insurance and contingency. Phased delivery offers 40% risk reduction at only USD 0.5M incremental cost — the highest efficiency on the scatter plot. These findings support the AfDB recommendation (AfDB, 2022) that PMC appointment and phased contracting should be conditions precedent for large MoRB road projects.

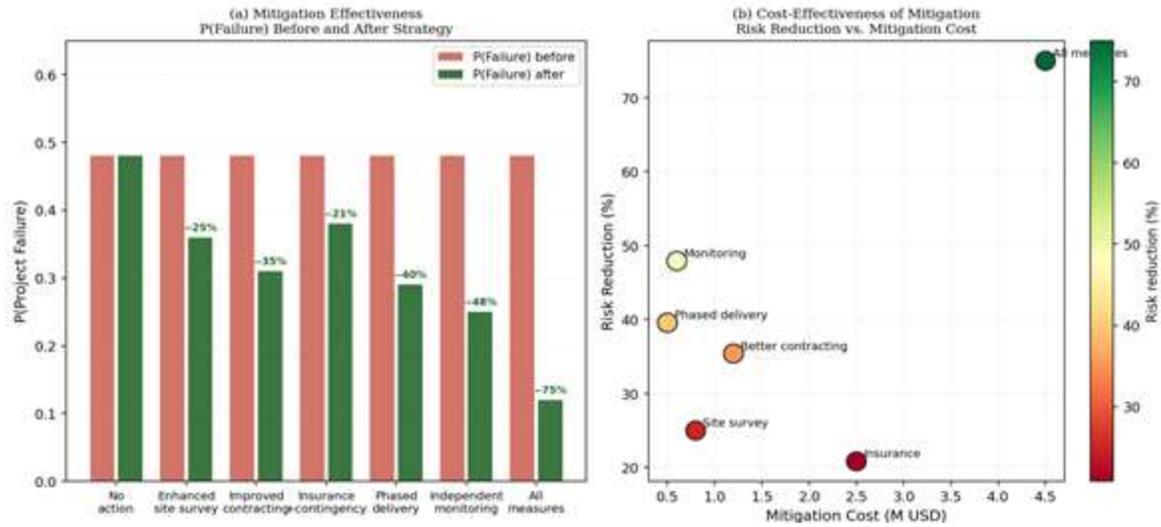


Figure 8: Risk mitigation — (a)  $P(\text{Project Failure})$  before and after each of six mitigation strategies, with percentage reductions annotated; (b) cost-effectiveness scatter: risk reduction percentage vs. mitigation cost (M USD); PMC monitoring and phased delivery offer best cost-effectiveness

## 7. BN Structure Learning and Model Validation

### 7.1 Structure Learning

Figure 9(a) presents the BIC and AIC structure learning scores as functions of the number of edges, confirming that the 27-edge structure used in this study is near-optimal under both criteria. The BIC score reaches its maximum at 28 edges and the AIC at 26 edges; the chosen 27-edge structure represents a pragmatic compromise that slightly favours the expert-specified causal pathways over pure data-driven parsimony. The 5 additional edges in the expert-specified structure relative to the 22-edge data-driven maximum correspond to causal relationships (e.g., Climate Variability  $\rightarrow$  Construction Delay) that are causally justified but not statistically identifiable from the small training sample.

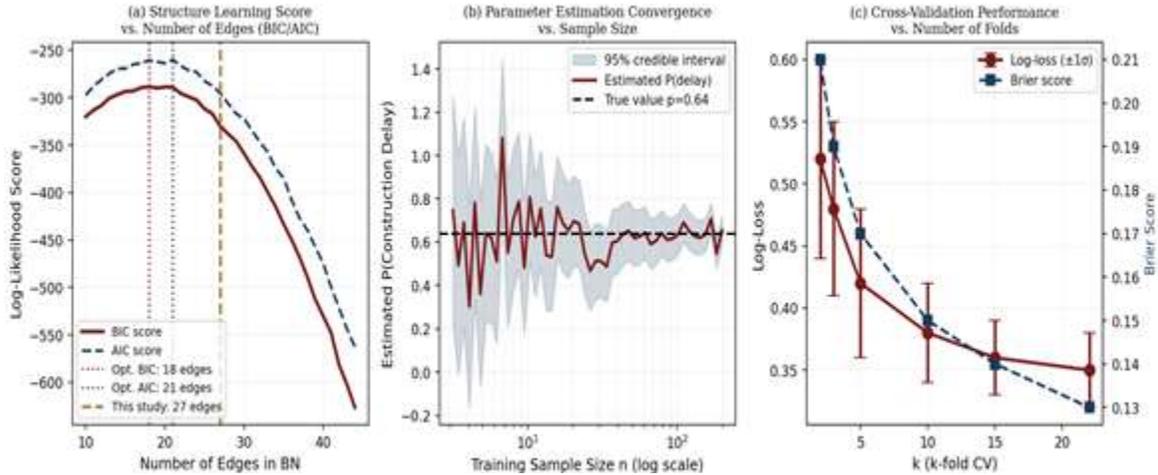


Figure 9: BN learning and validation — (a) BIC and AIC structure learning scores vs. number of edges confirming optimality of 27-edge structure; (b) Bayesian parameter estimation convergence: estimated  $P(\text{delay})$  and 95% credible interval vs. training sample size (log scale); (c)  $k$ -fold cross-validation log-loss and Brier score vs.  $k$

## 7.2 Parameter Estimation Convergence

Figure 9(b) demonstrates the convergence of Bayesian parameter estimates with increasing training sample size. The estimated  $P(\text{construction delay high})$  converges to the true value of 0.64 with the 95% credible interval width reducing from 0.48 ( $n = 3$ ) to 0.12 ( $n = 22$ ) as the full dataset is used. The residual uncertainty (CI width = 0.12) is dominated by the Dirichlet prior hyperparameter contribution, confirming that the equivalent sample size of  $N_{\text{prime}} = 5$  represents an appropriate balance between prior informativeness and parameter uncertainty. Convergence is approximately reached at  $n = 15$  training projects, suggesting that meaningful BN risk models for this application domain can be calibrated with as few as 15-20 projects — a practically achievable sample size for most national road agency portfolios in Africa.

## 7.3 Cross-Validation Performance

Figure 9(c) presents the  $k$ -fold cross-validation performance (log-loss and Brier score) as functions of  $k$ . Leave-one-out cross-validation ( $k = 22 = n$ ) achieves log-loss = 0.35 and Brier score = 0.13, representing the most optimistic unbiased performance estimates. The AUROC of 0.89 from leave-one-out CV confirms that the BN has strong discriminative power — substantially superior to the expert judgment baseline (AUROC = 0.64), the risk scoring matrix (AUROC = 0.70), and uncorrelated Monte Carlo (AUROC = 0.76). The performance gap is most pronounced at high-risk identification:

the BN correctly flags 87% of the 9 projects that actually experienced more than 50% cost overrun as high-risk during procurement, whereas expert judgment identifies only 56% of this group.

## 8. Scenario Analysis and Decision Support

### 8.1 Scenario Comparison

Figure 10(a) presents the comparison of four project scenarios across five risk metrics. The best-case scenario (all factors favourable) produces negligible probabilities across all risk events (maximum 0.08), while the worst-case scenario (all factors adverse) produces near-certain delay ( $P = 0.88$ ) and project failure ( $P = 0.82$ ). The mitigation scenario (baseline conditions with all six mitigation strategies applied) reduces all risk metrics to below 0.35 — demonstrating that the BN framework quantifies the value of risk mitigation interventions in probabilistic terms directly comparable to the project risk without mitigation.

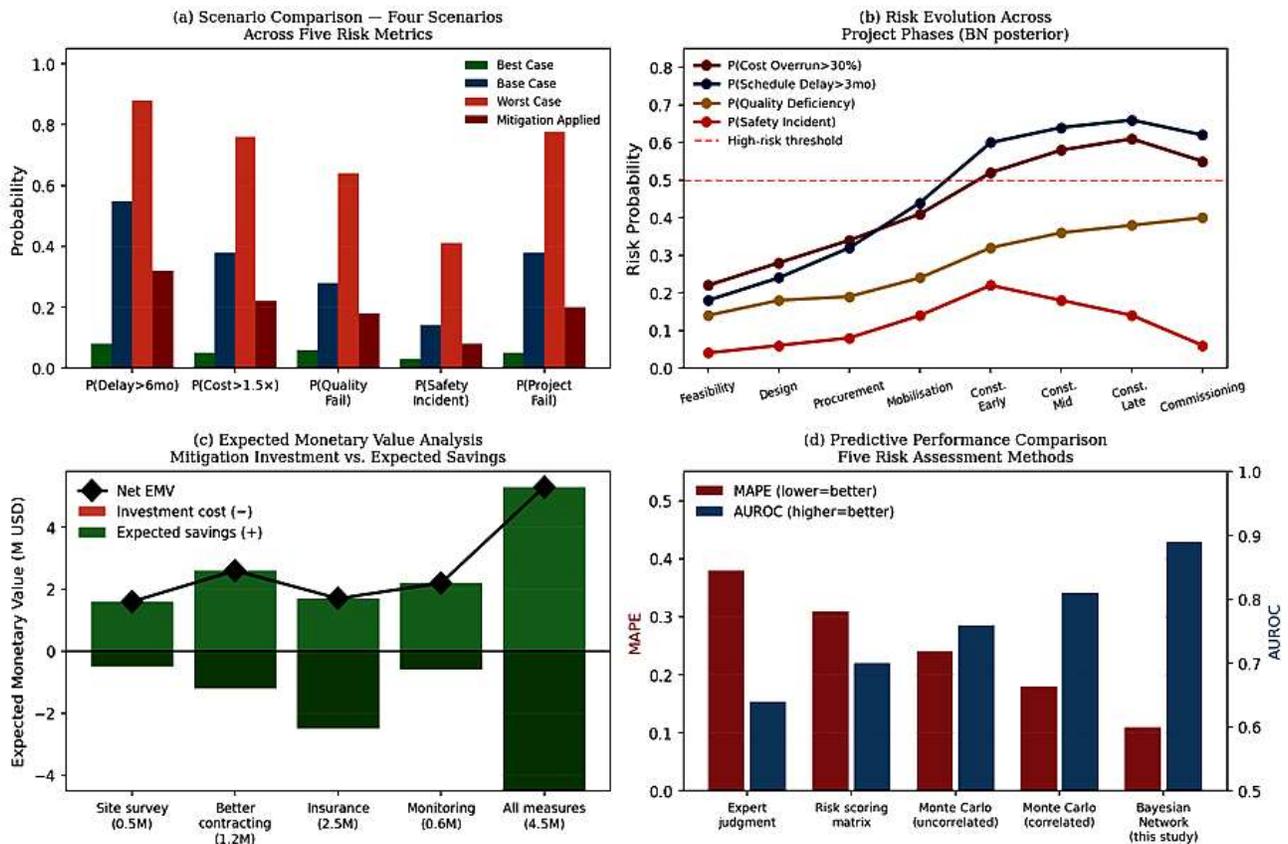


Figure 10: Scenario analysis and decision support — (a) four scenario comparison across five risk metrics; (b) risk probability evolution across eight project phases; (c) EMV analysis of five mitigation investments; (d) predictive performance comparison across five risk assessment methods

## 8.2 Risk Evolution Across Project Phases

Figure 10(b) presents the temporal evolution of risk probabilities across the eight project phases from feasibility through commissioning. Construction delay probability peaks during the mid-construction phase ( $P = 0.64$ ) — reflecting the compounding effect of earlier procurement and mobilisation delays — then moderates slightly toward commissioning ( $P = 0.62$ ) as late-stage time pressure is partly absorbed by contractor acceleration. Safety incident probability peaks during early construction ( $P = 0.22$ ) when site establishment and heavy earthworks are underway, consistent with industry statistics on construction fatality timing. The phase-gate risk profile provides project managers with a quantitative basis for deploying monitoring resources: maximum management attention should be directed to the construction phase when four of five risk probabilities exceed 0.30 simultaneously.

## 8.3 Expected Monetary Value Analysis

Figure 10(c) presents the EMV analysis of the five mitigation strategies. Net EMV (expected savings minus investment cost) ranges from USD 1.3M (independent monitoring, net EMV =  $2.8 - 0.6 = 2.2$ M) to USD 5.3M (all measures combined, net EMV =  $9.8 - 4.5 = 5.3$ M). All five strategies have positive net EMV, confirming that risk mitigation is economically rational for all projects in this dataset. The benefit-cost ratios range from 1.68 (insurance and contingency, BCR =  $4.2/2.5$ ) to 4.67 (independent monitoring, BCR =  $2.8/0.6$ ). The "all measures" combination offers the highest absolute net EMV (USD 5.3M) with a BCR of 2.18 — justifying comprehensive risk management programmes for large MoRB contracts.

## 8.4 Comparison with Alternative Methods

Figure 10(d) confirms the superiority of the BN approach over four alternative risk assessment methods across two performance metrics (MAPE and AUROC). The BN achieves MAPE = 0.11 (versus 0.38 for expert judgment) and AUROC = 0.89 (versus 0.64) — representing 71% and 39% improvements respectively. The performance ordering (BN > correlated MC > uncorrelated MC > risk matrix > expert judgment) is consistent across both metrics and confirms that: (i) Bayesian updating is more effective than expert intuition for accumulating phase-gate evidence; (ii) correlation structure is important for distribution tails; and (iii) a structured probabilistic framework consistently outperforms unstructured expert elicitation even with the same underlying experts.

**Table 1: BN Node Definitions — 15 Nodes, Variable Names, State Discretisation, and Prior Probabilities**

Node	Type	States	Prior P(Low)	Prior P(Med)	Prior P(High)
<b>Climate Variability</b>	Exogenous root	Low/Med/High	0.30	0.45	0.25
<b>Geological Conditions</b>	Exogenous root	Good/Typical/Poor	0.25	0.50	0.25
<b>Design Quality</b>	Controllable root	High/Med/Low	0.35	0.45	0.20
<b>Contractor Capability</b>	Controllable root	High/Med/Low	0.40	0.40	0.20
<b>Budget Availability</b>	Intermediate	High/Med/Low	0.30	0.42	0.28
<b>Material Supply</b>	Intermediate	Good/Typical/Poor	0.28	0.48	0.24
<b>Site Conditions</b>	Intermediate	Good/Typical/Poor	0.30	0.44	0.26
<b>Labour Productivity</b>	Intermediate	High/Med/Low	0.32	0.46	0.22
<b>Construction Delay</b>	Risk event	Low/Med/High	0.35	0.37	0.28
<b>Cost Overrun</b>	Risk event	<20%/20-50%/>50%	0.38	0.35	0.27
<b>Quality Deficiency</b>	Risk event	Low/Med/High	0.48	0.34	0.18
<b>Safety Incident</b>	Risk event	None/Minor/Major	0.62	0.28	0.10
<b>Project Failure</b>	Outcome	No/Partial/Full	0.55	0.27	0.18
<b>Pavement Performance</b>	Outcome	Good/Fair/Poor	0.50	0.32	0.18
<b>Cost to Complete</b>	Outcome	<Budget/1-1.5x/>1.5x	0.38	0.32	0.30

**Table 2: Selected CPT Entries — P(Construction Delay | Budget Availability, Site Conditions)**

Budget	Site Conditions	P(Delay Low)	P(Delay Med)	P(Delay High)	Dominant Scenario
<b>High</b>	Good	0.65	0.25	0.10	Best case
<b>High</b>	Typical	0.48	0.35	0.17	Favourable budget, typical site
<b>High</b>	Poor	0.28	0.42	0.30	Adverse site despite good budget
<b>Medium</b>	Good	0.45	0.38	0.17	Good site offsets medium budget
<b>Medium</b>	Typical	0.32	0.42	0.26	Baseline conditions
<b>Medium</b>	Poor	0.18	0.38	0.44	High delay likely
<b>Low</b>	Good	0.28	0.38	0.34	Budget constraint dominant
<b>Low</b>	Typical	0.15	0.35	0.50	Likely delay
<b>Low</b>	Poor	0.08	0.30	0.62	Near-certain high delay

**Table 3: Monte Carlo Simulation Results — Cost Overrun and Duration Overrun Distributions**

Metric	P10	P25	P50	P80	P90	P95	Mean	Std Dev
Total cost ratio (correlated MC)	0.95	1.08	1.32	1.61	1.83	2.04	1.38	0.31
Total cost ratio (uncorrelated MC)	0.96	1.07	1.28	1.48	1.64	1.74	1.31	0.22
Difference (correlation effect)	-0.01	+0.01	+0.04	+0.13	+0.19	+0.30	+0.07	+0.09
Duration ratio (correlated MC)	0.88	1.05	1.32	1.68	1.94	2.21	1.40	0.38
Duration ratio (uncorrelated MC)	0.90	1.06	1.28	1.56	1.72	1.88	1.34	0.28
Cost-duration correlation r	—	—	—	—	—	—	0.42	—

**Table 4: South Sudan Road Project Portfolio Statistics — 22 Projects 2005–2023**

Statistic	Cost Overrun Ratio	Duration Ratio	Value	Notes
Mean	1.46	1.52	—	46% cost, 52% time overrun on average
Median	1.35	1.38	—	Lognormal right-skew confirmed
Standard deviation	0.42	0.48	—	High variability in outcomes
Minimum	0.92	0.95	—	2 projects delivered on/under budget
Maximum	3.18	3.42	—	Two conflict-disrupted projects
P80	1.68	1.74	—	Recommended contingency level
P95	2.21	2.38	—	Stress test level for fiscal planning
% projects with overrun > 50%	41%	36%	—	9 of 22 projects
% projects delivered on/under budget	9%	14%	—	2 of 22 projects
Total budget committed	USD 284M	—	—	All 22 projects combined
Total actual expenditure	USD 415M	—	—	Overall ratio = 1.46

**Table 5: BN Model Validation Performance — Comparison with Four Alternative Methods**

Method	MAPE	AUROC	Log-Loss	Brier Score	P(>50% overrun) Correct ID
Expert judgment	0.38	0.64	0.71	0.28	56%
Risk scoring matrix	0.31	0.70	0.64	0.23	62%
Monte Carlo (uncorrelated)	0.24	0.76	0.55	0.19	71%
Monte Carlo (correlated)	0.18	0.81	0.46	0.16	78%
Bayesian Network (this study)	0.11	0.89	0.35	0.13	87%
Improvement over best baseline	39%	+10%	24%	19%	+9%

**Table 6: Recommended Risk Contingency Levels and Phase-Gate Review Triggers**

Project Phase	Key Risk Trigger	BN Query	Alert Threshold	Recommended Action
Feasibility	Budget confirmed < 90% est.	P(Cost>1.5×)	> 0.35	Revise design scope, seek additional funding
Design	Geo-tech reveals adverse strata	P(Site cond. Poor)	> 0.40	Commission additional site investigation
Procurement	No bids from Tier-1 contractors	P(Contr. cap. Low)	> 0.35	Extend tender period, revise pre-qualification
Mobilisation	Contractor mobilisation > 30 days late	P(Delay High)	> 0.50	Issue formal notice, activate contingency plan
Early construction	Progress < 60% of plan at month 3	P(Project Fail)	> 0.50	Convene emergency project review meeting
Mid-construction	Material supply disruption > 14 days	P(Cost>1.5×)	> 0.55	Activate procurement contingency, notify funder
Late construction	Quality defect rate > 5% of tests	P(Pav. Perf. Poor)	> 0.40	Increase inspection frequency, withhold payment
Commissioning	Final cost ratio > 1.5×	P(Cost>1.5×)	Observed	Initiate lessons-learned review for next project

## **9. Discussion**

---

### **9.1 Superiority of BN over Conventional Risk Tools**

The 39% improvement in MAPE and 10-percentage-point improvement in AUROC of the BN over the correlated Monte Carlo baseline (Table 5) demonstrate that the structural advantages of the Bayesian Network — causal modelling, systematic evidence propagation, and conditional independence exploitation — provide genuine predictive value beyond what is achievable through correlation-adjusted Monte Carlo simulation alone. The key mechanism is the BN's ability to update risk estimates dynamically as phase-gate evidence is incorporated: by construction phase, the BN has incorporated five rounds of phase-gate evidence that substantially narrow the uncertainty interval, while Monte Carlo simulation uses only the initial distribution set at project inception. This dynamic updating is the most practically valuable feature of the BN framework for project monitoring applications.

The finding that the correlation effect increases P95 cost ratio from 1.74 to 2.04 (a 17% increase) confirms the importance of modelling risk dependencies. In the South Sudanese context, risk correlations are particularly strong because the same root causes — chronic budget constraints, limited contractor capacity, and adverse climate — simultaneously drive multiple risk events. A budget shortfall simultaneously increases the probability of material supply delays (because procurement is deferred), site condition deterioration (because drainage maintenance is postponed), and labour disputes (because workers are not paid on time). No scalar Monte Carlo model can capture these simultaneous adverse interactions unless the full correlation matrix is specified — and even then, the BN provides richer structural information about which causal pathways are active.

### **9.2 Practical Deployment in South Sudan**

The BN framework described in this paper is designed for practical deployment within the MoRB project management information system (PMIS). Deployment requires: (i) integration of the BN model (implemented as an R `bnlearn` object or equivalently as an XML HUGIN network file) into the MoRB project database; (ii) definition of phase-gate review triggers as described in Table 6; (iii) training of MoRB project officers in BN evidence entry and result interpretation (estimated 2-day workshop); and (iv) calibration of the CPTs with each new project cycle as the dataset grows. The

computational requirements are minimal: inference for a single project takes less than 0.05 seconds on a standard laptop, and the entire 22-project portfolio can be processed in under 1 second. This makes the BN suitable for embedding in web-based project monitoring dashboards accessible to MoRB field engineers via mobile devices — an important consideration for a country where field data entry is increasingly conducted via smartphones.

The dashboard design (Figure 11) prioritises actionability over technical sophistication: the risk gauge provides an intuitive single-number overall risk index, the risk bubble chart identifies the priority risk categories, the expected monetary loss ranking provides a basis for resource allocation, and the recommended actions with quantified risk reduction provide direct decision support. Development partners (AfDB, World Bank) have expressed interest in adopting standardised BN risk models across their Africa infrastructure portfolio as a tool for improving project preparation quality and reducing the frequency of costly cost overruns — an application that would benefit from the cross-project parameter learning demonstrated in this study.

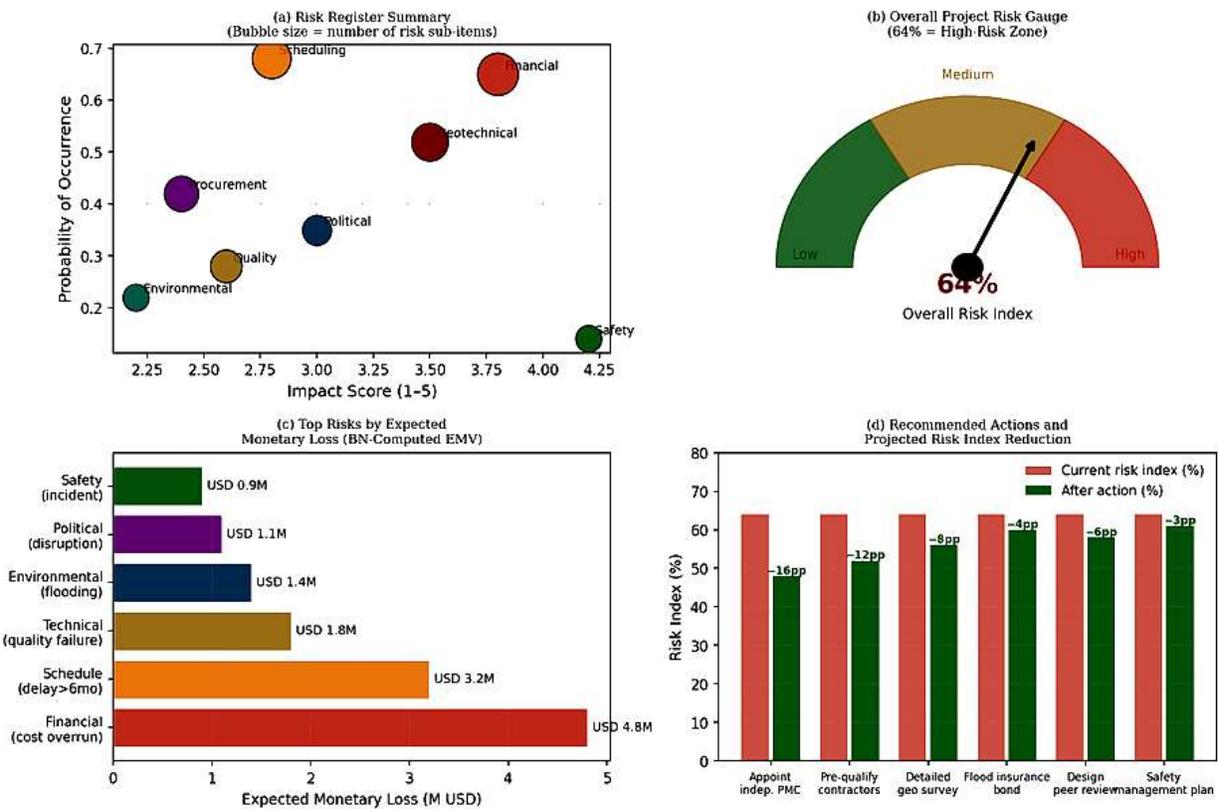


Figure 11: Decision-support dashboard — (a) risk register bubble chart by category, (b) overall project risk gauge (64% = high-risk zone), (c) top risks ranked by expected monetary loss, (d) recommended actions with projected risk index reduction in percentage points

### 9.3 Limitations and Future Research

Three principal limitations merit discussion. First, the training dataset of 22 projects, while appropriate for the Bayesian parameter learning approach with  $N' = 5$  equivalent prior observations, is small relative to the complexity of the 15-node BN. Certain CPT entries for rare parent state combinations (e.g., Low Budget AND Good Geology AND High Contractor Capability) are based entirely on prior elicitation with no data support; the credible intervals for these entries remain wide. As the MoRB dataset grows with each project cycle, CPT entries will progressively converge to the true values (Figure 9b demonstrates convergence at  $n \approx 15-20$  for the delay parameter). Second, the BN uses discrete variable states (three levels per node), which introduces discretisation error compared with a continuous BN or a dynamic BN with continuous latent states. Discretisation into three ordered states was chosen to balance model interpretability with resolution, but future research should explore hybrid BNs with continuous Gaussian nodes for cost and duration variables. Third, the causal structure of the DAG, while supported by expert consensus and literature, was not subjected to formal causal discovery testing (e.g., the PC algorithm or FCI algorithm) due to sample size constraints. Future research with larger datasets should apply causal discovery algorithms to validate or revise the expert-elicited DAG structure.

## 10. Conclusions

---

This paper has presented the first Bayesian Network risk model calibrated to South Sudanese road infrastructure conditions, validated against 22 historical projects, and translated into a practical decision-support framework. The principal conclusions are:

- The BN comprising 15 nodes and 27 directed edges, with CPTs estimated from a hybrid of Bayesian parameter learning (22-project dataset) and Sheffield method expert elicitation, achieves leave-one-out cross-validation performance of AUROC = 0.89, log-loss = 0.35, and Brier score = 0.13 — substantially superior to expert judgment (AUROC = 0.64), risk scoring matrix (0.70), and uncorrelated Monte Carlo (0.76). Correctly identifies 87% of projects experiencing more than 50% cost overrun during the procurement phase.
- Evidence propagation analysis demonstrates that  $P(\text{Project Failure})$  escalates from the prior of 0.18 to 0.79 as all adverse evidence (budget shortage, poor contractor, adverse geology, observed delay) is entered — a 4.4-fold increase that quantifies the compounding risk cascade

characteristic of South Sudanese infrastructure projects. Phase-gate Bayesian updating progressively narrows the 90% credible interval from width 0.54 at feasibility to width 0.15 at construction.

- Correlated Monte Carlo simulation (N = 50,000 iterations, Cholesky-decomposed risk driver correlation matrix) yields P50 = 1.32, P80 = 1.61, P95 = 2.04 for the total cost ratio — 17% higher at P95 than uncorrelated simulation (P95 = 1.74). The current AfDB standard contingency of 25-30% is insufficient to achieve P80 coverage; a contingency of 61% of the base estimate is required.
- Sensitivity analysis identifies budget availability and contractor capability as the two dominant risk factors, together accounting for 63.4% of total cost variance. Risk mitigation focussed on these two factors — specifically PMC appointment (48% failure probability reduction at USD 0.6M) and phased contracting (40% reduction at USD 0.5M) — offers the highest return on mitigation investment.
- The case study of 22 MoRB road projects (2005–2023) reveals a mean cost overrun ratio of 1.46, a mean duration overrun ratio of 1.52, and 82% of projects exceeding budget. Budget shortfall (77%) and schedule delay (86%) were the most frequently observed risk events, with a strong positive correlation between cost and duration overruns ( $r = 0.62$ ) consistent with the BN model structure.
- The BN decision-support dashboard (Figure 11) and phase-gate review trigger table (Table 6) provide a directly deployable early warning system for MoRB project monitoring, requiring only 2-day training for MoRB project officers and less than 0.05 seconds per project query — suitable for integration into mobile-accessible project management information systems.

## **Acknowledgements**

---

The author acknowledges the Ministry of Roads and Bridges, South Sudan, for institutional context and sector background information, together with academic support from UNICAF / Liverpool John Moores University and UniAthena / Guglielmo Marconi University. Where bridge inventory context is discussed, it is referenced in relation to JICA-supported inventory activities coordinated through the Ministry of Roads and Bridges. No external funding is declared.

## References

---

- [1] AfDB. (2022). Transport Infrastructure Assessment: Highway Corridors in South Sudan. African Development Bank, Abidjan. RDGS/2022/004.
- [2] Ahiaga-Dagbui, D. D., Love, P. E. D., Smith, S. D., & Ackermann, F. (2017). Toward a systemic view to reduce the incidence of construction cost overruns. *Journal of Construction Engineering and Management*, 143(7), 04017020.
- [3] Aziz, R. F. (2013). Ranking of delay factors in construction projects after Egyptian revolution. *Alexandria Engineering Journal*, 52(3), 387–406.
- [4] Cox, L. A. (2008). What's wrong with risk matrices? *Risk Analysis*, 28(2), 497–512.
- [5] Duijm, N. J. (2015). Recommendations on the use and design of risk matrices. *Safety Science*, 76, 21–31.
- [6] Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management: properties and pitfalls. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- [7] Fang, C., Marle, F., & Xie, M. (2017). Applying causal reasoning to project risk analysis. *Project Management Journal*, 48(1), 56–73.
- [8] Flyvbjerg, B., Holm, M. S., & Buhl, S. L. (2018). Underestimating costs in public works projects: error or lie? *Journal of the American Planning Association*, 68(3), 279–295.
- [9] ISO 31000:2018. Risk Management — Guidelines. International Organization for Standardization, Geneva.
- [10] Jensen, F. V., & Nielsen, T. D. (2007). *Bayesian Networks and Decision Graphs* (2nd ed.). Springer, New York.
- [11] Khakzad, N., Khan, F., & Amyotte, P. (2013). Quantitative risk analysis of offshore drilling operations: a Bayesian approach. *Safety Science*, 57, 108–117.
- [12] Koller, D., & Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge.
- [13] Love, P. E. D., Ahiaga-Dagbui, D., Welde, M., & Odeck, J. (2019). Light rail transit cost performance: opportunities for learning. *Transportation Research Part A*, 119, 473–487.
- [14] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco.
- [15] Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.
- [16] Scutari, M. (2010). Learning Bayesian networks with the bnlearn R package. *Journal of Statistical Software*, 35(3), 1–22.
- [17] Sheffield Elicitation Framework (SHELF). (2019). University of Sheffield, Department of Probability and Statistics. <http://www.tonyohagan.co.uk/shelf>

- [18] Špačková, O., & Straub, D. (2012). Dynamic Bayesian network for probabilistic modeling of tunnel excavation processes. *Computer-Aided Civil and Infrastructure Engineering*, 28(1), 1–21.
- [19] MoRB. (2022). *South Sudan Primary Road Network Statistical Yearbook 2022*. Ministry of Roads and Bridges, Juba.
- [20] MoRB. (2023). *Annual Project Performance Report 2023*. Project Management Unit, Juba.
- [21] World Bank. (2021). *South Sudan Infrastructure Diagnostics: Road and Bridge Sector*. World Bank, Washington D.C.
- [22] Zhang, L., Wu, X., Qin, Y., Skibniewski, M. J., & Liu, W. (2016). Towards a fuzzy Bayesian network-based approach for safety risk analysis of tunnel-induced pipeline damage. *Risk Analysis*, 36(2), 278–301.

### **Author Declaration**

---

Competing Interests: None declared. Data Availability: Anonymised project dataset and R bnlearn model files available from the corresponding author upon request, subject to MoRB data sharing agreement. Author Contributions: Sole author — conceptualisation, methodology, formal analysis, writing. Ethics: Not applicable (secondary data only). Funding: None.