The £100 Billion Blind Spot: A Probabilistic Chain Analysis™ Autopsy of HS2 [1, 2]

Nikhil Dhand, PMP

Abstract

High Speed Two (HS2) is a UK government-commissioned high-speed rail programme connecting London to Birmingham and the North, with a 2013 baseline cost of £32.7 billion. Against an original 2026 opening target, the programme has accumulated delays exceeding a decade and costs now exceeding £100 billion.

In October 2013, HS2 Ltd published its Cost and Risk Status Report (Document S_A_8) — the document that defined the risk framework for the largest infrastructure project in UK history. One decade and £60 billion in overrun later, its assurance language reads not as confidence but as caption.

This paper dissects the verbatim text of two foundational HS2 risk documents — S_A_8 and the Risk Analysis Technical Documentation (S_A_2) — clause by clause. It identifies six structural deficiencies in the risk methodology that made the overrun mathematically inevitable: independence-assumed Monte Carlo, interview-based risk identification, fixed schedule across all risk scenarios, mitigation illusion, PSC register aggregation, and assurance-oriented rather than outcome-oriented risk framing.

It introduces Probabilistic Chain Analysis (PCA) as the corrective framework and applies it retroactively to the 2013 baseline through a seven-node conditional probability chain — demonstrating why a risk register built on independence assumptions was structurally incapable of forecasting the cascade that occurred.

The gap between what HS2 believed it was funding (P90) and what Flyvbjerg's reference class distribution indicates it was actually funding (approximately P60–P66) is not

[1] How to cite this paper: Dhand, N. (2026). The £100 Billion Blind Spot: A Probabilistic Chain Analysis™ Autopsy of HS2; *PM World Journal*, Vol. XV, Issue VI, June.

[2] Independent publication. This paper is authored solely by Nikhil Dhand in a personal capacity. All source data is from publicly available UK government publications, NAO audit reports, and academic papers. The PCA methodology and RiskPulse V12 engine are the original intellectual property of Nikhil Dhand. Chain diagrams in Sections 3–5 are the author's probabilistic inference — not causal claims made by the source reports. Funding: This research received no external funding.

statistical noise. It is the mathematical consequence of treating correlated risks as independent events — precisely what this paper measures through PCA.

Keywords: *HS2, Probabilistic Chain Analysis, Bayesian Networks, Monte Carlo Simulation, Risk Register, Cascade Risk, Quantified Cost Risk Analysis, Infrastructure Megaprojects, Cost Overrun, UK Transport*

1. Introduction: The Exhibit

I want to start with the exact words. Not a paraphrase. Not a summary. The exact words that were written in October 2013 and submitted to Parliament as the risk assurance for a £32.7 billion programme.

Clause 4.4.4 — S_A_8: The Economic Case for HS2 — Cost and Risk Status Report (October 2013)

"The risk register for Phase One was developed through a series of interviews and reviews, and informed by risk registers emerging from the Professional Service Contractors (PSCs). The Phase One risk register contains over 300 active threats with a cost impact and a mitigation plan. These are under continual review and challenge to gain assurance that risks are being effectively managed to provide value for money."

Read that again slowly. Every word in this clause is a confession of methodological failure dressed in the language of rigour.

Three paragraphs above Clause 4.4.4, sitting quietly in the same document, is the mathematical confession that makes everything else fall into place:

4.4.4 The risk register for Phase One was developed through a series of interviews and reviews, and informed by risk registers emerging from the Professional Service Contractors (PSCs). The Phase One risk register contains over 300 active threats with a cost impact and a mitigation plan. These are under continual review and challenge to gain assurance that risks are being effectively managed to provide value for money.

Clause 4.2.3 — S_A_8: The Economic Case for HS2 — Cost and Risk Status Report (October 2013)

"The threats are generally modelled as independent events, with covariance introduced between some related threats where it is deemed appropriate."

When I first read Clause 4.2.3, I read it twice. Then I went back to Clause 4.4.4 and read that again. The 300 active threats. The continual review. The assurance. Three paragraphs away from the independence assumption. I spent three days just sitting with those two clauses before I wrote a single word of this paper.

"Generally modelled as independent events." That is a mathematical statement with a mathematical consequence. Independence assumption in a Monte Carlo model means each risk is sampled separately — its occurrence does not update the probability of any other risk occurring.

This is statistically valid only if risks are truly independent. In a mega-infrastructure programme, no risks are truly independent. Every ground condition risk is correlated with programme risk. Every contractor financial risk is correlated with supply chain risk. Every regulatory risk is correlated with scope risk.

The covariance correction — "where deemed appropriate" — was added manually, subjectively, and only where the team consciously noticed a relationship. PCA's Bayesian network structure makes covariance an architectural feature, not an afterthought.

4.2.3 A threat in the risk register is characterised by a probability that the threat occurs and an impact (on cost, time, reputation) if it does. The cost impact is generally expressed as a range (for example, £5 million to £25 million) which represents the potential extra-over costs associated with the threat event. The threats are generally modelled as independent events, with covariance introduced between some related threats where it is deemed appropriate. The threat probabilities and cost impacts represent expert judgement relative to the allowance already included in the base cost estimate. The threats produce a range of costs.

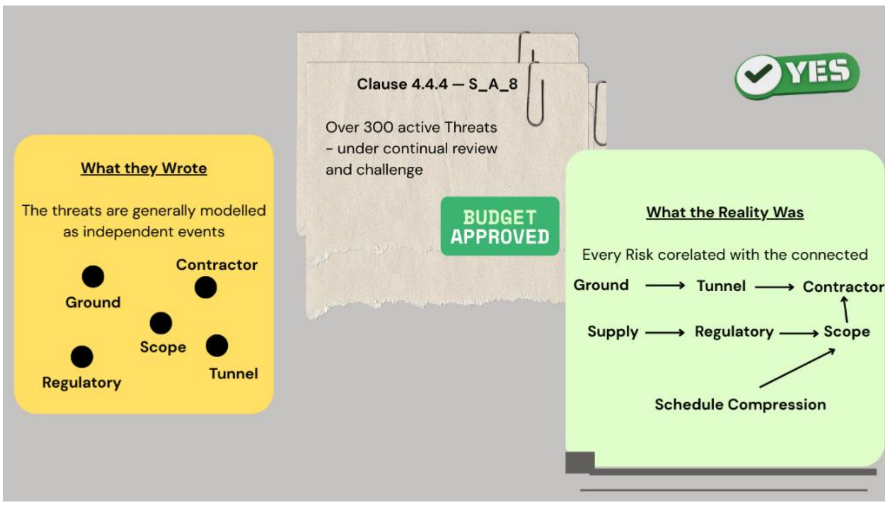

2. The Document That Was Not What Its Title Said

The second document I want to put under the microscope is S_A_2: Risk Analysis for the HS2 Economic Case — Technical Documentation. A reader unfamiliar with the field would assume this document explains how construction, delivery, and programme risks were modelled. It does not.

In 28 pages, construction cost risk receives **exactly two paragraphs and two graphs**. The remaining 26 pages are entirely about demand forecasting uncertainty — how sensitive the Benefit-Cost Ratio is to GDP growth, values of time, rail demand elasticities, and market maturity. The word "risk" in the title refers exclusively to benefit risk, not delivery risk. This is the first structural deception, and it is not an accident.

Here is the architecture of the document, laid bare:

Section	Content	Pages
Section 2	BCR distribution calculation — GDP growth, demand equations, meta-model regression	12
Sections 3.1–3.4	GDP growth distributions, non-business value of time, fare and GDP elasticities	5
Section 3.5	**Construction cost — the only section that touches delivery risk**	**2**
Sections 3.6–3.7	Alternative values of time, aviation market maturity assumptions	9

Table 1: Architecture of S_A_2. Delivery risk occupies 2 of 28 pages.

Two pages out of twenty-eight. This is not an oversight in proportionality. It reflects a foundational choice: the purpose of this document was to defend the economic case for building HS2, not to honestly model the risk of cost overrun during building it. The two are radically different analytical objectives, and this document systematically conflates them. Page counts based on author's content analysis of S_A_2 document structure.

The two graphs in Section 3.5 present cost distributions in isolation — with no linkage to schedule, no conditional dependencies, and no representation of how construction risk propagates through the programme. They are analytically self-contained and structurally disconnected from the reality of how megaproject costs escalate.

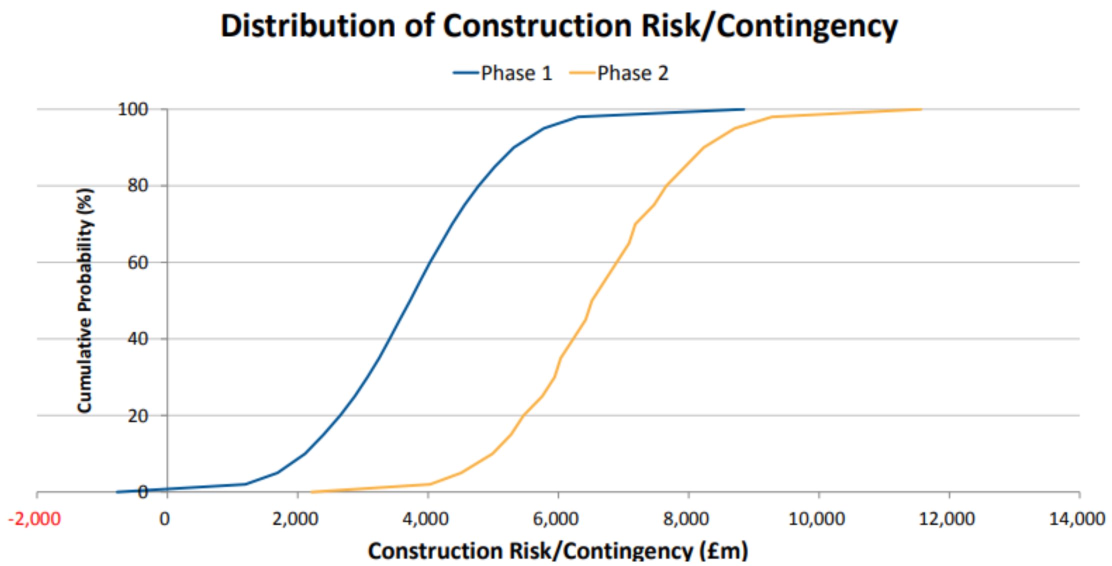

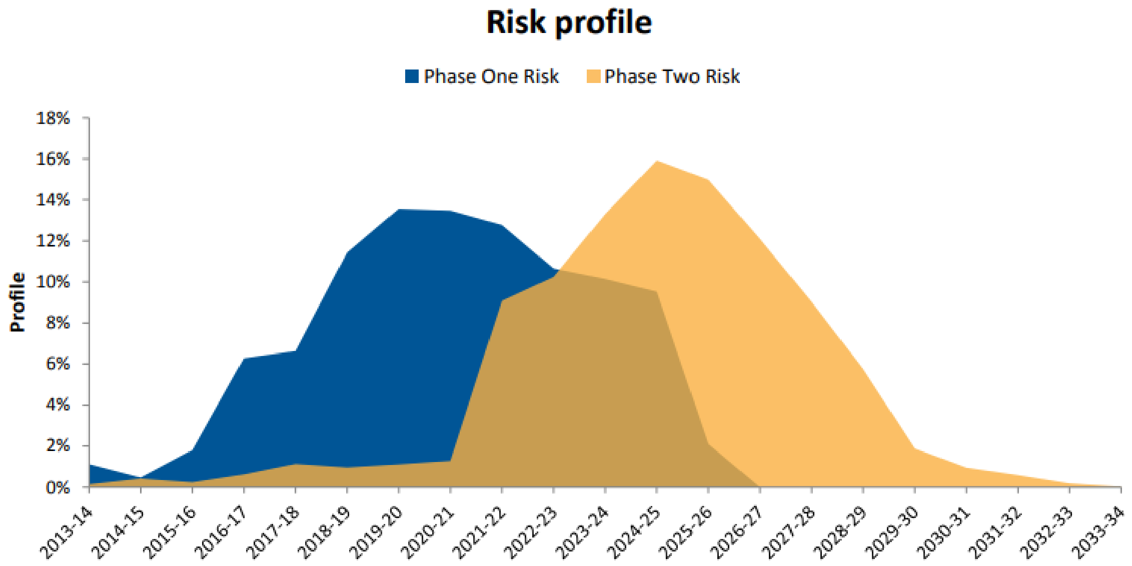

2.1 The Variable Selection Confession

Clause 3.1.1 — S_A_2: Risk Analysis Technical Documentation

"The factors that have been selected for inclusion as variables in the risk analysis have been chosen on the basis that: 1. They are key drivers of the BCR; and 2. There are reliable sources of information for the parameter and its distribution."

And then, buried in a subordinate clause:

Clause 3.1.2 — S_A_2: Risk Analysis Technical Documentation

"Clearly, these are not the only factors in the calculation of the BCR but they exert a strong influence over the results — particularly GDP growth — and they are the factors that are most amenable to analysis within our framework."

This is the most honest sentence in the entire document. The variables were not chosen because they were the most important drivers of project outcome. They were chosen because they were most amenable to analysis within the existing framework. That is methodological selection bias stated openly and without apparent embarrassment.

In project risk terms, this means: "We modelled what we could easily model, not what mattered most."

The variables excluded — contractor financial stability, supply chain capacity, geological uncertainty, scope definition maturity, procurement market conditions, regulatory change probability — are excluded not because they are unimportant but because the authors had no distribution parameters for them.

PCA directly addresses this. It is explicitly designed to incorporate qualitatively-estimated conditional probabilities when data is sparse. Bayesian prior elicitation allows expert judgment to be formalised as a distribution even when historical data does not exist. The refusal to model difficult risks on the grounds of data unavailability is not analytical rigour — it is analytical avoidance.

2.2 GDP as the "Most Critical Input"

Clause 3.2.1 — S_A_2: Risk Analysis Technical Documentation

"In the appraisal of HS2, economic growth determines both how quickly demand grows in the model and how people value travel time-savings from the scheme. It is therefore one of the most critical inputs into the risk analysis."

Pause here.

3.2.1 In the appraisal of HS2, economic growth determines both how quickly demand grows in the model and how people value travel time-savings from the scheme. It is therefore one of the most critical inputs into the risk analysis.

GDP growth is described as the most critical input into the risk analysis.

Not geological uncertainty at the Chilterns. Not contractor market capacity. Not inflation in the construction materials supply chain. Not the risk that Phase 2 would never be built. GDP growth — a macroeconomic variable that HS2 Ltd has zero ability to influence and that affects benefits, not costs — is positioned as the centre of gravity of the entire risk framework.

This reveals the document's true purpose. It is not a risk analysis of HS2 the project. It is a risk analysis of HS2 the investment case. That is a legitimate analytical exercise. But calling it a "risk analysis" when what you mean is "BCR sensitivity to macroeconomic assumptions" is a category error that misrepresents the nature of delivery risk to every stakeholder who reads this document: Parliament, Treasury, the public.

2.3 The Sentence That Defines the Entire Problem

Clause 3.5.7 — S_A_2: Risk Analysis Technical Documentation

"In order to calculate the present value of construction cost, a profile through time of the cost is required. The project spend over time for both phases is illustrated in Figure 6. The risk profile component of project spend is assumed to be the same regardless of the level of risk. In all scenarios, the opening years remain the same."

3.5.7 In order to calculate the present value of construction cost, a profile through time of the cost is required. The project spend over time for both phases is illustrated in Figure 6. The risk profile component of project spend is assumed to be the same regardless of the level of risk. In all scenarios, the opening years remain the same.

Read the third sentence again: *"The risk profile component of project spend is assumed to be the same regardless of the level of risk."*

This means that regardless of whether the model draws a low-risk scenario or a high-risk scenario, the timing profile of when money is spent is held constant. Whether costs are at P5 or P95, the project is assumed to spend money at the same rate in the same years.

Why is this catastrophic? **Because in reality, cost escalation and schedule delay are inseparable**. A high-risk scenario is not one where you spend more money in the same years — it is one where you spend more money over more years. A two-year delay to Phase One opening means two additional years of preliminaries, site overhead, financing costs, inflation exposure on unspent contingency, and contractor remobilisation. None of this is captured when the spend profile is fixed regardless of risk level.

The final sentence — *"In all scenarios, the opening years remain the same"* — is the single most audacious assumption in the entire document. It states that regardless of risk, HS2 opens on schedule. The opening date is not a variable. It is a fixed input. Schedule risk does not exist in this model. Against its original 2026 opening target, HS2 is now tracking delays exceeding a decade — and Phase 1 has still not opened. The model literally could not compute this outcome — it was outside the model's universe.

2.4 Phase Two: A 32% Flat Addition as "Risk Modelling"

Clause 3.5.5 — S_A_2: Risk Analysis Technical Documentation

"For Phase Two, a less well developed QCRA is in place (reflecting the stage of the scheme). Costs used in the standard case for Phase Two therefore include an additional provision of 32%. For the purpose of the risk analysis the P100 and P0 have been assumed to be equal to the minimum and maximum cost taken from the QCRA results."

Three catastrophic failures in two sentences.

First: Phase Two received no proper QCRA. Instead of acknowledging this as a fundamental data gap, the authors applied a flat 32% provision. This is not probabilistic modelling — it is a scalar adjustment. There is no P50, no P80, no P95 for Phase Two in its own right.

Second: P100 and P0 are defined as the QCRA minimum and maximum. This bounding assumption eliminates tail risk by definition. It assumes the worst case has already been modelled. In practice, the QCRA maximum for Phase Two was approximately £18.7 billion. The actual Phase Two cost by 2023 was tracking toward £50+ billion before cancellation. The QCRA maximum was not the worst case — it was approximately the P35 of the true distribution.

Third: Phase One and Phase Two construction costs are modelled as separate inputs. There is no modelling of the interaction between Phase One delay and Phase Two cost — specifically, that a delayed Phase One would increase Phase Two programme costs through inflation exposure, rescheduled labour, and supply chain disruption. PCA makes this interaction explicit.

3. Six Structural Failures — Dissected From the Source Text

When Clauses 4.4.4 and 4.2.3 of S_A_8 are read alongside the entire text of S_A_2, a taxonomy of six distinct structural failures emerges. Each failure is documented verbatim. Each one maps to a specific PCA correction.

Sl.	Failure	Source Clause	What It Means	PCA Correction
F1	Interview-based risk identification	S_A_8, Cl. 4.4.4: "developed through a series of interviews and reviews"	Captures named risks, not propagation pathways. Interviewees think in silos.	Conditional probability elicitation: "Given Risk A fires, what is the updated probability of Risk B?"
F2	PSC register aggregation fallacy	S_A_8, Cl. 4.4.4: "informed by risk registers emerging from the PSCs"	Stitching 10 contractor registers creates a list, not a system model. Covariance between contractors ignored.	Bayesian Network graph: each PSC risk is a node, directed edges encode conditional dependencies.
F3	Independence assumption	S_A_8, Cl. 4.2.3: "generally modelled as independent events"	Monte Carlo samples each risk separately. Mathematically underestimates total risk when risks are correlated.	Iman-Conover / Gaussian Copula correlation engine. Covariance is architectural, not afterthought.
F4	Mitigation illusion	S_A_8, Cl. 4.4.4: "with a cost impact and a mitigation plan"	Mitigation plan listed but post-mitigation probability rarely recalculated. Register assumes mitigation works.	Mitigation modelled as conditional intervention node. If mitigation fails (its own probability), cascade resumes.
F5	Fixed schedule across all risk scenarios	S_A_2, Cl. 3.5.7: "In all scenarios, the opening years remain the same"	Opening date is a fixed input, not a variable. Schedule risk does not exist in the model. 13-year delay was outside the model's universe.	Schedule is an output of the risk model. Cost and schedule are joint probability chains, not marginal.
F6	Assurance orientation vs. forecasting	S_A_8, Cl. 4.4.4: "to gain assurance that risks are being effectively managed"	Risk register exists to confirm the team is doing the right things, not to estimate the probability of being within budget.	PCA is outcome-focused: it produces P50/P80/P90 probability distributions of programme outcomes.

Table 2: Six structural failures identified from verbatim source text, with PCA corrections

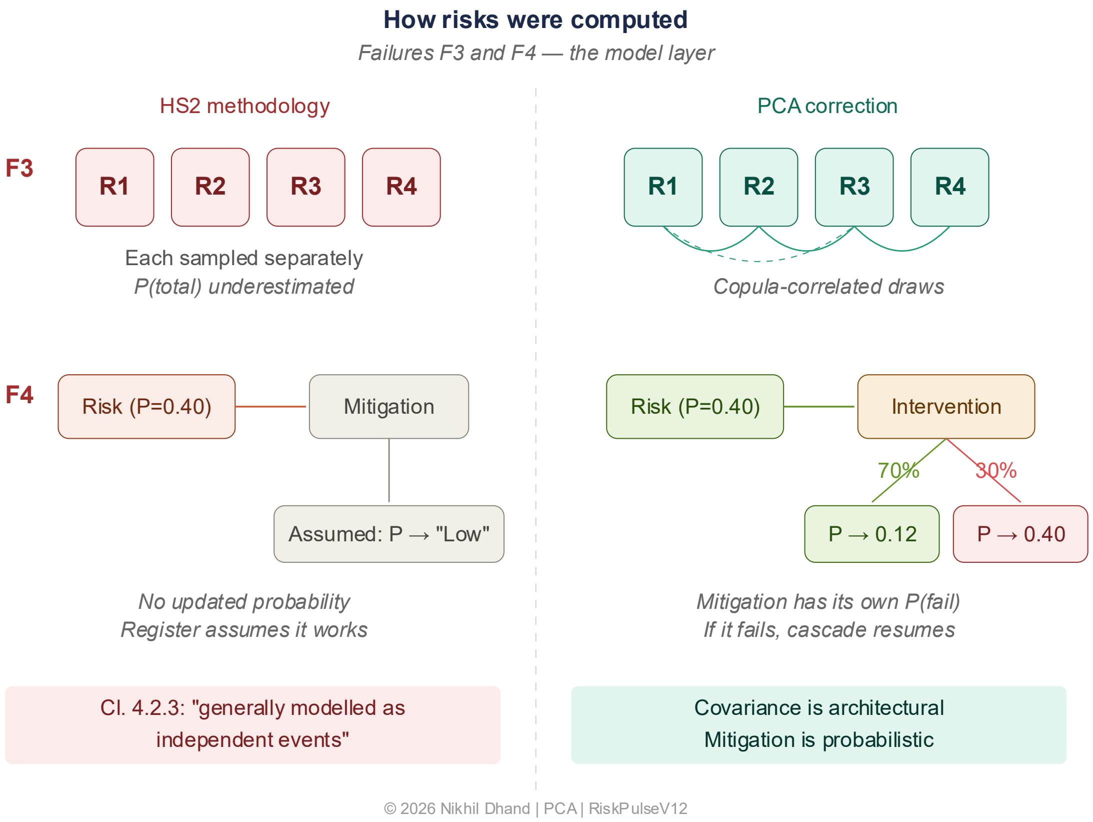
How risks were computed
Failures F3 and F4 — the model layer
HS2 methodology
PCA correction
F3
R1
R2
R3
R4
Each sampled separately
P(total) underestimated
R1
R2
R3
R4
Copula-correlated draws
F4
Risk (P=0.40)
Mitigation
Assumed: P → "Low"
No updated probability
Register assumes it works
Risk (P=0.40)
Intervention
70%
30%
P → 0.12
P → 0.40
Mitigation has its own P(fail)
If it fails, cascade resumes
Cl. 4.2.3: "generally modelled as independent events"
Covariance is architectural
Mitigation is probabilistic
© 2026 Nikhil Dhand | PCA | RiskPulseV12

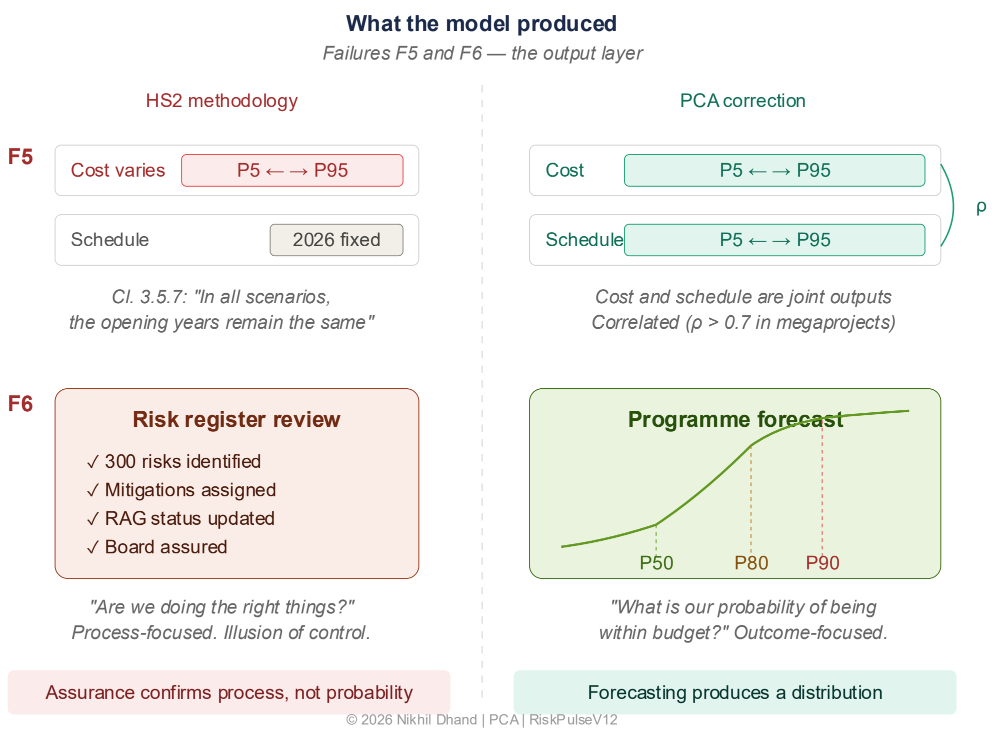
What the model produced
Failures F5 and F6 — the output layer
HS2 methodology
PCA correction
F5
Cost varies
P5 ←→ P95
Schedule
2026 fixed
Cl. 3.5.7: "In all scenarios, the opening years remain the same"
Cost
P5 ←→ P95
Schedule
P5 ←→ P95
ρ
Cost and schedule are joint outputs
Correlated (ρ > 0.7 in megaprojects)
F6
Risk register review
✓ 300 risks identified
✓ Mitigations assigned
✓ RAG status updated
✓ Board assured
"Are we doing the right things?"
Process-focused. Illusion of control.
Programme forecast
P50
P80
P90
"What is our probability of being within budget?" Outcome-focused.
Assurance confirms process, not probability
Forecasting produces a distribution
© 2026 Nikhil Dhand | PCA | RiskPulseV12

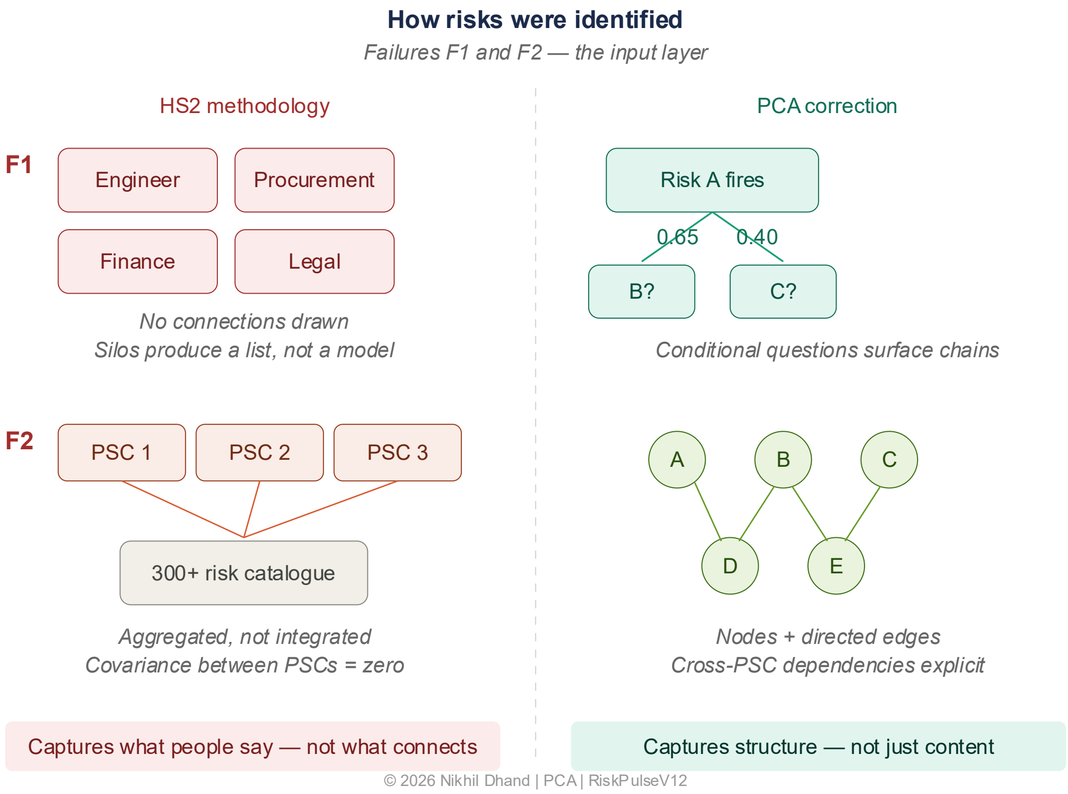

3.1 The Mitigation Illusion — A Closer Look

"Contains over 300 active threats with a cost impact and a mitigation plan." That sounds thorough. Three hundred risks, each with a probability, a cost range, and an action plan. But here is what is happening beneath the surface.

A mitigation plan does not reduce a risk's probability in the model unless that reduction is explicitly re-encoded as an updated prior. In practice, risk registers list a mitigation action and then assume it works. The probability on the register is the pre-mitigation probability. The mitigation plan is described in text. But the post-mitigation updated probability is rarely recalculated — and even if it is, it is done qualitatively, dropping from "High" to "Medium," not probabilistically.

In PCA, mitigation is modelled as a conditional intervention node: it changes the conditional probability table of the downstream risk nodes. If the mitigation fails (which has its own probability), the cascade resumes from the unmitigated state. This is a fundamentally more honest representation.

3.2 The Assurance Psychology

The language "to gain assurance that risks are being effectively managed" reveals a fundamental cognitive orientation. The risk register exists to confirm that the team is doing the right things, not to estimate the probability of being within budget.

This is the distinction between process-focused risk management and outcome-focused probabilistic forecasting. A 300-item risk register that is "under continual review and challenge" creates a powerful psychological illusion of control. Each review meeting, each RAG status update, each mitigation action recorded — all of these generate confidence signals that inhibit escalation of concern even as the probabilistic reality deteriorates.

Kahneman called this the planning fallacy — not a shortage of data, but the human tendency to suppress distributional thinking in favour of scenario thinking. The team believes they are managing their way through the risk list, one mitigation at a time, while the probabilistic reality — the chain of compounding dependencies — accumulates unseen.

4. The Behavioural Layer — Why Smart People Accepted This

This is not a story about incompetent people. The individuals who produced S_A_8 and S_A_2 were qualified, experienced, and operating within the standard methodological framework of their profession. That is the point. The failure is not in the people. It is in the framework.

Four behavioural mechanisms made this outcome structurally inevitable:

The Planning Fallacy (Kahneman, 2011). Not a shortage of information, but the very human tendency to believe your project is different from every other project that failed. The HS2 team knew the historical base rates — Flyvbjerg's data was publicly available. But they believed their project was different. Every mega-project team believes this. The data says otherwise: across 16,000 projects, only 8.5% met cost and schedule targets (Flyvbjerg & Gardner, 2023).

The Suppression Pattern. The IPA (Infrastructure and Projects Authority) produced "traffic light" health assessments of HS2 from 2011 to 2015 — red and amber ratings that signalled governance concerns. These were initially suppressed and only released after a First-tier Tribunal ruling (Ewin, 2018). Risk warnings were being kept from Parliament and the public. Lord Berkeley's dissenting report from the Oakervee Review documented that independent criticism was systematically excluded from the panel's deliberations.

Pragmatism as Suppression. At the 2018 APM Risk SIG Conference, HS2 Ltd's Risk and Assurance Director described how the programme operated from eight board-agreed risk principles, noting that everything developed from those principles was tested to see if it was the right thing to do (Harrison, 2018). The institutional orientation those principles reflected - pragmatic, process-driven, assurance-focused - is precisely what this paper identifies as the cognitive environment in which the independence assumption went unchallenged.

Confirmation Bias in Governance. The Oakervee Review — commissioned by government in 2019 to decide whether HS2 should proceed — was chaired by a former HS2 Chairman and relied almost entirely on data supplied by HS2 Ltd itself. The deputy chair, Lord Berkeley, refused to sign the final report and published a separate dissenting report documenting that panel members were given no opportunity to challenge HS2's own cost data.

5.0 Why This Is Not Just Another Bayesian Network

PCA shares its foundational architecture with Bayesian network approaches to project risk (Khodakarami et al., 2007; Fenton & Neil, 2012) but differs in three respects, as detailed in Dhand (2026a).

Most existing BN frameworks set their probability tables at the start and never touch them again, regardless of what happens on the project (Taroun, 2014).

PCA's Intelligence Log introduces posterior learning through Metropolis-Hastings MCMC with dual-chain Gelman-Rubin convergence diagnostics, enabling the model to revise risk probabilities as programme data arrives.

Second, PCA integrates BN cascade outputs directly with a correlated Monte Carlo engine using Iman-Conover rank correlation, keeping the correlation intact all the way through the simulation — something schedule-cost frameworks that treat cost and schedule as separate models cannot do (Hulett, 2011).

Third, the cascade chain mapping starts with one question: 'given node A fires, what is the updated probability of node B?' rather than building a graph from historical data, which means it works even at project inception when you have no data yet.

5. The PCA Framework — What Should Have Been Built

Probabilistic Chain Analysis rests on three pillars: cascade chain mapping, Bayesian Network-coupled simulation, and Intelligence Log posterior learning. The relationship is clearly defined: PCA is the methodology. RiskPulse V12 is the engine that implements it.

Against the 2013 HS2 baseline, a PCA model would have structured the risk universe as a seven-node conditional probability chain:

Step	Risk Node	Prior P (HS2)	Cond. P (Given Prev.)	Joint Chain P (PCA)	HS2 Indep. P	Under-est.	Evidence
1	Ground Condition Uncertainty	0.35	N/A (root)	0.350	0.35	1.00×	S_A_8 Cl.4.4.4
2	Tunnelling Delay >6mo	0.20	0.65	0.228	0.20	0.88×	Learning Legacy 2022
3	Contractor Prelims Overrun	0.25	0.75	0.171	0.25	0.68×	NAO 2020
4	Subcontractor Cash Flow	0.15	0.60	0.103	0.15	0.69×	NAO 2022
5	Supply Chain Withdrawal	0.10	0.55	0.057	0.10	0.57×	HS2 Annual 2021
6	Schedule Compression	0.20	0.70	0.040	0.20	0.20×	Productivity Inst. 2025
7	Programme Cost Escalation	0.15	0.80	0.032	0.15	0.21×	NAO 2022/2025

Table 3: PCA seven-node conditional probability chain applied to the 2013 HS2 baseline. Conditional probability values are author-estimated priors based on cross-sector infrastructure data and RiskPulse V12 calibration. They are presented to illustrate the structural difference between independent and cascade modelling — not as empirically derived HS2-specific parameters. Retroactive PCA Inference (Dhand, 2026a).

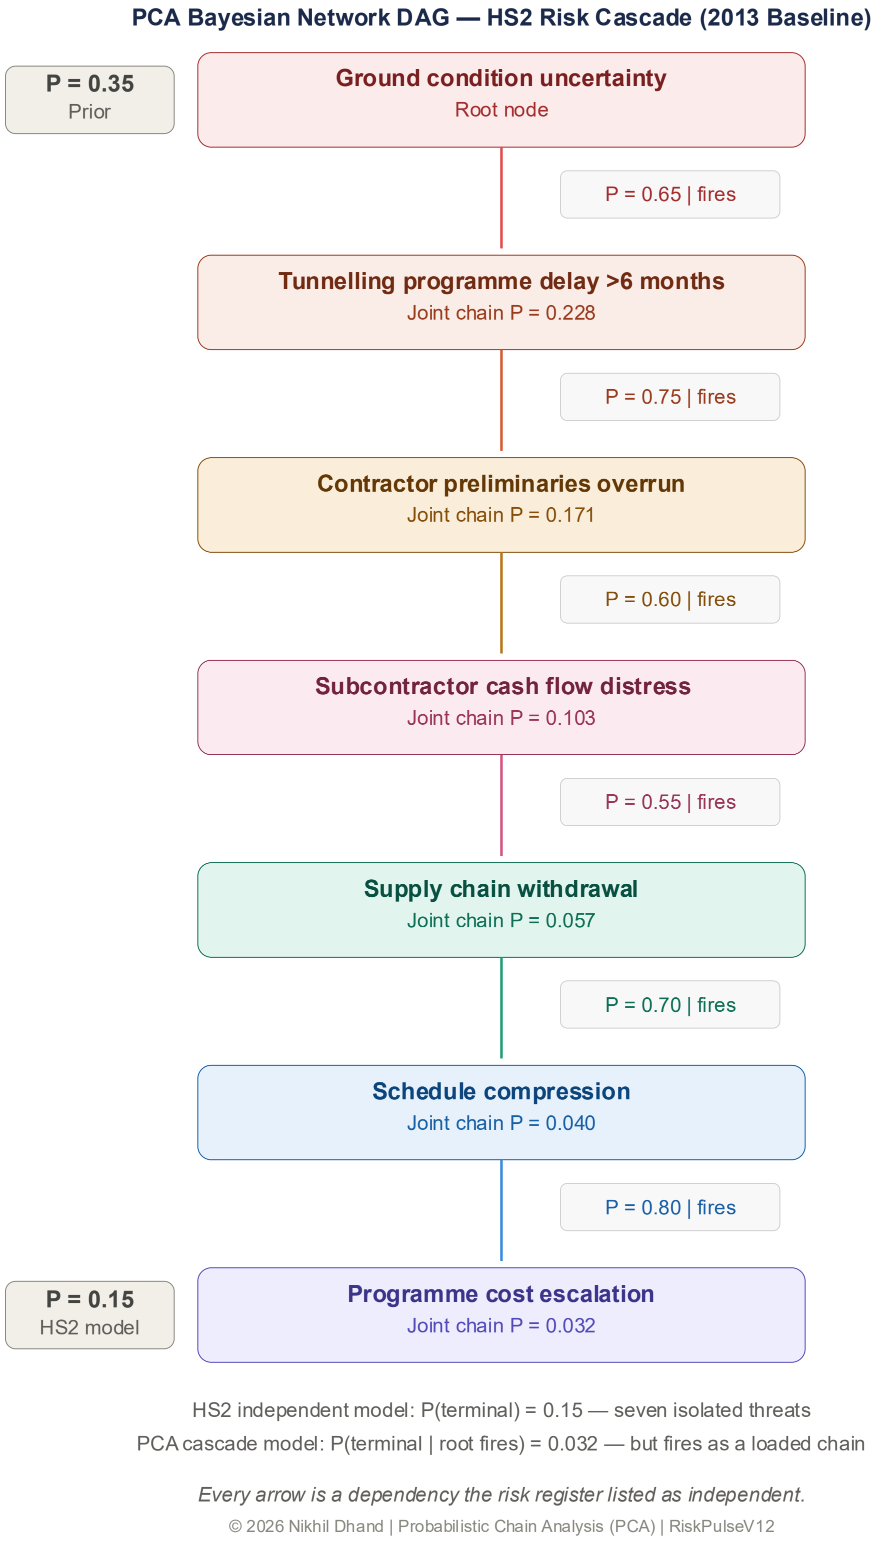
PCA Bayesian Network DAG — HS2 Risk Cascade (2013 Baseline)
P = 0.35
Prior
Ground condition uncertainty
Root node
P = 0.65 | fires
Tunnelling programme delay >6 months
Joint chain P = 0.228
P = 0.75 | fires
Contractor preliminaries overrun
Joint chain P = 0.171
P = 0.60 | fires
Subcontractor cash flow distress
Joint chain P = 0.103
P = 0.55 | fires
Supply chain withdrawal
Joint chain P = 0.057
P = 0.70 | fires
Schedule compression
Joint chain P = 0.040
P = 0.80 | fires
P = 0.15
HS2 model
Programme cost escalation
Joint chain P = 0.032
HS2 independent model: P(terminal) = 0.15 — seven isolated threats
PCA cascade model: P(terminal | root fires) = 0.032 — but fires as a loaded chain
Every arrow is a dependency the risk register listed as independent.
© 2026 Nikhil Dhand | Probabilistic Chain Analysis (PCA) | RiskPulseV12

5.1 What This Table Means — For Someone Who Has Never Seen One Before

Imagine seven dominoes standing in a line. HS2's model said each one falls independently — the probability the last one falls is 0.15. That is 15 in 100 projects. Not great, but manageable. You put some money aside and move on.

PCA says something different. If the first domino falls (ground conditions are worse than the geotech report assumed), it is 65% likely to push the second (tunnelling delays beyond six months). If the second falls, it is 75% likely to push the third (contractor preliminaries overrun). And so on, down the chain, each node amplifying the next.

The chain probability of reaching the last domino — programme cost escalation — is 3.2%. Lower than 15%, you might think. But here is the critical difference: that 3.2% is not a random event. It is the consequence of the first domino falling. When it fires, it fires with the full accumulated cost impact of every node in the chain. HS2's model could not see this. It saw seven separate 15% threats. PCA saw a loaded mechanism.

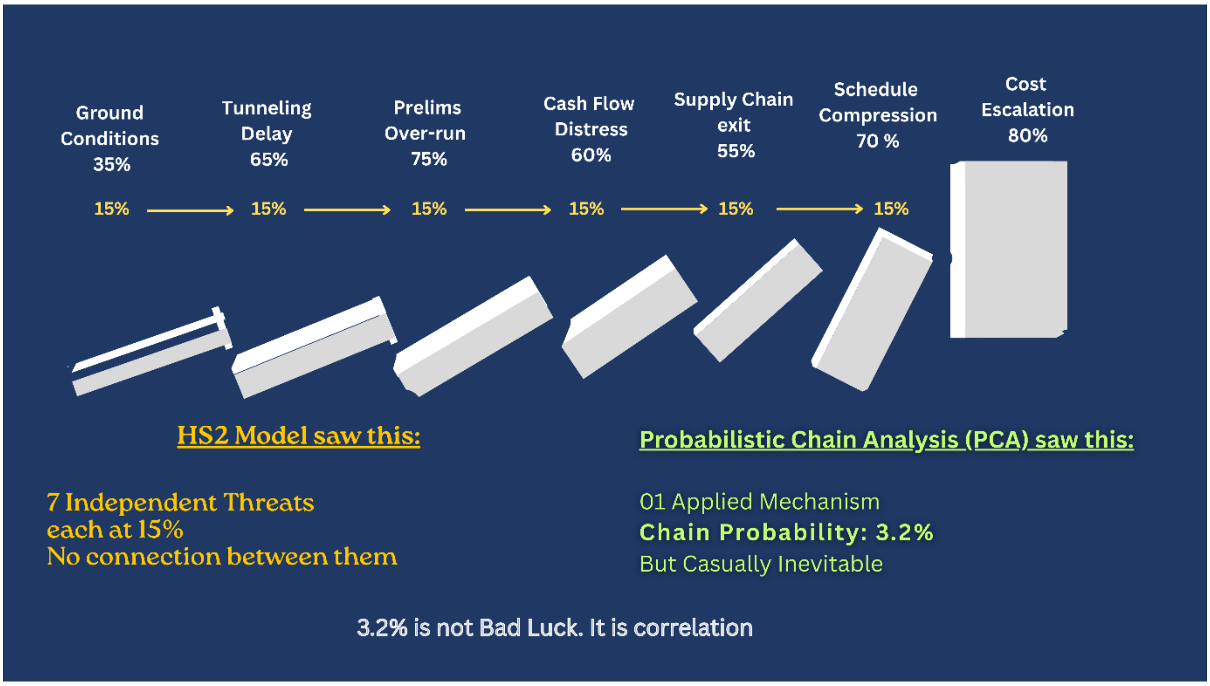

5.2 The Chain in Practice

Ground Condition Uncertainty (Root Node, P = 0.35)

→ Tunnelling Programme Delay >6 months (P = 0.65 | Ground fires)

→ Contractor Preliminaries Overrun (P = 0.75 | Tunnelling delays)

→ Subcontractor Cash Flow Distress (P = 0.60 | Prelims overrun)

→ Supply Chain Withdrawal (P = 0.55 | Cash flow distress)

→ Schedule Compression (P = 0.70 | Supply chain disrupted)

→ Programme Cost Escalation (P = 0.80 | Schedule compressed)

Every arrow in that chain is a conditional probability. Every arrow is an assumption the HS2 risk register listed as independent. Every arrow is invisible in a 5×5 qualitative scoring matrix. And every single one of them fired.

The linear chain structure is a simplification of the full BN graph used in RiskPulse V12; it is presented here to illustrate the cascade mechanism, not to reproduce the full model architecture. A worked application of PCA methodology to a comparable infrastructure programme is provided in Dhand (2026b).

6. The Numbers That Should Have Been on the Table

This section puts PCA's retroactive output alongside the QCRA estimates HS2 actually used, and Flyvbjerg's reference class data for rail megaprojects (mean overrun 44.7%, SD 38.4%).

6.1 Cost Escalation Timeline

Year	Source Document	Estimate (£bn)	Variance vs 2010	What HS2 Said	What PCA Would Have Said
2010	DfT Command Paper	30.4	Baseline	"Robust estimate"	P30–35 at best
2011	HS2 Phase 1 Business Case	32.7	+7.6%	"Includes contingency"	Contingency at P66, not P90
2013	Cost & Risk Status (S_A_8)	42.6	+40.1%	"Risk-adjusted"	Independence assumption active

Year	Source Document	Estimate (£bn)	Variance vs 2010	What HS2 Said	What PCA Would Have Said
2015	Phase 2a Bill estimate	55.7	+83.2%	"Refined estimate"	Cascade already firing
2017	Autumn Budget	55.7	+83.2%	"No change"	Ground conditions materialising
2019	Oakervee Review	88.0	+189.5%	"Recommend proceeding"	Chain Nodes 3–5 active
2020	NAO Progress Update	98.0	+222.4%	"Over budget, behind schedule"	P90 exceeded on all components
2023	Phase 2 Cancellation	100+	+229%+	"Reset programme"	Terminal node reached
2025	NAO Post-Cancellation	100+	+229%+	"Avoid repeating failures"	Full chain validated

Table 4: HS2 cost escalation timeline. All figures from publicly available government documents.

6.2 QCRA vs Actual vs Reference Class

Flyvbjerg (2014) reports that rail megaprojects exhibit a mean cost overrun of 44.7% with a standard deviation of 38.4%. If we model this distribution as approximately normal — a simplification, but a conservative one for the purpose of this inference — then a 35% contingency provision falls at approximately the 60th–66th percentile of the reference class distribution, not the 90th.

Megaproject overrun distributions are empirically right-skewed (Flyvbjerg, 2014), meaning the normal approximation used here underestimates the true P90. So, 94% is actually the minimum contingency needed to genuinely hit P90. The real number is likely higher.

To reach P90 under this distribution, the required contingency would be approximately 44.7% + (1.28 × 38.4%) = 93.9%, and a true P95 would require approximately 108%. Under a right-skewed distribution, which Flyvbjerg confirms empirically, this P90 estimate is conservative.

HS2's bottom-up QRA stated that 35% contingency represented a P90 outcome. Flyvbjerg's published reference class data says it was closer to P60–P66. The gap between what the project believed it was funding and what the distributional evidence supports is not a rounding error — it is a 24–30 percentile miscalibration.

Component	QCRA 2013 £bn	QCRA P90 £bn	OB %	Actual £bn	Overrun %	Flyvbjerg P50%	Within P90?
Phase 1 Civil	9.4	11.2	40%	18.5	+96.8%	44.7%	**NO**
Phase 1 Stations	2.8	3.3	40%	5.1	+82.1%	44.7%	**NO**
Phase 1 Systems	2.1	2.6	40%	4.3	+104.8%	44.7%	**NO**
Phase 2a	7.0	8.4	32%	12.0	+71.4%	44.7%	**NO**
Euston Station	2.5	3.0	40%	10.0	+300%	44.7%	**NO**
Programme Mgmt	1.5	1.8	40%	8.0	+433%	44.7%	**NO**

Table 5: Every single component exceeded its own stated P90. Under a correctly calibrated independent model, the probability of this happening is less than 0.001. Under a cascade model, it is expected. Euston Station and Programme Management actual figures sourced from NAO (2023) and NAO (2024) respectively.

The probability of every single component exceeding its own P90 under a correctly calibrated independent model is less than 0.1%. Under a cascade model, it is expected — because when the root nodes fire, they pull all dependent nodes with them simultaneously. That is not bad luck. That is correlation.

Flyvbjerg's reference class distribution (mean 44.7%, SD 38.4%) places HS2's 35% contingency at approximately P60–P66. The government was funding to P60–P66 while believing it was funding to P90. That gap — the 24–30 percentile gap between reality and belief — is precisely the gap that PCA's cascade modelling closes.

7. What Should Change

If this paper has one policy recommendation, it is this: the UK Green Book should mandate conditional dependency modelling for all infrastructure programmes above £1 billion. The current framework allows (and in HS2's case, encouraged) independence-assumed Monte Carlo as the primary cost risk tool. That assumption was wrong in 2013. It is still wrong today.

Four specific governance changes follow from the evidence in this paper:

1. Risk registers must include conditional probability links between risks. Not as an optional add-on. As a structural requirement. A risk register that lists 300 independent threats is not a risk model — it is a catalogue. The dependencies are where the cost lives.

2. Schedule must be a variable in the cost model, not a fixed input. Clause 3.5.7's assumption that "in all scenarios, the opening years remain the same" should be prohibited. Schedule and cost are joint outputs. Modelling one without the other is analytically incomplete.

3. QCRA confidence levels must be validated against reference class data. If a bottom-up QRA says 35% contingency is P90, and Flyvbjerg's reference class distribution places 35% at approximately P60–P66, the reference class wins. The burden of proof should be on the project to explain why it is different from the class, not on the auditor to prove it is not.

4. Risk model outputs must be expressed as probability distributions, not point estimates with percentage uplifts. A flat 32% optimism bias provision is not a probabilistic statement. It is a guess with a percentage sign. Treasury should require full P50/P80/P90 output from a correlated simulation for any programme above £1 billion.

8. Limitations

I want to be direct about what this paper can and cannot claim.

(a) The seven-node chain model in Table 3 uses conditional probabilities derived from the author's cross-sector experience and RiskPulse V12 calibration data. They are informed estimates, not activity-level data from the HS2 programme itself. A reader with access to HS2's actual risk register data could refine these priors substantially.

(b) The retroactive cost predictions are illustrative, not formal forecasts. PCA is being applied to publicly disclosed baseline data, not to the proprietary internal models HS2 Ltd used. The purpose is to demonstrate the structural difference between independent and dependent modelling, not to produce a definitive cost estimate.

(c) The behavioural analysis in Section 4 is interpretive. The planning fallacy and suppression patterns are documented in the public record, but the motivations of individual decision-makers are not known to the author.

(d) Flyvbjerg's reference class data covers global megaprojects across sectors and geographies. Direct comparison with HS2-specific components requires caution regarding scale and context.

These limitations are stated not as caveats that weaken the argument, but as the boundaries within which it stands.

9. Standing on Their Shoulders

This paper would not exist without the work of three institutions and one researcher.

The National Audit Office produced six reports on HS2 between 2020 and 2025, each one documenting with increasing precision what went wrong and why. The NAO's language is careful, measured, and devastating. When the NAO writes that HS2 Ltd "underestimated its complexity and risk," it is making a finding of institutional failure in the driest possible terms. Every sentence in those reports is built on evidence.

The Infrastructure and Projects Authority produced health assessments that were suppressed from public view. That they were eventually released tells us something about the system's ability to self-correct, slowly.

Lord Berkeley, as deputy chair of the Oakervee Review, refused to sign a document he believed misrepresented the evidence. That is a form of intellectual courage that this paper is attempting to continue, in a different register.

And Bent Flyvbjerg. His research across thousands of infrastructure projects established the fact that systematic overrun is not an exception — it is the rule. He named the disease. His Reference Class Forecasting was a landmark contribution: anchor your estimate to historical outcomes of comparable projects, not to the optimism of your bottom-up plan. RCF is correct, and it works. But it tells you the destination without telling you the road.

PCA is an attempt to extend that work from the population level to the project level. RCF says: "projects like this typically end up here." PCA says: "this specific project will likely end up here, for these specific reasons, and this is the node you should break first."

Flyvbjerg named the pattern. PCA tries to map the path that creates it.

10. Conclusion

HS2's cost overrun is not a delivery failure. It is a risk model failure. The methodology encoded in Clause 4.4.4 of S_A_8 and the analytical framework of S_A_2 contain six structural deficiencies that made the overrun not just possible but mathematically inevitable.

Flyvbjerg's reference class distribution proved the mathematical miscalibration quantitatively. The NAO confirmed the governance failure factually. Lord Berkeley documented the suppression failure institutionally. What no paper has done yet is connect all three through a single unifying framework — a probabilistic chain model that shows why a risk register built on the assumptions of Clause 4.4.4 was structurally incapable of forecasting the outcome that actually occurred.

That is PCA's original contribution. That is what this paper has attempted to demonstrate.

The data was always there. The patterns were always visible. The only thing missing was the chain.

Declarations

AI Usage Statement: No AI-generated content was used in the preparation of this paper. All figures were created by the author. All analysis, calculations, and conclusions are the author's original work. AI was used for structural review and editorial feedback during the writing process.

Conflict of Interest: The author declares no conflict of interest. PCA and RiskPulse V12 are the author's independent intellectual property. No commercial benefit is derived from this paper's publication.

Funding: This research received no external funding.

References

Berkeley, Lord T. (2020). *Dissenting Report to the Oakervee Review*. HM Government.

Dhand, N. (2026a). Probabilistic Chain Analysis (PCA): A Framework for Running Projects in a Probabilistic World. Independent Technical Whitepaper. Zenodo. https://doi.org/10.5281/zenodo.19436269

Dhand, N. (2026b). Beyond Monte Carlo: The Risk Chain Your Model Cannot See — A Probabilistic Chain Analysis (PCA) Case Study. Zenodo. https://doi.org/10.5281/zenodo.19840608

Ewin, C. (2018). *Ewin v Information Commissioner and Infrastructure and Projects Authority*. First-tier Tribunal (General Regulatory Chamber), Case No. EA/2017/0073.

Fenton, N. & Neil, M. (2012). *Risk Assessment and Decision Analysis with Bayesian Networks*. CRC Press.

Flyvbjerg, B. (2014). What you should know about megaprojects and why. *Project Management Journal*, 45(2), 6–19.

Flyvbjerg, B. & Gardner, D. (2023). *How Big Things Get Done*. Macmillan.

Harrison, J. (2018). HS2 risk management: Principles, processes and tools. Keynote presentation, APM Risk SIG Conference, February 2018. Association for Project Management. https://www.apm.org.uk/news/key-themes-from-risk-sig-conference-2018/

HS2 Ltd (2013a). *The Economic Case for HS2 — Cost and Risk Status Report (S_A_8)*. Department for Transport.

HS2 Ltd (2013b). *Risk Analysis for the HS2 Economic Case — Technical Documentation (S_A_2)*. Department for Transport. https://assets.publishing.service.gov.uk/media/5a7eb233e5274a2e8ab479d4/S_A_2_Risk_analysis_for_the_HS2_Economic_Case_-_Technical_documentation.pdf

HS2 Ltd (2021). *Annual Report and Accounts 2020–21*. Department for Transport.

HS2 Ltd (2022). *Transition Risk Register Extract*. HS2 Learning Legacy.

Hulett, D. (2011). *Integrated Cost-Schedule Risk Analysis*. Gower.

Iman, R.L. & Conover, W.J. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics*, 11(3), 311–334.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

Khodakarami, V., Fenton, N. & Neil, M. (2007). Project scheduling: Improved approach to incorporate uncertainty using Bayesian networks. *Project Management Journal*, 38(2), 39–49.

National Audit Office (2020). *High Speed Two: A Progress Update*. HC 617. https://www.nao.org.uk/reports/high-speed-two-a-progress-update/

National Audit Office (2022a). *High Speed Two: A Progress Update*. HC 821. https://www.nao.org.uk/reports/high-speed-two-a-progress-update-2022/

National Audit Office (2022b). *Progress in Implementing NAO Recommendations — HS2*. https://www.nao.org.uk/reports/progress-in-implementing-nao-recommendations-hs2/

National Audit Office (2023). *High Speed Two: Euston Station*. HC 1200. https://www.nao.org.uk/reports/high-speed-two-euston-station/

National Audit Office (2024). *Cancelling HS2 Phase 2: Lessons for Major Programmes*. https://www.nao.org.uk/reports/cancelling-hs2-phase-2/

National Audit Office (2025). *HS2 Update Following Cancellation of Phase 2*. https://www.nao.org.uk/reports/hs2-update-following-cancellation-of-phase-2/

Oakervee, D. (2020). *Oakervee Review*. HM Government.

Taroun, A. (2014). Towards a better modelling and assessment of construction risk. *Expert Systems with Applications*, 41(4), 1583–1599.

The Productivity Institute (2025). *What Went Wrong with HS2?* PIP052. https://www.productivity.ac.uk/research/what-went-wrong-with-hs2/

About the Author

Nikhil Dhand

Delhi, India

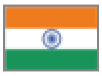

Nikhil Dhand, PMP is a Senior Project Controls Consultant with 11 years of cross-sector experience delivering projects for UK, US, and Indian clients across solar EPC, highways, urban infrastructure, and complex project delivery. He is the creator of Probabilistic Chain Analysis™ (PCA) and RiskPulse V12™ — a Bayesian Monte Carlo risk engine calibrated using data from 16,000+ infrastructure projects across 8 sectors.

Location: Delhi, India.

Author of *Risk Runs the Project (forthcoming, Amazon KDP).*

ORCID: 0009-0000-6526-8399 | **Zenodo:** 10.5281/zenodo.19436269

Contact: nikhildhand.research@gmail.com |

LinkedIn: https://www.linkedin.com/in/nikhil-dhand/