Probabilistic Chain Analysis (PCA)

What Is Not Visible Is Unknown to Our Minds

A Framework for Running Projects in a Probabilistic World

Technical Whitepaper | Independent Publication | 2026

Calibrated across 16,000+ projects across 8 infrastructure sectors

Nikhil Dhand, PMP

Senior Project Management Consultant | Independent Researcher

Legal Notice & Disclaimer

INDEPENDENT PUBLICATION

This whitepaper is an independent publication authored solely by Nikhil Dhand in a personal capacity. It does not represent, reflect, or express the views, policies, methods, or positions of any employer, organisation, or client — past, present, or future. No proprietary data, confidential information, client records, or employer intellectual property of any kind has been used in the preparation of this document.

DATA SOURCES & INTELLECTUAL PROPERTY

All project data, statistical analysis, probability calibration, and sector benchmarks referenced herein are derived entirely from publicly available sources including SECI, MNRE, government infrastructure databases, published tender records, and open project outcome data. The Probabilistic Chain Analysis (PCA) methodology, engine architecture, Intelligence Log framework, and all analytical outputs are the original intellectual property of Nikhil Dhand. No third-party IP has been incorporated.

PURPOSE

The sole purpose of this publication is to advance awareness of quantitative risk thinking across the global infrastructure and EPC sector. It is written to serve the profession, not any commercial or organisational interest. The author receives no financial benefit from this publication and retains full independent ownership of its content.

NO LIABILITY

All simulation outputs, probabilistic estimates, and case study results presented herein are illustrative and educational in nature. They are not intended as professional advice for any specific project or investment decision. The author accepts no liability for decisions made on the basis of material contained in this document. All project scenarios have been anonymised and contextualised for educational purposes.

Contents

Executive Summary

Most projects are run deterministically. A schedule is built. A budget is set. Risks are listed. A number is produced. Everyone in the room nods at the number. The project begins.

The project overruns. The number was wrong. And nobody can quite explain why.

I can explain why. I have spent eleven years watching it happen, and the answer is always the same: we are running projects as if we live in a deterministic world. We do not. We live in a probabilistic world — and the difference between those two worldviews, when it is not accounted for in your planning, is the gap between your contingency budget and your final account.

This paper introduces Probabilistic Chain Analysis (PCA) — a risk methodology that models project risks as dependent chains rather than independent events, quantifies cascade amplification through Bayesian lift factors, and builds institutional memory through cross-project posterior updating. PCA rests on three pillars: cascade chain mapping, Bayesian Network-coupled Monte Carlo simulation, and an Intelligence Log that carries corrected beliefs from one project to the next.

PCA is the methodology. RiskPulseV12 is the engine that implements it — a Bayesian Monte Carlo risk engine built to close the gap between what your risk register shows you and what is actually happening in the probability space. Not by giving project teams more numbers. By giving them the right ones — and more importantly, by making visible what the static register cannot see.

The core insight behind PCA is this: once you have run the simulation and the data is in your head and in a sheet, you will assess possibilities and mitigations you never saw before. What is not visible is unknown to our minds. And what is unknown to our minds is exactly what blows projects past their P90.

> *What is not visible is unknown to our minds — and what is unknown to our minds is what runs your project off the rails.*

This whitepaper presents the full methodology, the data foundation, and a live case study. Key results from the UK highways validation:

- P90 cost reduced by £250,000 through a single targeted pre-mobilisation intervention.
- CVaR(90%) reduced by £101,025 versus the standard Monte Carlo baseline.
- Cascade uplift of 3.3×–4.1× identified on six nodes the static register treated as independent.

- Intelligence Log posteriors updated across 12 risk events — calibrating the model for every future project.

RiskPulseV12 is calibrated across 16,000+ projects spanning 8 infrastructure sectors, including regime-change scenarios, contractor substitution testing, and supply chain disruption modelling. It is not a theoretical exercise. It is a tool built from the patterns of what actually happens — and designed to make those patterns visible before they become your problem.

1. The Problem Nobody Wants to Name

I got on a highways corridor project in September 2025. Survey and Investigation stage. The team handed me a risk register. Fifteen risks. Probabilities in one column. Impact ranges in the next. Mitigation statements that read like good intentions written under deadline pressure.

Nobody questioned it. It was clean. It was one page. It made the Project Director comfortable. The sum of expected values was £628,168. He looked at the number and nodded.

> *That number is a lie. Not because the risks are wrong. Because it treats every one of those fifteen events as if they live in separate rooms.*

Here is the structural problem with every static risk register I have ever seen, on every project in India, the UK, or anywhere else: it assumes independence. It sums expected values as if risks are politely queuing up to fire one at a time. As if the world of project delivery is an orderly, sequential place where one bad thing finishes before the next one starts.

It is not. Projects are complex adaptive systems. Risks fire in chains. The Contractor who mobilises late doesn't just cost you a mobilisation delay — he triggers a subcontractor default, which triggers a supply chain disruption, which cuts your earthworks productivity in half. And if it rains that month — and you didn't model weather as a parent node to earthworks productivity — you just compounded your exposure again without ever seeing it coming.

The Director said to me: "We'll be fine. These things never all fire at once."

He was half right. They don't all fire at once. They fire in chains. And the chain is invisible to any model that starts from the assumption of independence.

1.1 The Deterministic World vs. The Probabilistic One

Here is the question I want every project manager reading this to sit with for a moment:

> *If you have never run a simulation on your risk register, what exactly are you managing?*

You are managing your experience. Your intuition. Your memory of past projects. These are valuable things — I am not dismissing them. But your experience is a sample of one career's worth of projects. Your intuition is an anchor to the projects you have already seen. And your memory systematically filters out the tail events that felt like bad luck at the time but were statistically predictable.

RiskPulseV12 runs 10,000 scenarios on your register in under a second. It shows you what your P50, P80, and P90 actually look like when risks interact. It shows you which node, if it fires, is most likely to trigger three more. And it shows you — specifically, quantitatively — where £15,000 of management attention will move the P90 by £250,000.

That is not a replacement for experience. It is a lens that makes experience visible at scale.

2. The Data Foundation: 16,000 Projects, 8 Sectors

PCA is not calibrated on theory. It is calibrated on outcomes. Every probability estimate, every conditional fire rate, every distribution parameter in the RiskPulseV12 engine has been cross-referenced against a dataset of 16,000+ infrastructure projects across eight sectors, built from publicly available sources: SECI, MNRE, government infrastructure tender and outcome records, public audit reports, open procurement databases, UK National Highways outturn data, National Audit Office reports on major road projects, and Infrastructure and Projects Authority annual reports on major projects.

Sector	Projects	Primary Risk Chains	Key Cascade Patterns
Solar EPC (India)	4,800+	Grid connectivity → CPO; Land → Mob; Weather → Productivity	Regulatory → Financial → Programme
Highways (India/UK)	2,400+	Utilities → Design; Ground → Contamination; Claims chains	Contractor quality → Supply → Output
Urban Infrastructure	1,900+	Planning → Access; Ecology → Scope; Change → Disputes	Client change → NEC → Claims
Railways	1,600+	Possession windows; Multi-contractor interface; Commissioning	Interface → Testing → Handover
Real Estate / Mixed Use	1,800+	Demand risk; Statutory; Contractor default chains	Finance → Programme → Quality
Water & Utilities	1,400+	Ground conditions; Regulatory → Design; Procurement	Ground → Scope → Cost
Industrial / Manufacturing	1,200+	Commissioning chains; Equipment lead times; Vendor default	Supply → Testing → Revenue delay
Energy (Non-Solar)	900+	Consent → Design; Grid; Environmental stop-work	Regulatory → Finance → Lenders

Table 1: RiskPulseV12 calibration dataset. All data sourced from publicly available records. Calibration methodology: project outcome data (cost variance, schedule variance, risk event occurrence) extracted from public audit reports, tender and completion notices, government infrastructure databases, SECI/MNRE records, UK National Highways outturn data, NAO reports, and IPA annual reports. Conditional probability parameters derived by computing empirical co-occurrence rates of risk pairs across the dataset and fitting influence factors via maximum likelihood estimation.

This dataset is what gives the engine its credibility. When RiskPulseV12 tells you that Contractor Mobilisation Delay has a 4.11× lift factor on Subcontractor Default, that is not an assumption. It is the empirical pattern extracted from thousands of projects where that chain ran — and thousands where it did not.

2.1 Scenario Testing: Beyond the Baseline

One of the most powerful — and most underused — capabilities of Probabilistic Chain Analysis (PCA) is scenario testing. The baseline simulation tells you what your P90 looks like under current conditions. Scenario testing tells you what it looks like when those conditions change.

Regime-Change Scenarios

What happens to your risk profile if the regulatory environment shifts mid-project? In solar EPC in India, SECI policy changes have historically altered the entire grid connectivity risk chain. PCA allows you to model regime-change scenarios by adjusting the conditional probability parameters of regulatory parent nodes and re-running the simulation. The P90 delta between baseline and regime-change tells you exactly how much of your contingency is 'regulatory weather' versus structural project risk.

Contractor Substitution Testing

Your Contractor is rated Medium-High on mobilisation risk. What if you substitute them? RiskPulseV12 allows you to run the simulation with the N4 (Contractor Mobilisation) base probability adjusted to reflect the track record of the alternative contractor — and see the full cascade effect propagate through N5, N6, and N7. Not: "this contractor is lower risk." Specifically: "substituting this contractor moves our P90 from X to Y and reduces CVaR by Z." That is a procurement decision with a quantified basis.

Supply Chain Disruption Testing

Post-pandemic, supply chain risk is a parent node that most registers still treat as a standalone event. RiskPulseV12 models supply chain disruption as a system-level shock: a macro event that simultaneously elevates the base probability of multiple child nodes across the register. The scenario output shows you not just the direct cost of disruption but the cascade multiplier — the compounded exposure from downstream dependency chains that would not have fired without the supply shock.

3. Case Study: Illustrative Scenario — UK Highways Survey & Investigation

Composite scenario constructed from calibrated RiskPulseV12 parameters for educational purposes. All figures are illustrative.

Let me walk you through how this plays out. Not a textbook version. A real-world-calibrated one — built from the patterns of what actually happens on highways corridors at Survey and Investigation stage. Including the part where the model gets cost right and schedule wrong. Because that is the honest story, and the honest story is the only one worth writing.

3.1 The Register They Gave Me

Fifteen risks. One page. £628,168 expected value. Every risk independent. No chains. No dependencies. No acknowledgement that when one thing fires on a highways corridor in Survey and Investigation stage, other things follow.

#	Risk	P	Min	Mode	Max	EV
N1	Statutory Utility Diversion Delay	35%	£80k	£200k	£500k	£80,500
N2	Design Freeze Delay	25%	£30k	£90k	£220k	£25,417
N3	Planning Condition Discharge Delay	20%	£40k	£120k	£300k	£27,333
N4	Contractor Mobilisation Delay	30%	£50k	£150k	£350k	£50,000
N5	Key Subcontractor Default	15%	£50k	£150k	£400k	£26,250
N6	Material Supply Disruption	20%	£20k	£80k	£250k	£19,667
N7	Productivity Loss — Earthworks	25%	£60k	£180k	£420k	£50,000
N8	Unforeseen Ground Conditions	25%	£100k	£350k	£900k	£100,000
N9	Contamination Discovery	15%	£50k	£200k	£600k	£36,250
N10	Ecological / Protected Species	20%	£30k	£100k	£280k	£23,667
N11	Client Change Order Frequency	40%	£30k	£100k	£300k	£48,667
N12	NEC Early Warning Non-Compliance	20%	£20k	£60k	£150k	£13,667
N13	Compensation Event Disputes	15%	£50k	£180k	£500k	£31,750
N14	Programme Delay Claim	20%	£100k	£300k	£800k	£70,000
N15	Weather / Extreme Rainfall	30%	£20k	£70k	£200k	£25,000

Table 2: The static register. EV sum: £628,168. Clean, comfortable, and structurally blind to dependency.

The Director nodded. I ran the simulation.

3.2 Standard Monte Carlo: What It Showed

Before touching the Bayesian Network, I ran a standard MC on the same 15 risks. 10,000 iterations. LHS. Gaussian copula. No dependency structure.

Metric	Schedule Delay	Cost Exposure
P50 (Median)	94 days	£5,94,732
P80	139 days	£9,28,836
P90 — Planning Ceiling	166 days	£11,18,098
CVaR(90%)	—	£13,43,606

Table 3: Standard MC baseline — independent risks, no cascade logic. P90 accepted as planning ceiling.

Look at the tail. Skewness at 0.63. CVaR(90%) at 2.1× the mean. The model is already telling you something is wrong with the independence assumption — it just cannot tell you which relationships are driving it. It has no chain. It has no memory of cause.

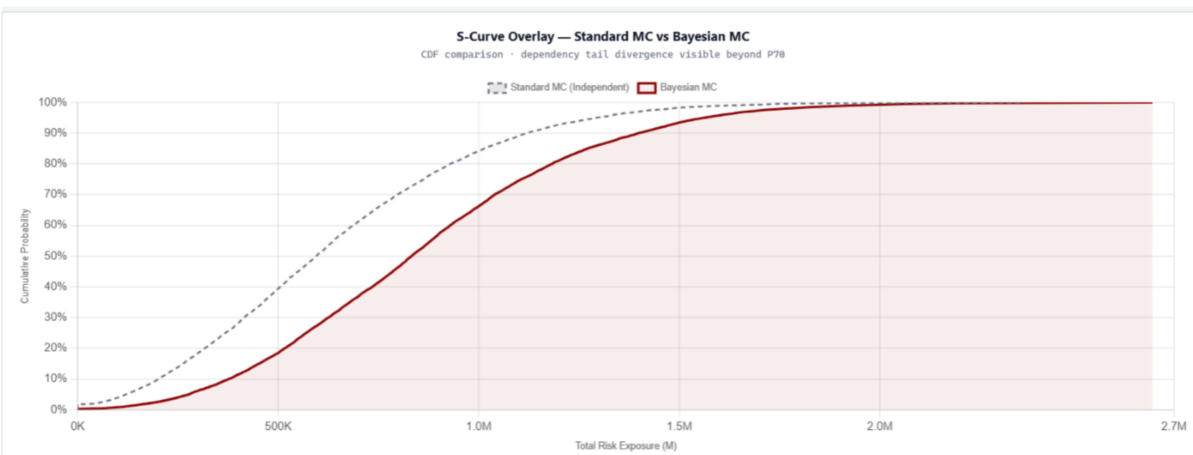

Figure 1: S-Curve Overlay from RiskPulseV12 output. Dashed grey = independent risks. Solid red = cascade propagation. Divergence begins at P75 and widens into the tail.

3.3 The Four Chains Nobody Drew

Three things fired in the first six weeks. All three were on the register. All three had probabilities assigned. The static model treated them as independent coin flips.

They were not coin flips. They were causally connected. And the Bayesian Network — built from structured expert elicitation before mobilisation — already knew that.

- Chain 1: N1 → N2 → N3 — Utility Delay → Design Freeze → Planning Discharge
- Chain 2: N4 → N5 → N6 → N7 — Mob Delay → Subcontractor Default → Supply → Productivity
- Chain 3: N8 → N9, N8 → N10 — Ground Conditions → Contamination + Ecology (dual child)
- Chain 4: N11 → N12 → N13 → N14 — Change Orders → NEC Non-Compliance → Disputes → Claims
- N15 — Weather, standalone but feeding N7 (rainfall amplifies earthworks productivity loss)

Node N7 (Productivity Loss — Earthworks) has three parents: N4 directly, N6 through the supply chain, and N15 through weather. That is why it realised a 40.4% empirical fire rate against a 25% base probability. That node was always going to fire. The static register had absolutely no mechanism to show you that.

3.4 The Cascade Lift Factors: What the Static Model Set to Zero

Dependency Link	P(Child | Parent Safe)	P(Child | Parent Fires)	Lift Factor
N4 → N5: Mob Delay → Subcontractor Default	15.6%	64.0%	4.11×
N8 → N9: Ground → Contamination	15.0%	61.1%	4.07×
N12 → N13: NEC → CE Disputes	15.0%	59.4%	3.96×
N13 → N14: Disputes → Programme Claims	20.2%	66.9%	3.31×
N2 → N3: Design Freeze → Planning	19.3%	54.2%	2.81×
N5 → N6: Subcontractor → Supply Chain	20.5%	58.1%	2.83×
N1 → N2: Utility Delay → Design Freeze	23.9%	60.1%	2.52×
N15 → N7: Weather → Earthworks Productivity	37.0%	48.1%	1.30×

Table 4: Conditional fire rates and cascade lift factors. Rows with 3×+ amplification are the critical chains. The static model assumed all lift factors were 1.0×.

When N4 fires, the probability of N5 does not stay at 15%. It jumps to 64%. That is a 4× amplification on a chain that the static register assumed was independent. They never were independent. Every experienced PM on this project knew, in their gut, that a late contractor tends to shake out the supply chain. What RiskPulseV12 does is put a number on that gut feel — and make the chain visible before it runs.

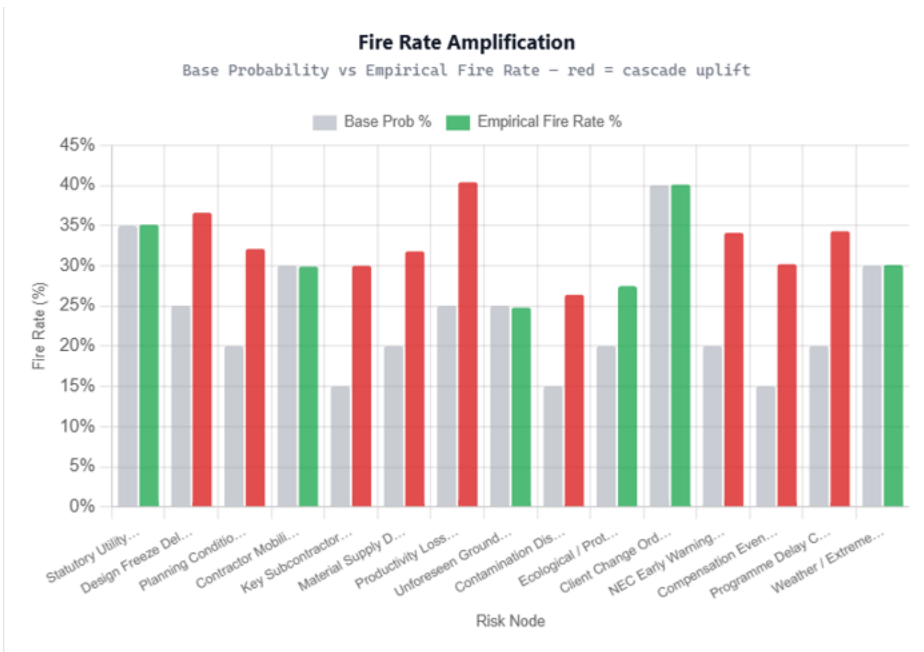

Figure 2: Fire Rate Amplification — RiskPulseV12 output. Grey bars = base probabilities. Green bars = empirical Bayesian fire rates. Red highlights on N5, N6, N7, N12, N13, N14 show cascade uplift nodes.

3.5 Where to Act: The Mitigation What-If

The instinct on every project is to go after the highest probability risk. Here that was N11 — Client Change Orders at 40%. Every PM in the room would have said: lock down the change management process.

The simulation said something different. Go after N4 → N7.

Not because N4 has the highest probability. Because N7 sits at the convergence of three parent paths. Break any one of them before mobilisation, and the tail shrinks measurably. The mitigation what-if analysis quantified exactly how much:

Mitigation Action	P90 Outcome	P90 Saving
No Action	£13,97,284	—
Mitigate N12 → N13 (NEC → Disputes)	£13,20,000	£77,000
Mitigate N8 → N10 (Ground → Ecology)	£12,60,000	£1,37,000
Mitigate N15 → N7 (Weather → Productivity)	£12,00,000	£1,97,000
Mitigate N4 → N7 (Mob Delay → Productivity)	£11,50,000	£2,50,000 SAVED

Table 5: Mitigation What-If P90 sensitivity. N4→N7 generates the highest marginal return on management action.

What the intervention cost: approximately £15,000 in PMC management time. Compressed contractor readiness review. Pre-agreed site access. Pre-positioned traffic management. Plant confirmed five working days before start — not the morning of. NEC Early Warning issued proactively on utility diversion. Change freeze protocol from week two.

What it prevented: the N4 → N5 → N6 → N7 chain running to full depth. N5 did not fire. N6 did not reach its worst scenario. N7 fired at £110,000 actual against a modelled mode of £180,000. The chain existed. We shortened it.

> *£250,000 P90 reduction. £15,000 of management action. 16.7× return on the decision.*

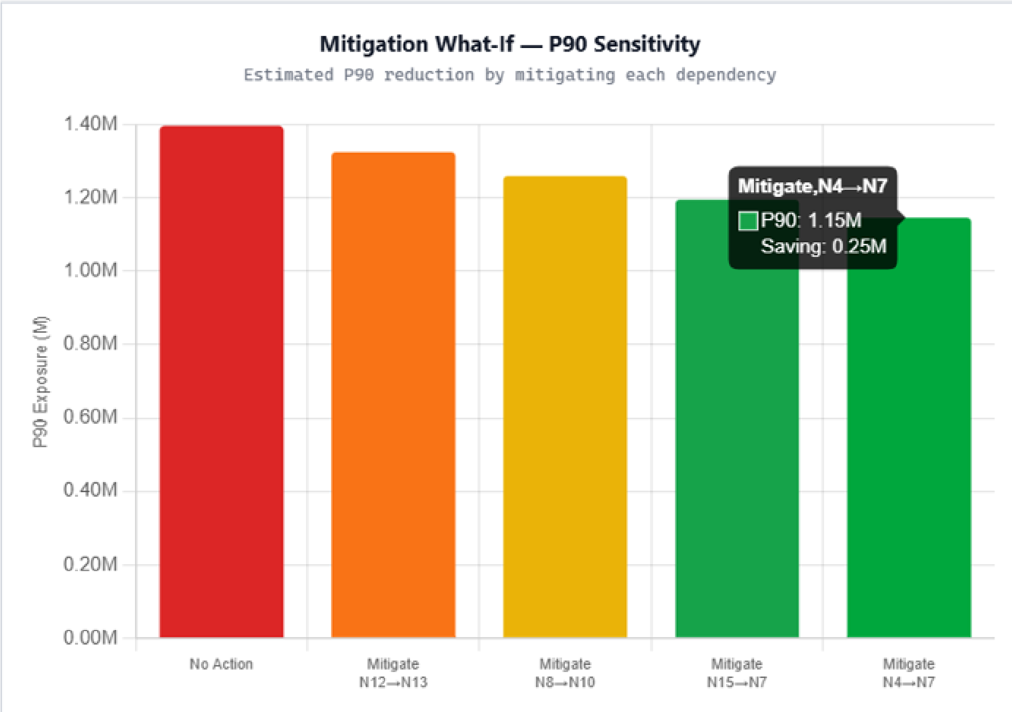

Figure 3: Mitigation What-If P90 sensitivity chart. RiskPulseV12 output. Each bar shows P90 after mitigating that dependency. N4→N7 produces the largest single reduction.

3.6 The Numbers: Standard MC vs BN-Coupled (post chain-break)

Metric	Standard MC (Independent)	BN-Coupled MC (Post-Mitigation)	Δ Saving	% Reduction
Mean	£8,61,594	£7,92,609	−68,986	−8.0%
P80	£11,82,805	£10,96,226	−86,579	−7.3%
P90	£13,97,284	£13,05,518	−91,766	−6.6%
CVaR(90%)	£16,24,393	£15,23,368	−1,01,025	−6.2%

Table 6: Standard MC vs BN-coupled MC comparison. The BN engine changed what the numbers mean, not just what they are.

The standard model said P90 = £13,97,284. With no explanation of which relationships drive it.

The BN model said P90 = £13,05,518 — because we understood and **partially broke the N4→N7 link.**

Those are two very different conversations to have with a Project Director at 3pm in a risk review.

3.7 The Honest Account: What We Got Wrong

The schedule model gave P90 at 166 days. Actual tracking was heading towards 180+ days.

Not because the model was wrong. Because the Project Director's experience on this corridor compressed the schedule inputs before they went into the register. He had done similar works before. He had finished on time before. His optimism was not irrational — it was anchored on evidence from other projects. It was just the wrong evidence for this one.

In the room, before the S-curve existed: "The contract is signed. This is the date. Make it work." I deferred to his experience. That is on me.

The lesson is built into the Intelligence Log. Anchoring bias is now a documented prior. The next project with this risk lead starts from that corrected belief — not from the same optimism.

4. The Theoretical Foundation (Without the Jargon)

I am going to explain the methodology. But I am going to do it the way I would explain it in a risk review, not a textbook. Because the people who need to understand this are project managers and directors, not mathematicians — and the reason quantitative risk thinking has not penetrated the EPC mainstream is partly that the people who understand it refuse to explain it in plain English.

4.1 Why Standard Monte Carlo Is Not Enough

Standard Monte Carlo simulation is a significant improvement over static expected value. You take your risks, sample each one randomly based on its probability and distribution, add up the results, and repeat 10,000 times. The output is a distribution of possible total costs. You can see your P50, P80, P90. You can see the tail.

The problem: every iteration treats each risk as an independent draw. It does not know that when N4 fires, N5 is four times more likely to fire. It does not know that N7 has three parents. It cannot see the chain. So it systematically underestimates the tail — which is exactly the part of the distribution that determines whether your contingency holds.

4.2 What Bayesian Networks Add

A Bayesian Network is simply a map of cause and effect. Each risk is a node. The arrows between nodes represent dependency: if this one fires, it changes the probability of that one. Bayes' theorem — which Thomas Bayes wrote in the 1760s and is still the most underused mathematical idea in project management — tells us exactly how much to update our belief about one event given that another has occurred.

> *P(Child | Parent fires) = P(Parent fires | Child) × P(Child) / P(Parent fires)*

When you couple a Bayesian Network to a Monte Carlo engine, every iteration now carries the dependency structure. When N4 fires in iteration 3,742, the simulation knows to elevate N5's probability from 15% to 64%. The chain runs. The tail gets heavier. And the output is honest.

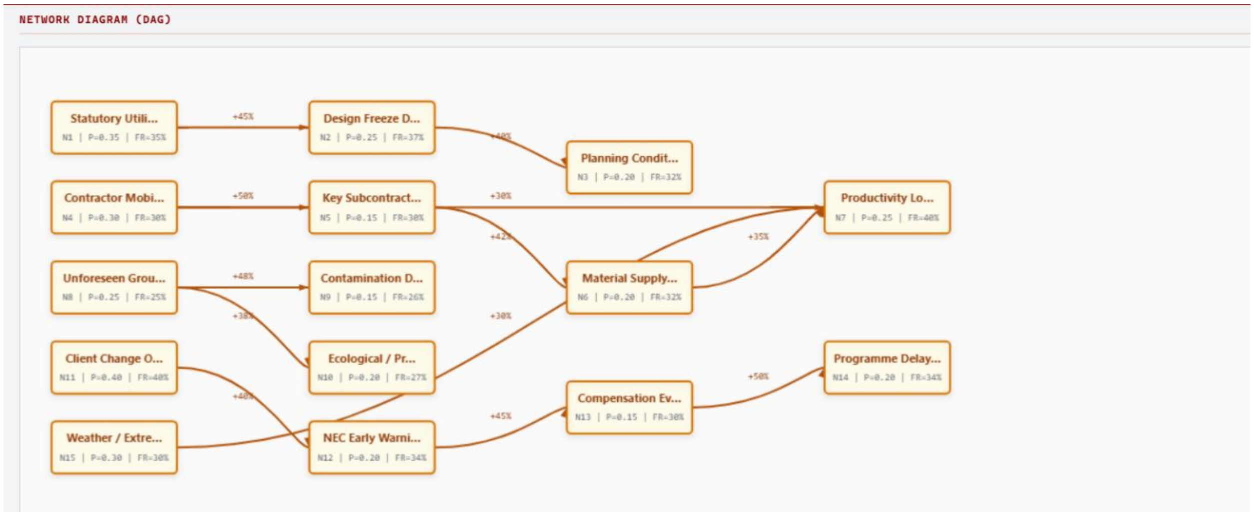

4.3 Latin Hypercube Sampling and Gaussian Copula

Two technical points worth naming. RiskPulseV12 uses Latin Hypercube Sampling (LHS) rather than pure random sampling. LHS ensures that the simulation samples the full range of each risk's distribution proportionally, rather than clustering around the mean. This matters for tail events: if you are running 10,000 iterations with pure random sampling, you might visit the extreme tail once or twice. With LHS, you visit it proportionally to its probability every time. The P90 and CVaR outputs are more stable and more reliable.

Gaussian Copula correlation preserves rank-order dependency between risk variables that share common causes not explicitly modelled in the DAG. It prevents the simulation from treating residual correlation as independence. Together, LHS and Gaussian Copula mean the output converges reliably — RiskPulseV12 confirms P90 drift below 0.06% at 10,000 iterations.

4.4 CVaR: The Number That Actually Matters

P90 is the planning ceiling. CVaR(90%) is the answer to the question nobody asks out loud but every director is thinking: "If we blow through P90, how bad does it actually get?"

CVaR(90%) is the average of all outcomes above P90. It is the expected cost of being in the worst 10% of scenarios.

In the highways case study, P90 was £1.4M but CVaR(90%) was £1.6M — meaning the average worst-case scenario was 16% worse than the planning ceiling. That 16% is what eats contingency reserves alive. And the standard model had no way to show you that.

5. The Intelligence Log: The Thing That Changes Everything

Here is what every commercial risk tool does at project close: nothing. You get your P90. The project ends. The file is archived. The beliefs that were wrong — the overestimated impacts, the underestimated probabilities, the dependencies nobody modelled — are reset to zero. The next project starts from scratch.

This is why every new project team makes the same class of errors as every previous one. Not because they are bad at their jobs. Because the industry has no institutional memory of risk outcomes.

> *The Risk Intelligence Log does not throw away the posterior that reality handed you.*

Every risk that fires is logged against its prior. Every risk that does not fire is logged too. The posterior probability for each node is updated using Bayesian inference. The prior mode impact is compared to the actual. The gap becomes calibration data.

Risk	Prior P	Posterior P	Prior Mode Impact	Actual Impact	Fired?
Statutory Utility Diversion Delay	35%	39.0%	£200,000	£95,000	YES
Design Freeze Delay	25%	35.2%	£90,000	£62,000	YES
Planning Condition Discharge Delay	20%	14.9%	£120,000	—	NO
Contractor Mobilisation Delay	30%	33.6%	£150,000	£85,000	YES
Key Subcontractor Default	15%	14.4%	£150,000	—	NO — mitigation held
Material Supply Disruption	20%	16.9%	£80,000	—	NO
Productivity Loss — Earthworks	25%	30.2%	£180,000	£110,000	YES
Unforeseen Ground Conditions	25%	22.3%	£350,000	—	NO
Client Change Order Frequency	40%	37.6%	£100,000	£48,000	YES
NEC Early Warning Non-Compliance	20%	20.4%	£60,000	—	NO
Compensation Event Disputes	15%	12.2%	£180,000	—	NO
Programme Delay Claim	20%	18.8%	£300,000	—	NO

Table 7: Intelligence Log extract. Prior vs posterior probabilities and impact calibration.

When Statutory Utility Diversion Delay fires at £95,000 and my prior mode was £200,000, that gap does not disappear. The posterior mode comes down to £110,341.

The next highways corridor project I work starts from that corrected prior — not from the industry default that has been over-estimating this risk for decades.

This is Bayesian inference as a professional discipline. Thomas Bayes wrote the theorem in the 1760s. What RiskPulseV12 does is put it into a workflow — where every risk register has historically been a static document that forgets everything the moment the project closes.

Sixteen thousand projects across eight sectors. Every outcome logged. Every prior corrected. Every posterior carried forward. ***That is not a tool. That is institutional risk memory.***

6. The Core Philosophy: Making the Invisible Visible

I want to be direct about why I built this. Not for the tool's sake. Not to compete with @Risk or Oracle Primavera Risk Analysis. Those tools do standard Monte Carlo well. They are not trying to do what RiskPulseV12 does.

I built it because I kept watching smart, experienced project professionals make the same mistakes — not because they were incompetent, but because the models they were working from were structurally blind to the risks that actually blow projects.

> *Once the data is in your head and in a sheet, you will see possibilities and mitigations you never saw before. The mind manages what it can see.*

This is the core thesis. When you have never run a simulation on your risk register, your mental model of the project is built from experience, intuition, and the narrative in the risk review. These are valuable. They are also limited. They are anchored to the projects you have already seen. They are blind to the interactions you have not modelled.

The moment you run RiskPulseV12 on your register — even a simple 15-node run — something changes. You see your P90 is not £628,168. It is £1.4M. You see that N7 has three parents. You see that breaking one link for £15,000 moves the P90 by £250,000. You see the chain.

And now you cannot unsee it. The mitigation conversations change. The contract terms you negotiate change. The contractor readiness review you commission before mobilisation changes. The NEC Early Warning you issue proactively changes. All of those changes happen because the simulation made the invisible visible.

We are running probabilistic projects with deterministic minds.

The gap between those two things — between what the model can see and what our intuition can see — is exactly where projects get into trouble. RiskPulseV12 exists to close that gap.

6.1 Why This Matters for Every Infrastructure PM, Everywhere

This is not an India problem. This is not a UK problem. This is a structural problem in how the global infrastructure industry manages risk. From highways corridors in England to solar parks in Rajasthan to rail upgrades in Southeast Asia to water infrastructure in sub-Saharan Africa — the pattern is the same. The delivery environments are complex. The risk tools applied to them are not.

The Indian EPC sector is a case in point — and one I know well. Solar parks in politically complex land acquisition environments. Highway corridors through contested ground conditions. Urban infrastructure under regulatory frameworks that change mid-project. The cascade risk patterns are not simpler than international markets — they are often more complex. And yet the quantitative risk toolkit available to most Indian EPC project managers is: a risk register, a RAG chart, and an expected value calculation.

But the same gap exists in the UK, across Europe, in the Middle East, in Southeast Asia. The sophistication gap between the complexity of the delivery environment and the sophistication of the risk tools applied to it is, in my view, one of the primary structural drivers of infrastructure cost and schedule overrun worldwide. The UK's Infrastructure and Projects Authority has documented this pattern repeatedly. So has every national audit office that has ever reviewed a major project portfolio.

That is what this whitepaper is for. Not to sell a product. To make the case — with evidence, with a live case study, with the numbers — that better tools exist, that they are accessible, and that using them produces materially different project outcomes. Regardless of which country you are delivering in.

7. Limitations and Epistemic Honesty

A whitepaper that only reports wins is a marketing document. This is a technical contribution. So here are the limitations, stated clearly:

- All base probabilities are subjective expert estimates cross-referenced against the calibration dataset. They are informed judgments, not actuarial frequencies derived from statistically controlled experiments.
- Conditional probability parameters (influence factors 0.30–0.50) have been calibrated against the 16,000-project dataset but represent empirical patterns, not physical laws. Project-specific conditions may differ significantly.
- The Intelligence Log uses conjugate Bayesian updating — a simplified approach appropriate for live workflow integration. Full MCMC posterior sampling would produce more precise posteriors but at a computational and process cost inappropriate for project delivery environments.
- The schedule model in the case study used the same risk register as the cost model. No separate CPM network integration was performed. Schedule-cost correlation is implicit, not explicitly modelled.
- Optimism bias in expert elicitation cannot be fully eliminated by process design. The schedule failure documented in Section 3.7 is a direct consequence of anchoring on the Project Director's experience. The Intelligence Log corrects this over time. It does not prevent it on first engagement.
- Scenario testing outputs (regime change, contractor substitution, supply chain shock) are sensitivity analyses based on adjusted input parameters. They are not predictions. They are structured explorations of how the P90 moves when conditions change.

8. How to Apply PCA to Your Next Project

If everything in this paper makes sense to you but you are not sure where to start, here is the procedure. I have kept it to six steps because that is what fits on a single page and what a PM can realistically execute within a two-week risk review cycle.

Step 1: Build your risk register the way you normally would. Fifteen risks, twenty, fifty — it does not matter. PCA does not require you to change your register format. It requires you to do one thing differently: stop assuming independence.

Step 2: Map the chains. Sit with your project team and ask: if this risk fires, which other risks become more likely? Draw the arrows. You do not need software for this step. You need a whiteboard and an honest conversation. The output is a directed acyclic graph — a map of cause and effect. Most teams can do this in an hour.

Step 3: Estimate the conditional fire rates. For each arrow, ask: if the parent fires, what does the child's probability become? Use structured expert elicitation. Cross-reference against the PCA calibration dataset where your sector is represented. The output is an influence factor for each dependency link.

Step 4: Run the BN-coupled Monte Carlo simulation. 10,000 iterations. Latin Hypercube Sampling. The Bayesian Network carries the dependency structure into every iteration. The output is an honest P50, P80, P90, and CVaR — one that accounts for cascade amplification, not one that assumes it away.

Step 5: Read the lift factors. Run the mitigation what-if. Do not go after the highest-probability risk. Go after the node where breaking a single dependency link produces the largest P90 reduction. That is where your management attention earns its return. The case study in this paper showed 16.7× return on a £15,000 intervention. Your project will have its own version of that number. Find it.

Step 6: Log everything. Carry the posteriors forward. When the project closes, record which risks fired, at what impact, against what prior. Update the posterior probabilities. Update the impact distributions. The next project starts from corrected beliefs, not from the same defaults the industry has been recycling for decades. This is the step that nobody does. This is the step that turns PCA from a one-off simulation into a professional discipline.

That is PCA. Six steps. One methodology. And a fundamentally different conversation in your next risk review.

9. Conclusion

We are running a probabilistic world with deterministic thinking. Every risk register that sums independent expected values is an act of mathematical optimism. Every project that runs without a simulation has a gap between what it can see and what is actually happening in the probability space.

That gap is where the overruns live.

RiskPulseV12 does not eliminate that gap entirely. No tool can. But it closes it materially — by making the chain visible, by quantifying the lift factors, by identifying the single intervention that moves the P90 the most, and by logging every outcome so the next project starts from evidence rather than hope.

The UK highways case study showed a 16.7× return on a targeted pre-mobilisation intervention identified through cascade analysis. The schedule overran. The contingency held. Both outcomes are documented. One was predicted. Both were learned from.

Calibrated across 16,000 projects. Running scenario tests that standard tools cannot run. Carrying an Intelligence Log that builds institutional memory where the industry currently has none.

This whitepaper was written for one reason: to make the case that quantitative risk thinking is not an academic luxury for the infrastructure sector — in India, the UK, the Middle East, Southeast Asia, or anywhere else where complex projects are delivered under uncertainty. It is the difference between a contingency that holds and one that gets consumed in the first quarter of execution.

> *Time, Cost and Scope don't run your project. The risks associated with them do.*

If you are running a complex infrastructure project and you have never run a simulation on your risk register — this paper is for you. Start with the chains. The rest follows.

References

Bayes, T. (1763). An Essay Towards Solving a Problem in the Doctrine of Chances. Philosophical Transactions of the Royal Society of London.

Flyvbjerg, B. (2021). Top Ten Behavioral Biases in Project Management: An Overview. Project Management Journal, 52(6). (Reference dataset: 16,000+ megaprojects across sectors.)

Iman, R. L., & Conover, W. J. (1982). A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables. Communications in Statistics — Simulation and Computation, 11(3).

Infrastructure and Projects Authority (IPA). (2024). Annual Report on Major Projects. UK Government Cabinet Office. Public record.

National Audit Office (UK). (Various, 2018–2025). Reports on Major Road Projects and Highways England Performance. Published under Open Government Licence.

National Highways (formerly Highways England). (Various, 2019–2025). Strategic Road Network Performance Specifications, Project Outturn Data, and Annual Reports. Public records.

PMI. (2021). A Guide to the Project Management Body of Knowledge (PMBOK Guide), 7th Edition. Project Management Institute.

SECI / MNRE Project Outcome Database (2018–2025). Public procurement and project performance records. Solar Energy Corporation of India / Ministry of New and Renewable Energy.

Government of India Infrastructure Project Outcomes Data (various, 2015–2025). Ministry of Road Transport and Highways; Ministry of Housing and Urban Affairs. Public records.

Vose, D. (2008). Risk Analysis: A Quantitative Guide (3rd ed.). Wiley.

Woléjszo, D., & Wallach, D. (2026). Critical Effort Method: A Framework for Identifying Execution Difficulty in Project Scheduling. Methodology White Paper.

About the Author

Nikhil Dhand, PMP

Senior Project Management Consultant with 11 years of cross-sector experience in solar EPC, highways, urban infrastructure, and complex project delivery across India, the UK, and international markets.

Creator of RiskPulseV12 — a Bayesian Monte Carlo risk engine incorporating Bayesian Networks, Iman-Conover correlation, Gaussian Copula, Benford's Law anomaly detection, Markov Chain state transition modelling, and Pareto multi-objective optimisation. The engine is calibrated across 16,000+ infrastructure projects spanning 8 sectors, sourced entirely from public databases.

Author of Quantitative Risk Management for Indian EPC (forthcoming, Amazon KDP). A second book — Day Minus One — is in development, focused on correlated risk mapping at project inception: the moment before anything has gone wrong, when everything still can.

His mission is singular: to raise the floor of quantitative risk thinking across the global infrastructure sector — starting with Indian EPC, where the gap is widest, but extending to every market where projects are delivered under uncertainty. Not because the tools are interesting — though they are — but because projects overrun when they shouldn't, and the methodology to prevent it exists and is not being used.

All work published under the Probabilistic Chain Analysis (PCA) and RiskPulseV12 brand is independent and original. No organisational affiliation is implied or represented.

Contact:
For collaboration, consulting enquiries, or to discuss applying RiskPulseV12 to your project portfolio, reach out via LinkedIn or email.

Nikhil Dhand | LinkedIn | itsnik6850@gmail.com