CONFESSIONS OF A BID MANAGER

AKERLOF

February 2026

Confessions of a Bid Manager

For those of you hoping that this may be another instalment in the Timothy Lea collection, I'm afraid to disappoint. Neither Robin Askwith nor any bawdy humour make an appearance here.

Instead, this is a reflection on a few years spent inside the preconstruction department of a Tier 1 contracting organisation. Think the Secret Barrister but less interesting. More Diaries of a Wimpy Kid - if the kid enjoyed studying ITTs, spreadsheets and BIM models.

Mostly, it's an attempt to make sense of how otherwise sensible people (myself included) can make some odd decisions. For long stretches of my career, I felt like a part-time anthropologist: watching how people behaved in pipeline meetings, governance reviews and board meetings.

And during that period, I regularly saw decisions that made little sense on paper and yet were entirely predictable once you understood the context. I was regularly reminded that human behaviour doesn't run on logic any more than a horse runs on petrol[1].

None of this would surprise anyone who has worked in construction for more than a few years. We operate in a sector defined by risk and uncertainty: major projects, delivered on thin margins, creating both predictable surprises and the unexpected in equal measure. For a sector heavily dominated by engineers and an engineering mindset[2], the language of irrationality and bias is rarely used. And almost never in the context of bidding and preconstruction, where our plans, risk allocation and contracts are set. Not in the rooms where a handful of early choices quietly fix the trajectory for whether a project thrives or fails.

This piece tries to correct or at least highlight that blind spot.

On paper, most businesses have defined processes for bidding and preconstruction management – templates, gateways, settlement reviews, etc. But the psychology around a few key decision points often has far more influence on outcomes than any of the formality. The map is procedural but the journey is often psychological.

Behavioural science offers a framework to name the quiet contortions behind decisions that are often already half made. In many sectors, this thinking has gone mainstream - behavioural 'nudge' units in government, specialist consultancies and entire bookshop sections dedicated to the simple idea that people are predictably irrational.

Yet, in construction, we rarely talk about this science. We run unconscious bias training for recruitment and behavioural safety on site but almost nothing on the biases that shape which projects we pursue, how we price them, and how we persuade ourselves they're deliverable. We obsess over design and commercial risk but rarely pause to reflect on decision risk.

So this piece isn't an academic tour through behavioural science. It's something more grounded: an account of how I've seen biases and heuristics play out.

Applied behavioural science with rough edges.

1 It's a line I like but one that should be credited to Rory Sutherland.

2 An Engineering Mindset is something brought to life by Diego Gambetta and Steffen Hertog in their book 'Engineers of Jihad: Curious Connection between Violent Extremeism and Education'. Its not a sweeping statement to suggest that most engineers like clear rules and have a low tolerance for ambiguity.

Prelude

Before I go any further, it's worth stating that I'm wary of the narrative that people are predictably irrational. Far from it.

Herbert Simon, the Nobel Prize-winning economist, coined the term 'satisficing' to describe how decisions actually get made. His view was that decision makers can satisfice either by finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic one.

Most of the time, we are not optimising in some abstract, perfectly understood environment. We are satisficing in a messy one. We look for options that are good enough while juggling competing constraints - time, information, risk, politics. A simple rule that fits the environment usually outperforms a complex model built on uncertain inputs.

Context is everything. Bidding is a world of incomplete information and assumptions, competing priorities, and finite time. In this world, rules of thumb are essential. There are certainly patterns where things go wrong, but it is equally wrong to treat every deviation from textbook rationality as a mistake. Often, it's a perfectly sensible trade where seeking optimal decisions isn't sensible: it's naïve. What often gets dismissed as irrationality is, in fact, a kind of rough, local rationality.

The problem isn't that we satisfice. In a sector defined by uncertainty, we have to. The problem is when the shortcuts we use, no longer match the reality in front of us: when we make the circumstances fit the rule of thumb rather than testing whether the rule still applies. Benchmarks for one project applied to another, without recognition of the abnormals. A "strategic client" designation, once meaningful, becomes a standing excuse for poor commercial decisions.

That seam between necessary satisficing and dangerous self-deception, is where most of what follows sits. Unexamined satisficing: when rules of thumb that arise from habit, hierarchy, or optimism rather than evidence.

What I've written is unapologetically anecdotal. My observations are drawn from meetings, reviews, and tenders rather than lab experiments. But they point to a few recurring truths: that perception often matters more than reality; that the logic of psychology needs to be understood[3] and that small interventions in how we frame, test, and challenge decisions can often outperform grand structural solutions.

To provide some structure, I've organised what follows as a series of biases, loosely ordered against the chronology of a tender, you can click to view, or scroll - entirely up to you.

1. ABILENE PARADOX

2. FRAMING EFFECT

3. CONFIRMATION BIAS

4. UNIQUENESS BIAS

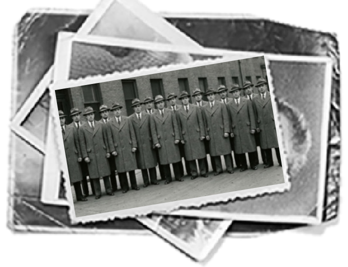

5. AVAILABILITY BIAS

6. AUTHORITY BIAS

7. STRATEGIC FRAMING

8. OPTIMISM BIAS

9. ANCHORING BIAS

10. OVERCONFIDENCE BIAS

11. LAW OF SMALL NUMBERS

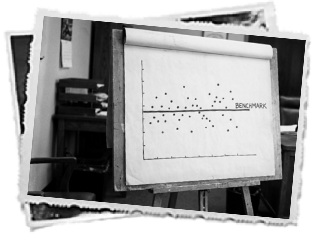

12. FOCUS BIAS

13. BIKESHEDDING

14. HINDSIGHT BIAS

15. OUTCOME BIAS

The list isn't exhaustive (I certainly have more notes about the second stage of a two-stage tender if anyone is interested), nor does it pretend to be a full taxonomy.

These are recurring patterns that, once you've seen them, are hard to unsee. Some are very difficult to offset. But in some cases, your business systems may already exist (at least on paper) to do exactly that.

3 Trying to work without an understanding of the psychology is like playing golf with only a couple of clubs (another Rory Sutherland quip)

ABILENE PARADOX: When everyone agrees but nobody believes

"It looks like a good job."

I've come to regard that phrase as one of the most dangerous in bidding. It sounds reassuring, rational even. But in reality, it carries no objective weight at all, it's usually a convenient shortcut to justify moving forward or an implied motivation to the team.

Selecting opportunities is rarely black and white. Even with a carefully designed bid/no-bid protocol, most assessments deliver the same conclusion: amber. Proceed, but with caveats. And it's in that murky middle ground where decisions are easily swayed by the loudest voice in the room or more significantly, the most senior one.

Personal attachments play their part too. The person who brings the opportunity in. A prior relationship with the client. A faint familiarity with the location ("I know that area well — my tailor's just around the corner"). Suddenly, what began as a cautious, balanced assessment tips into enthusiasm.

This is close to the Abilene Paradox: the story of a Texan family sitting on a porch on a hot afternoon. Someone suggests driving to Abilene for dinner. Nobody really wants to go, but each assumes the others are keen. No one objects. They drive for hours, eat a mediocre meal, and only on the way back does it emerge that none of them actually wanted to go. They all went to keep everyone else happy.

Like this, bid teams can drift into pursuing work they probably shouldn't. There's no moment of active persuasion, just a quiet slide toward consensus. Everyone nods, but privately few are convinced.

The phrase "it's a good job" has no empirical standing, yet it's often bandied around like a trump card. It reassures the room that we're making a sound investment, when in truth, it's a story we tell ourselves to make the decision feel right.

Two countermeasures help.

First, make the reasons not to bid explicit and visible. Write them down. Naming the downsides is itself recognition of risk and forces the room to weigh them rather than wish them away.

Second, adopt a discipline: challenge, then commit. Debate hard in the room; invite dissent explicitly, especially from those closest to delivery. Use pre-mortems (something I come back to later) to encourage and enable the dissent.

Once the decision is made, you stop reopening it. The point isn't to avoid Abilene altogether. It's to make sure that if you do decide to go, it's because you meant to, not because nobody dared to say no.

And certainly not because someone said "it looks like a good job".

Context is Everything

I was once taught that perception of risk is inversely proportionate to how much work you have on the books.

Organisations with a solid and stable pipeline can afford to be choosy. They can walk away from marginal opportunities, hold firm on terms, wait for something better. For those chasing growth targets or worse, watching their pipeline shrink, their evaluation lens is often less selective. Driven by necessity, the need for work can easily drive perception of what is or isn't a good job. It is rarely the project characteristics that determine its attractiveness alone but often instead the context for the bidder and what it looks like relative to other available opportunities.

FRAMING EFFECT: When language does the deciding

"This is a strategic opportunity"

The way we talk about opportunities shapes the decisions we make about them. Language sets the frame, and the frame sets the appetite for risk.

Bids can often be presented internally as "strategic growth opportunities" - this phrase alone can tilt the room. Risks become "manageable constraints". Thin margins become "acceptable for the relationship." Delivery challenges become "innovation opportunities.". Once a bid is framed as 'must win' it stops being assessed and starts being justified. The room shifts from asking should we pursue this? to how do we make this work? which is a different question entirely.

The pipeline effect compounds when growth targets are in play. Senior figures, under pressure to demonstrate momentum, often back "exciting new deals with new clients" over profitable rebids with established relationships, despite 80% of margin typically coming from 20% of clients. The cold bid, pursued because something needs to fill the forecast, is rarely acknowledged as such. It's dressed up as "strategic positioning" or "market entry." But the psychology is simpler: we needed a story to tell the board.

To offset this, separate the framing from the analysis. When an opportunity is described in loaded language - strategic, relationship-critical, must-win, pause and ask: what would we call this if it came from a client we'd never worked with? Strip the history. Assess the facts. If the opportunity doesn't stand on its own merits, the framing is doing too much work.

The words we use aren't neutral. They're decisions dressed as descriptions.

Back to clickable contents page

CONFIRMATION BIAS:
When the evidence fits the story

"The client loves our approach."

One of the quieter traps in bidding is how quickly we stop listening.

It usually begins with a win plan. A structured strategy built around key messages, differentiators, and positioning. In theory, it's a guide. In practice, once it's set, it becomes a filter. We stop testing assumptions and start defending them. Information that fits the story gets amplified. Information that doesn't gets explained away or quietly ignored.

I've seen bids where we became so locked into their own story that they lost sight of what the client actually wanted. Supplementary activities – features, showcase innovations or nice to have additions - gradually overtook the core value drivers that really mattered to the decision. We told the story we wanted to tell, not the one the client was asking to hear.

The same distortion plays out internally. Because we've invested huge effort into crafting the bid, we assume we must be on the right track, forgetting, of course, that every competitor is working hard. Effort feels like evidence, but it's not competitive advantage.

Client encouragement is particularly potent. A throwaway comment - "I really hope you bid this" can anchor a team's entire perception of their chances. What follows is confirmation bias: positive signals reinforces the belief; negative ones are rationalised away. The warm mid-bid becomes proof we're favoured. The formal Q&A feels colder: "They have to be neutral in writing." A competitor's advantage emerges: "But remember what they said to us."

Three things help.

First – ensure the win plan is both clear, quantifiable but short. Three or 4 focal points that will secure the work. Write the executive summary from this early and keep returning to it. Not as a document, but as a discipline. What is the client trying to achieve? What are we actually offering? If the answer has drifted from the question, you've lost the thread.

Second, peer review. The outside view. Someone who hasn't been marinating in the narrative for weeks can see what the room has stopped seeing.

Third, ask the uncomfortable question explicitly: what would have to be true for us to be wrong? If nobody can answer, the story has calcified into faith.

UNIQUENESS BIAS: We're different

"We'll manage it through"

Most projects are unique. The team, the client, the site, the circumstances.

But the ways projects get into trouble are remarkably consistent. And teams routinely assume those patterns won't apply to them.

It shows up most clearly when the risk profile is high. Heavily amended T&Cs. Missing relevant events. Terms that favour the client but leave exposure undefined. These aren't unusual, they appear regularly and regularly cause problems.

And yet teams convince themselves they'll be the exception. We'll manage the relationship. We'll document properly. The risk is theoretical.

This is uniqueness bias: the belief that what went wrong elsewhere won't go wrong here, not because the conditions have changed, but because we're different. Better. More capable.

What gets forgotten is that every team before felt the same way. They also believed their approach was robust. They also thought good practice would see them through. They weren't planning to get caught out.

Uniqueness bias lets us discount evidence we'd rather not face. It turns the outside view, what usually happens, into someone else's problem.

Two things help.

First, start from the pattern. Before assessing why this will be different, ask what usually happens when these terms appear. How often do those risks crystallise? Anchor in the evidence before adjusting for specifics.

Second, treat "we'll manage it" as a claim, not an answer. Every team that accepted aggressive terms believed they could handle them. The question is: what's structurally different this time, not in confidence, but in actual protection?

Back to clickable contents page

AVAILABILITY BIAS: Haunted by the last project

Context is everything and often that context is dominated by the issue of the day.

Bids recently lost on price can make the next pricing review more aggressive, even if the circumstances are entirely different. A live project struggling with a package specific issue can turn that trade into a risk on every tender, regardless of context. A quality issue can become a standing agenda item long after it's relevant.

This is availability bias: the tendency to overweight what comes easily to mind. Recent events, vivid experiences, and unresolved problems loom larger than they should, not because they're more important, but because they're more present.

In an effort to address lessons learned, features can easily receive disproportionate attention on every subsequent bid. Resources shift. Reviews intensify. The logic feels sound: we lost because of X, so we must fix X. But the next bid isn't the last bid. The circumstances may be different. The lesson may not transfer.

The same distortion works in reverse. A recent win inflates confidence in whatever approach delivered it. Teams repeat the strategy, assuming it was the decisive factor, when in truth it may have been incidental. Success is a poor teacher when availability bias is doing the analysis (more of this near the very end).

The problem isn't that recent experience informs decisions, it should. The problem is when it dominates them, crowding out broader patterns, and the specific facts of the opportunity in front of you.

Two things help.

First, ask explicitly: is this concern driven by this bid, or by the last one? If a risk is being flagged because of a recent bruise rather than current evidence, name it. The experience may still be relevant but it shouldn't be automatic.

Second, keep a longer view visible. Track loss reasons over years, not months. Patterns matter more than incidents. A single loss to price is an event. A repeated pattern is a signal. Availability bias collapses the two. Discipline separates them.

The most recent lesson isn't always the most important one. Sometimes it's just the loudest.

Back to clickable contents page

AUTHORITY BIAS:
When seniority trumps insight

Then there's the gold reviewer.

We've just advocated for the outside view. And in theory, senior leaders represent this, bringing fresh perspective. Big-picture thinking. Strategic challenge. However in practice, they often know the least about the client and the detail yet their comments carry disproportionate weight.

I've seen entire sections rewritten because a director engaged late in the process offered a "must-change" opinion based on instinct rather than evidence. Teams defer, assuming status equals insight. It can but often doesn't. In some cases, the gold review introduces more risk than it mitigates: answers become less aligned, less client-specific, and less grounded in delivery reality.

This isn't malicious; it's structural. Senior voices aren't challenged because we conflate hierarchy with expertise. And once their narrative lands, confirmation bias kicks back in - with teams either seeking evidence to support the new view, regardless of its validity or spending effort defending their position,

The fix isn't to exclude senior input, it's to structure it properly. Gold reviewers should see the bid early enough to shape strategy, integrated into pink and green reviews rather than parachuting in at the end. Late-stage authority input is the worst of both worlds: too senior to ignore, too late to integrate.

Challenge should be invited on logic, not accepted on authority. And when a senior comment contradicts the evidence, someone needs to be able to say so, without it being career-limiting.

The question to ask: is this change based on something the reviewer knows, or something they feel? If the answer is the latter, it needs to earn its place like any other opinion.

There's a deeper problem. Flyvbjerg cites research showing that people in power are more susceptible to cognitive biases than those without power. They rely more heavily on intuition, trust ease of retrieval over content, and show higher planning fallacy. The mechanism seems to be that deliberative rationality feels too slow for people accustomed to decisive action. So they shortcut it and in doing so, amplify the very biases that distort decisions.

This means authority bias compounds twice: we defer to seniors assuming they have superior insight, while seniors are themselves operating with amplified cognitive distortion.

The gold reviewer isn't just less informed than the team, they may be more biased.

STRATEGIC FRAMING: Telling the story that gets approved

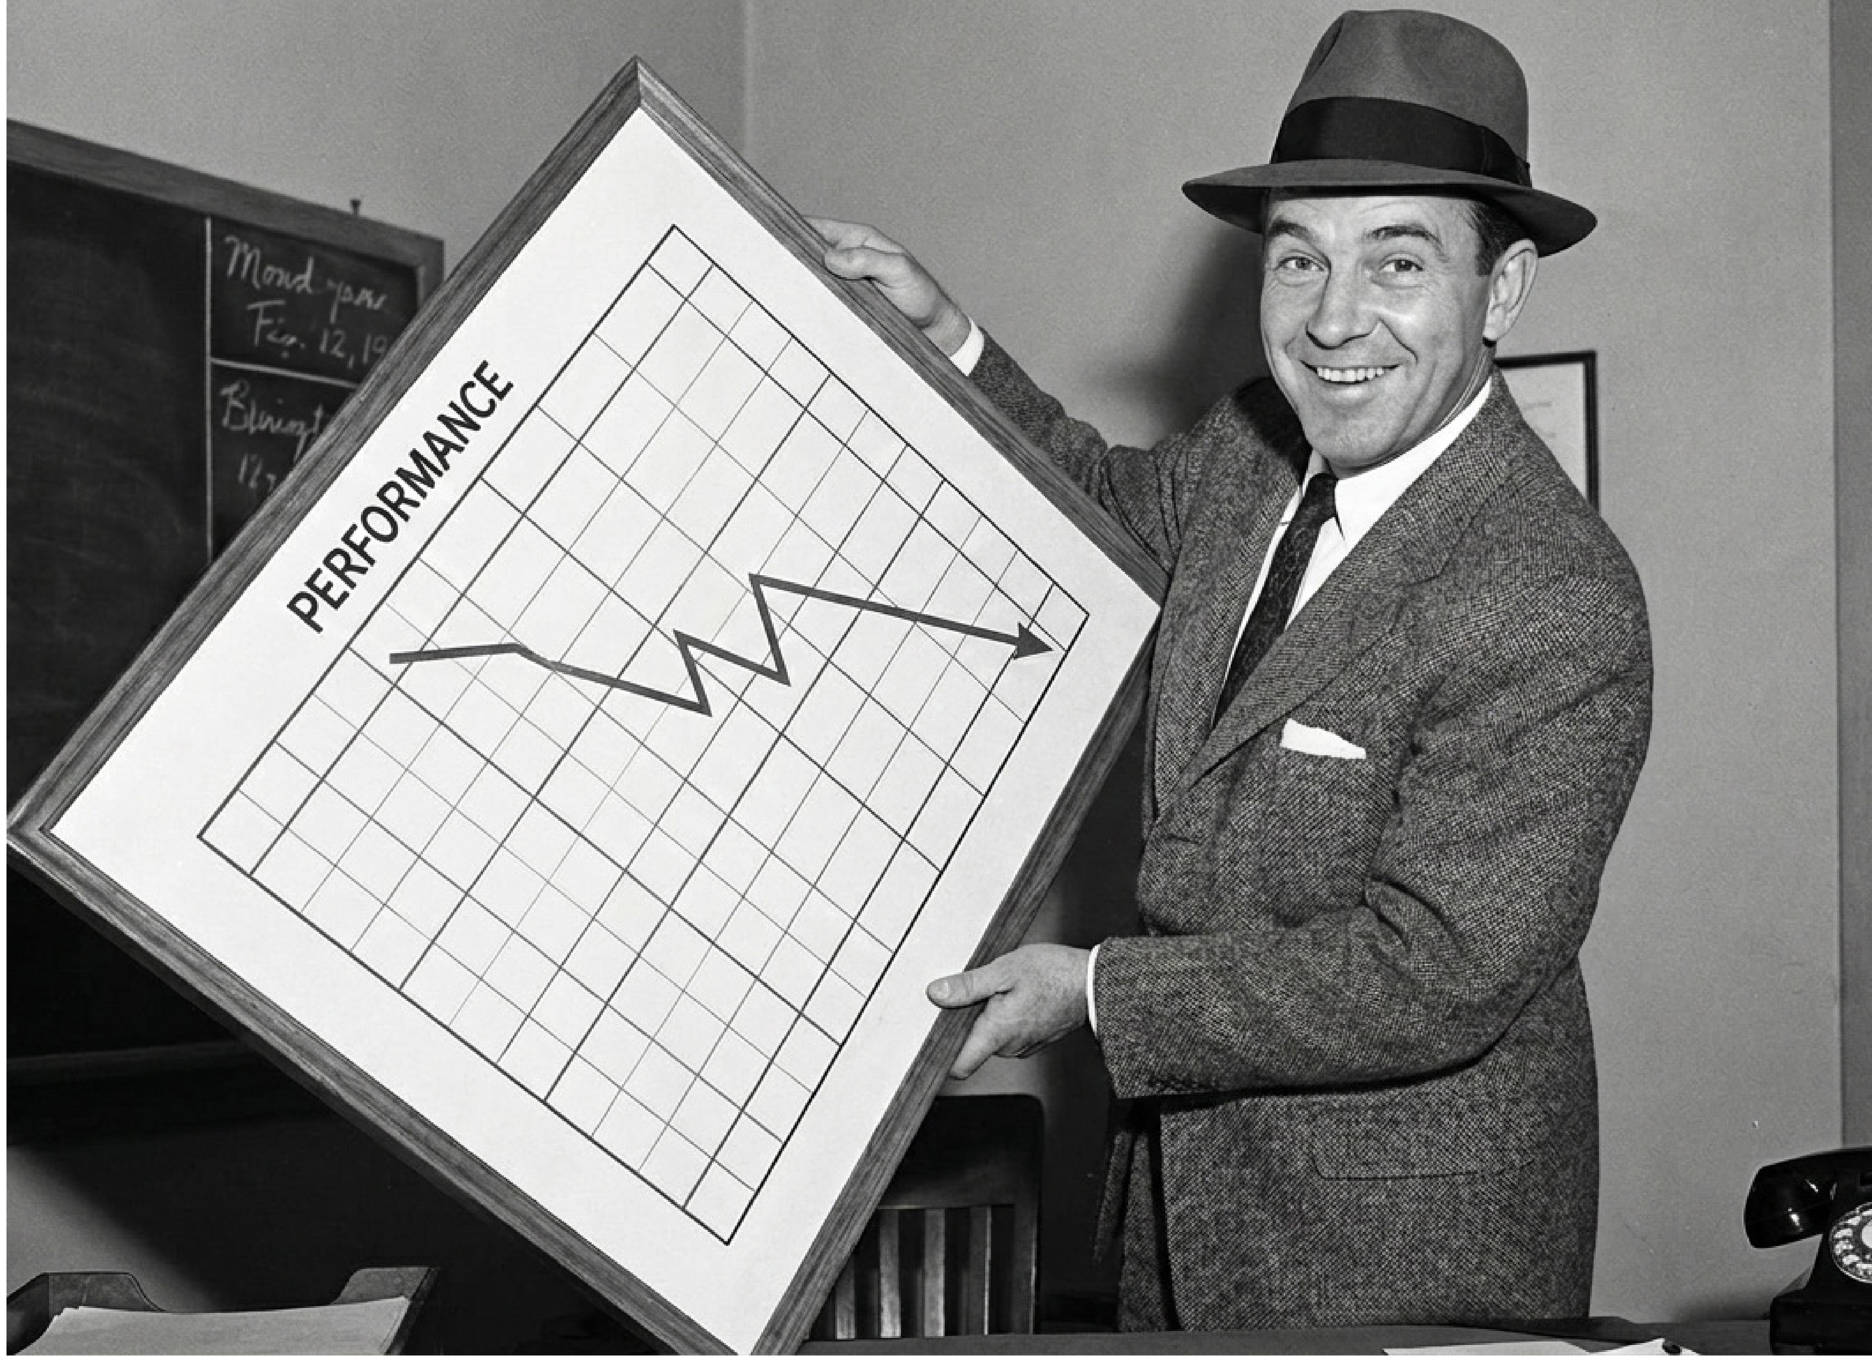

"We just need to get it through risk committee"

In his paper "Top Ten Behavioural Biases in Project Management", Bent Flyvbjerg[4] talks about strategic misrepresentation, the deliberate manipulation of forecasts to get projects approved.

It's a serious charge. And yet I've rarely seen teams set out to deliberately mislead. Most of the time, what happens in bid governance isn't deliberate distortion. It's tactical framing and it emerges from a structural tension that's rarely discussed openly.

Bid teams exist to win work. Risk and governance teams exist to protect the business from bad work. These objectives aren't aligned. Every screening matrix, every approval threshold, every required sign-off represents a potential delay, a blocker, or a no. For bid teams working to tight deadlines, the incentive is clear: frame the opportunity to mitigate the risk of rejection.

So risks aren't ignored but they are framed carefully. And framing shapes perception. Something with an 80% chance of success sounds far more attractive than something with a 20% chance of failure, even though they're the same thing. Governance committees are no more immune to this than anyone else. How a risk is presented can often matter more than what the risk actually is. Consequently commercial leads present mitigations with polished confidence.

The problem however is that tactical framing creates a feedback loop. When risks are shaped to secure approvals, everyone in the chain - governance, leadership, the delivery team - starts believing the story. Sometimes that's harmless. But occasionally, optimism tips into something closer to Flyvbjerg's concern: forecasts are quietly bent to fit internal expectations, targets.

It's rarely malicious. It's structural. Bid teams are rewarded for winning. Organisations reward confidence. So bids are shaped to look confident. But beneath the screening assessments and RAG-rated registers, there's often more fragility than anyone wants to name.

The countermeasure is twofold.

First, the structural answer is to separate approval from assurance. Governance assessing risk should be distinct from the operating unit seeking approval. When the same conversation must achieve both - permission to proceed and honest acknowledgment of exposure - one will dominate. Usually it's the former.

Second, reframe what good risk reporting looks like. Reviews should reward recognition of uncertainty, not performance of confidence. If every bid arrives at governance looking clean and controlled, the process isn't working. It's being worked.

4 https://journals.sagepub.com/doi/10.1177/87569728211049046

OPTIMISM BIAS:
We need to be more commercial

Optimism bias is probably the most familiar distortion in tendering but also an essential characteristic of a successful bid team. It's what gets people through the long hours, the setbacks, the losses.

The goal isn't to suppress it, it's to calibrate it. Faith in the outcome, discipline about the facts. It typically shows up in two places: how we plan and manage the bid, and what we offer in it.

Optimism in planning the bid

Bid teams often underestimate the effort required to hit deadlines (perhaps they need to, they might never put them themselves through the challenge if they were realistic).

But the same patterns repeat: last-minute scrambles, late finishes, submissions that never quite got the review time they needed. There will always be mitigating circumstances - client delays, late information, juggling of diaries / governance reviews. These are all real. But equally they're also always there. Optimism bias remains the constant.

Optimism in what we offer

This is the more damaging form. It's not about how we plan the bid, it's about what we promise in it.

Programme dates can easily be compressed because they need to be competitive. Costs get trimmed to maintain competitiveness. T&C's get dropped to reflect 'market standard'. A competitive scenario can easily become the prisoners dilemma.

There's a psychological driver here: hyperbolic discounting. We value the immediacy of winning over the distant pain of delivery. The award matters now. The consequences are months or years away. So we tell ourselves it'll work out, we'll find efficiencies, the team will step up, the supply chain will deliver. The future feels manageable in a way the present doesn't.

This gets reinforced by culture. When someone questions the realism of an offer, the response is predictable: they're not being commercial. They're not challenging themselves. They lack ambition. The language of competitiveness is used to shut down legitimate concerns. And so aggressive positions get normalised — not because they're justified, but because no one wants to be the person who wasn't "hungry enough."

The result is that we win work we shouldn't have priced that way, then spend years trying to recover margin that was given away before the job even started.

it's compounded by the prisoner's dilemma. If your competitors are sharpening programmes and trimming risk, standing firm on realistic positions can feel like unilateral disarmament. The rational response to irrational pricing is often... more irrational pricing.

And it shows up in how we remember. Successes reinforce the belief that our judgment is sound. Failures get explained away as bad luck or external factors. The feedback loop tightens: confidence grows, but calibration doesn't.

The challenge is that confidence is rewarded. Caution looks like weakness. Raising concerns can feel like disloyalty to the bid. So optimism goes unchecked, not because people are reckless, but because the culture makes it hard to say "we might not win this" or "this programme isn't realistic."

Unexamined optimism isn't conviction. It's exposure dressed as belief.

Two things help.

First, external reference. Not "how long do we think this will take?" but "how long have similar projects actually taken?" Not "what's our chance of winning?" but "what's the base rate for this type of bid, with this level of relationship, at this stage of engagement?" Optimism survives by avoiding comparison. Reference classes force it into the open. They don't resolve the tension but they make it visible. You can still choose to take an aggressive position. You just can't pretend it isn't one.

Second, make it safe to be realistic. If the only acceptable answer is confidence, that's what you'll get — regardless of whether it's warranted. Teams need permission to name risks without it being heard as negativity. The goal isn't to kill ambition. It's to ensure that when we say "we can win this," we actually believe it and have reasons beyond hope. It's to make sure that when someone says "this programme is tight" or "we're not well-positioned," the response is engagement,

It's a fine line between optimism and being commercial. Being able to identify understand the distinction is the key.

Back to clickable contents page

ANCHORING BIAS:
Locked into the first number

Cost plans and deadlines are always anchors.

Once a number lands on the table, it sets the frame for everything that follows. The client hints at a budget, and even if your model says it's higher, the conversation changes instantly: it's no longer "what will it cost?" but "how do we make it work?"

The client's parameters are real. They may be poorly founded, developed from early-stage estimates, outdated benchmarks, or optimistic advice but they exist. They're the constraints you're working to. That's not the problem.

The problem is when teams skip a step.

The disciplined approach is two stages. First, establish what the project should cost and how long it should take based on evidence: scope, specification, market conditions, delivery risk. Develop an informed view independent of the client's expectation. Second, if there's a gap, work out how to close it: descope, re-sequence, value-engineer, adjust phasing, or have an honest conversation about what's achievable.

What actually happens can be different. Teams collapse the two steps into one. They start from the anchor and work backwards, retrofitting logic to justify parameters they've been given. Assumptions get trimmed. Risks get absorbed. The estimate converges on the target, not because the evidence supports it, but because the constraint demands it.

Some call this coherent arbitrariness. The anchor may be arbitrary, but once it's set, organisations build increasingly coherent rationalisations around it. The more effort poured into justifying the number or the date, the harder it becomes to challenge either. Similarly, internal targets - turnover ambitions, benchmarked margins, "win price" assumptions - exert the same pull. The source matters less than the effect: once a number is established, evidence bends toward it.

By the time governance sees the bid, the gap has disappeared. But has it been closed through defined steps or through narrative and trying to make it work?

Three things help.

First, protect the first step. Before any discussion of "how do we fit the budget," require a clear answer to "what does the evidence say it should cost and how long it should take?" Document that figure. Keep it visible. If the final bid is significantly lower, the gap should be explainable through specific scope or risk decisions. Not through generalised optimism.

Second, separate the assessment from the bid strategy. The conversation about the baseline (e.g what the project should cost and how long it should take) is often distinct from what needs to be submitted. Once an evidence-based position is established then a bid team can decide whether to sharpen, where to take risk, and what gap they're consciously accepting. When both conversations blend, competitive pressure risks shaping the baseline. The anchor can easily win before anyone notices it's moved.

Third, build in re-anchoring points. Early numbers harden quickly. Once effort is poured into justifying a figure, challenging it feels like wasted work. Governance should create structured moments- at stage gates, at key milestones - to reset assumptions and ask whether the original anchor still holds. Anchors will always exist. The question is whether they ossify into dogma or get replaced with better ones as information improves.

Without this discipline, tenders become exercises in fitting constraints rather than solving problems. Clients think they're evaluating competitive offers. What they're often seeing is how well teams have contorted their proposals to fit parameters that were questionable from the start.

OVERCONFIDENCE BIAS: We'll make It work

Optimism bias is about believing the future will be kind. Overconfidence is about believing we're good enough to handle it when it isn't.

In bidding, overconfidence doesn't show up as swagger. It's rarely bravado, instead it's something subtler, quieter. It's the unexamined faith in our own tools, methods, and ability to close gaps that shouldn't have been opened.

I've seen teams convince themselves they can deliver to targets that were never realistic. A cost ceiling gets set, a programme date locked in, and suddenly the conversation shifts from "what's achievable?" to "how do we make it fit?" Assumptions become solutions. Ingenuity becomes the answer to everything.

It gets worse when benchmarking enters the room. I've seen carefully built prices, developed over weeks of engagement with supply chain partners, pulled apart because they didn't align with "the norm." Discounts were agreed against packages that had been properly scoped and market-tested, based solely on benchmarks that weren't interrogated.

The irony is that often nobody paused to ask why those benchmarks were lower or whether they were even comparable at all. Were they drawn from projects of a different scale? In different geographies? Under different procurement strategies? Often the seniority of origin was inversely proportional to the extent of interrogation. We just assumed the dataset was definitive, and adjusted our own numbers accordingly.

This is overconfidence at its worst: believing our forecasting, benchmarking, and pricing are more accurate than they are and worse, believing we can always "solve" the gaps by working harder, pushing the programme, or squeezing the supply chain. The result is a kind of collective self-deception, where we reassure ourselves we've cracked the problem without ever truly challenging whether the problem was framed correctly.

Two things help.

First, interrogate the benchmark before applying it. What's the source? What were the conditions? How comparable is it really? A number without context is just a number.

Second, treat "we'll make it work" as a red flag, not a reassurance. If the only answer to a gap is effort and ingenuity, the gap hasn't been closed, it's been deferred.

LAW OF SMALL NUMBERS: Three benchmarks and a prayer

We benchmarked costs and programmes on almost every major bid. At least, we told ourselves we did. In reality, our ability to leverage data effectively wasn't great - intelligence wasn't shared properly, and what we had was patchy at best. More often than not, we were working with fragments: a handful of past projects, often distorted by sector variances and market dynamics that quickly evolve.

This is where we repeatedly fell into the law of small numbers: assuming that a small, selective dataset could reliably predict the future. Broad-brush benchmarking - m² rates, cost-per-bed figures - has its place to encourage challenge and questions. But treating fragments as if they were patterns is dangerous.

Major schemes are rarely alike: site conditions, procurement strategies, planning constraints, client behaviours they all vary, often significantly. Yet we treated benchmarks as transferable and, in doing so, overestimated the accuracy of our forecasting.

Philip Tetlock's work in Superforecasting explains exactly why this happens. Even seasoned experts, he found, consistently overestimate the precision of their predictions. They confuse plausibility with probability. Numbers look robust but beneath the surface, confidence can be easily misplaced.

This is compounded when benchmarking and commercial strategy sit in the same hands, creating a natural gravitational pull toward numbers that support the position. Analysis that "proves" the price is competitive or that we can build it cheaper. None of this requires dishonesty. It's human nature. People validate their own work. Optimism gets baked into the dataset before anyone else sees it, not through bad faith, but through the quiet influence of wanting to win.

And once you believe the numbers are right, every subsequent decision flows downstream from that illusion.

Three things help.

First, use the data you have, properly. Most established contractors have delivered hundreds of projects. The data exists: final accounts, programme records, post-completion reviews. The barrier has never been volume; it's been fragmentation and access. Modern tools now make it possible to aggregate, normalise, and interrogate this information at scale. The excuse that "we don't have enough data" is harder to sustain than it used to be.

Second, separate the benchmarking from the incentive to win. This doesn't mean removing review by estimators, in fact their judgment is essential. But independent review, whether from a central function, a parallel team, or external reference, breaks the feedback loop. The question isn't whether people are trustworthy. It's whether the structure allows optimism to go unchecked.

Third, be honest about what you don't know. Small samples feel like evidence. They aren't. The discipline is knowing the difference. If comparables are scarce, say so. If the dataset is thin, acknowledge the uncertainty at governance. Confidence should be earned by evidence, not assumed by default.

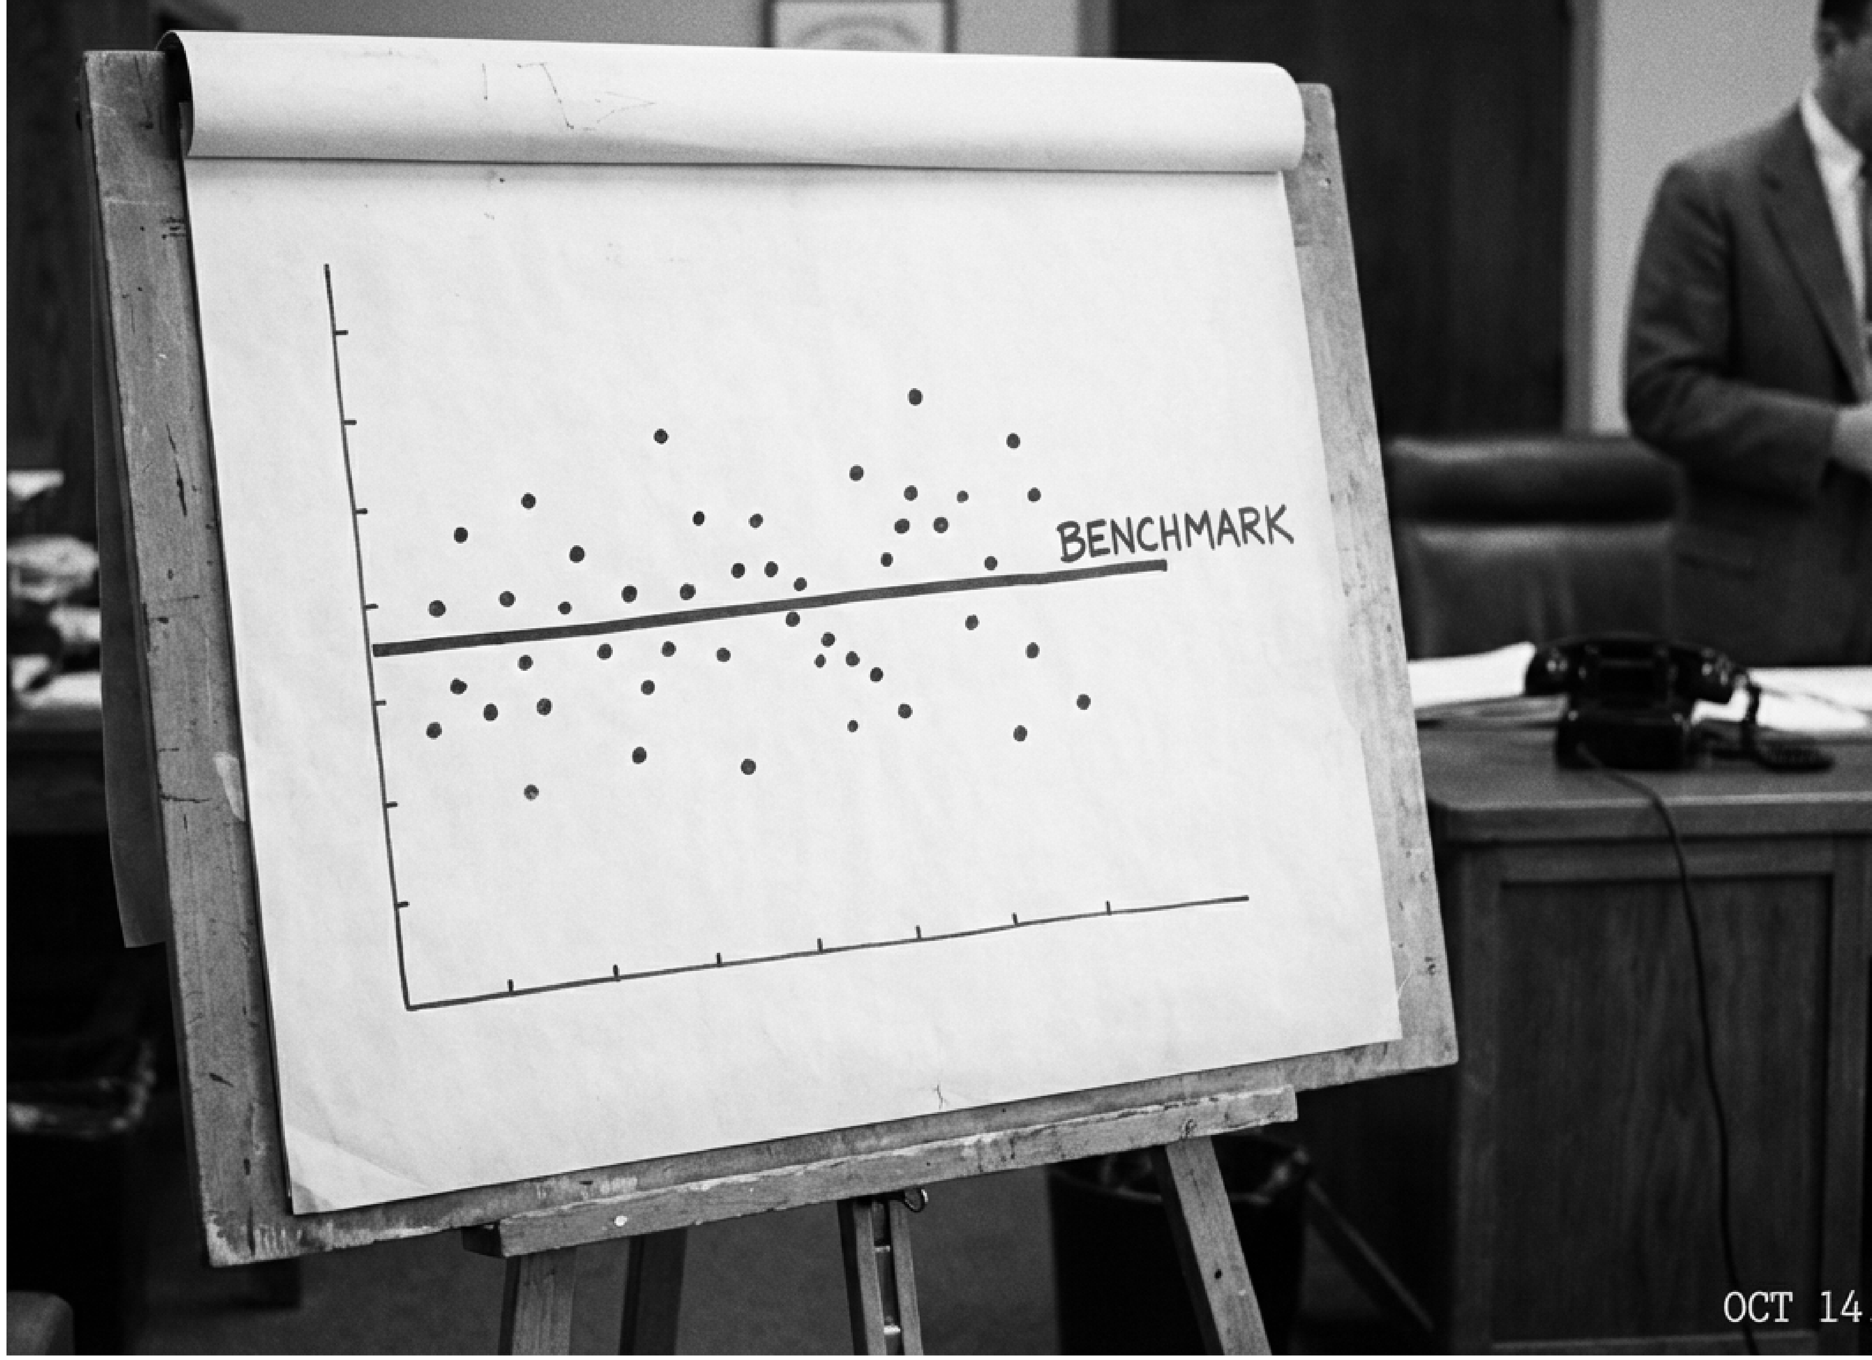

FOCUS BIAS:
The illusion of precision

Margins developed at bid stage are projections made at the point of maximum uncertainty. And yet business processes often demand they be finite and precise.

Once a project moves to site, forecasting becomes more nuanced. Delivery teams work with ranges - best case, likely, worst - and update them as information improves. But at tender stage, the margin is a single number. Absolute. Definitive. Carried forward as if it were fact.

This creates a false confidence that compounds through the organisation. A margin estimate born from incomplete design, provisional pricing, and assumed productivity gets treated as a baseline to defend rather than a hypothesis to test. Variances are measured against it as if the original figure were solid. But it wasn't.

The problem isn't that estimates are made - they have to be. The problem is that the uncertainty behind them disappears the moment the number enters a spreadsheet. Precision implies certainty. And certainty, in a tender, is almost always misplaced.

The countermeasure is to carry the uncertainty forward, not just the number. Present margins as ranges at governance, not points. A single figure is ambiguous: is it the minimum acceptable? The most likely? The target? It's open to interpretation, and different people in the room will interpret it differently.

A range - best case, likely, worst - forces a different conversation. It acknowledges volatility. It makes the potential swing visible. And it prompts sharper questions: what drives the upside, what threatens the floor, and what's the commercial plan for protecting the margin if things move against us? Governance can then be explicit about what it's signing off. The minimum threshold. The expected case. Or all three, with clear conditions attached. That clarity doesn't exist when a single number gets carried forward as if it were fact.

Track how estimates evolve from bid to delivery. If the spread narrows as information improves, the process is working. If it doesn't, the original precision was theatre.

BIKE SHEDDING:
The law of triviality

Settlement reviews can often interesting places for both the right and wrong reasons. But one feature that can often arise is bikeshedding.

Ther term comes from C.Nortcotes Parkinson's anecdote of a planning committee asked to review and approve the application of a nuclear power plant, who instead spent an inordinate amount of time debating the design of the bike shed[5]. The bike shed is something everyone feels qualified to have an opinion on. The reactor is too complex, too technical, too uncomfortable to challenge.

It reflects a phenomenon where people spend disproportionate time of trivial matters while important ones are put aside. In tenders, it often manifests in experienced people drilling into the risks they know while skipping past the ones they don't. The unfamiliar risk doesn't get challenged. It gets nodded through, not because it's been assessed, but because nobody in the room feels equipped to question it.

This is how blind spots become institutionalised. The scrutiny follows expertise, not exposure.

The problem compounds when zero risk bias enters. Governance often seeks certainty where none exists. A risk is flagged, and the response is predictable: can we eliminate it? The review fixates on a specific clause, a particular subcontractor risk, a single assumption, demanding resolution. The result is a strange inversion: governance becomes most rigorous on the things that are familiar and nameable, and most permissive on the things that are unfamiliar or diffuse. If a review is allowed to drift it can easily settle on comfortable ground, the risks people know, the certainties people want.

It can be easily corrected.

Set a clear agenda - identify the three or four risks that actually determine whether this bid succeeds or fails; not the ones that are easiest to discuss, but the ones that carry the most exposure. Put them at the top. Frame them clearly and develop considered solutions in advance. Force the room to engage before energy dissipates into familiar territory.

If a risk is unfamiliar to the room, flag it. If certainty doesn't exist, say so and explain how the uncertainty has been priced or mitigated. Awareness of what you don't understand is a form of risk management.

The risks that damage projects are rarely the ones the room spent an hour debating. They're the ones that got two minutes or none at all.

5 Parkinson's Law: Or The Pursuit of Progress (1957)

HINDSIGHT BIAS: When every loss was obvious

Once a tender result lands, everyone becomes a little wiser.

The outcome, whatever it was, suddenly looks explainable. Wins confirm that the strategy was right. Losses reveal flaws that were apparently visible all along — the price was too high, the client always preferred a competitor. What felt uncertain before the result gets described as inevitable after it.

This is hindsight bias: the tendency to reconstruct the past so that it appears predictable. It distorts learning.

The danger isn't that people ask "why didn't we see this coming?" It's that they skip the question entirely and jump straight to "that was obvious" — describing the past as if they'd known the outcome all along, even though they hadn't.

Losses are particularly prone to this.

After the result, the room fills with people who saw it coming. The warning signs were clear. The flaws were evident. Except nobody said so at the time. The confidence is retrospective. The clarity is invented.

There's a reason this happens. It's called regret aversion: post-rationalisation as psychological protection. Accepting that our judgment is fallible, that effort doesn't guarantee results, is uncomfortable. Explaining the loss as "obvious" avoids that discomfort. The story we tell protects us from the harder truth: that we tried, and it didn't work, and we may not fully understand why.

The problem is that what protects us emotionally blocks us from learning. If every loss was obvious in hindsight, we stop asking what we actually knew at the time. We stop examining whether our early-stage work - capture, engagement, client understanding - was ever strong enough to give us a realistic view of our chances. Instead, we invent explanations that make the present feel more controlled than the past ever was.

A simple correction is to capture predictions before the result. We used to run a quick anonymous survey asking: how do you think it went, and what's your confidence level? When we compared those forecasts to outcomes, the pattern was clear. We were poor at predicting results, far worse than post-result narratives suggested.

It raised an uncomfortable question: if we can't reliably forecast outcomes, have we really done the early work well enough? This is often where the real issue lies – the need for better foresight. Earlier engagement. Stronger capture. A clearer view of where we actually stand before the tender lands.

Hindsight bias lets us feel wise after the fact. The discipline is to capture what we believed before we knew and measure our judgment against that, not against the story we tell afterwards.

Back to clickable contents page

OUTCOME BIAS:
The Loser's Paradox

Hindsight bias rewrites the past. Outcome bias goes further: a win is taken as proof of competence. A loss as evidence of failure. In reality, neither is reliable.

The old cliche used to be that the winner was who got most wrong. That's not necessarily correct, but it's equally wrong to assume the winner understood the job best. They can also be the bidder who was most optimistic about risk, most aggressive on price or programme, or most willing to adopt a change strategy, post selection.

Sometimes the disciplined bidder — the one who priced the risk, challenged the programme, or refused to concede critical terms — loses. And sometimes that loss is the correct outcome. The work wasn't viable on equitable terms. Walking away preserved value, even if it damaged the win rate.

The danger is that teams rarely recognise this distinction. Losses can be post-rationalised as the market taking a flyer. Wins are celebrated as validation. Both can be right and wrong. This is how bad wins become institutionalised.

A bad win is not a project that later goes wrong, that's hindsight talking. A bad win is one where, at the point of commitment, the evidence did not support the promises made. Where competitiveness was achieved by absorbing uncertainty rather than resolving it. Where risk was accepted not because it was understood, but because it was not seen. As a rule of thumb, 80% of bad jobs go wrong in precon, so bad wins are real.

And these wins encourage repetition, quietly resetting what the organisation considers "normal." They create issues 18 or 24 months, or even longer down the line. And because the immediate outcome was positive, the lesson learned is precisely the wrong one.

Good losses rarely receive the respect. But occasionally they should. A good loss is one where the team held discipline in the face of competitive pressure, where the decision not to proceed was the right one.

Over 25 years ago Sir Clive Woodward made this point about sport: the stereotype is to celebrate a win and analyse a loss. He argued it should be the reverse: have a drink to drown your sorrows after a loss, but analyse the win to understand whether it was deserved.

The corrective is easy to say but harder to do: separate decision quality from outcome quality.

After every tender, the question shouldn't be "did we win?" but "given what we knew at the time, would we make the same decision again?" If the answer is yes after a loss, the process worked. If the answer is no after a win, that success is a warning.

The BEST MITIGATION Strategy

Many of these biases share a common root: decisions made too late, with too little information, under too much pressure. The structural countermeasure isn't better governance at bid stage - it's moving upstream. Two things help

First, earlier engagement.

Jeremy Brim at Growth Ignition rightfully champions capture management to move decisions upstream: qualifying opportunities before they become commitments, shaping requirements before they become constraints, building relationships before they become competitive disadvantages. His view (and I tend to agree) is that 60% or more of procurement decisions are effectively made before the tender is even released.

The best way to avoid Abilene, for instance, is to choose your destination before you're already in the car.

Second, pre-mortems not post-mortems.

Most internal reviews ask: how do we make this better? A pre-mortem asks the opposite: assume the bid has already failed - why?

Rather than diagnosis after the result, be proactive. Run a pre-mortem mid-bid. Assume the tender was lost. Or assume it was won but became a problem job - margin gone, relationships strained, delays incurred. Ask: why did that happen? What did we miss? What needs to change in the plan, the price, the risk position to protect the business and / or sharpen the offer?

This reframing changes the conversation. In a normal review, raising concerns can feel like disloyalty — you're undermining confidence, slowing things down. In a pre-mortem, surfacing risks is the point. It's psychologically safer to say "we might fail because..." when everyone in the room is being asked to imagine failure.

The point is to create space for doubt, structured, time-limited, but real. What emerges should inform the bid, not just decorate the governance pack. Pre-mortems don't guarantee success. But they force teams to confront issues while there's still time to act on them.

Back to clickable contents page

FINAL NOTE

There's a bias I haven't mentioned: the curse of knowledge. Once you understand something, you forget what the world looked like before you did. The patterns become obvious. The mistakes look avoidable. You lose sight of how hard it was to see clearly when you were in the middle of it. Equally, you can assume your team sees things you don't whilst you may lose sight of things they have clarity on.

This piece is full of that risk. It's easy to name biases after you've watched them play out for years. Harder to spot them in the moment, when you're new to bid leadership, when the deadline is real and the pressure is on and decisions are needed.

Three small disciplines help.

First, keep a record of what you believed before you knew. It's easy to rewrite history over time. A diary keeps you honest.

Second, pay attention to the questions that get asked at gateways and the metrics that get reported. These reveal the rules of thumb your organisation and seniors actually use: not the ones in the process manual, but the ones that shape decisions. Watch for patterns, both across the business and in specific individuals. What gets challenged? What gets waved through? The answers tell you what the system rewards and what it isn't looking at.

Third, encourage views from your team, especially junior members. Experience brings pattern recognition but it also brings blind spots. People newer to the process can often see things that veterans have stopped noticing. Their questions are sometimes more valuable than their answers.

Not every shortcut is a trap. Some are genuinely useful: the trick is knowing which ones to keep and which ones to consciously offset.

And finally: nothing is as important as you think it is while you're thinking about it. Results are a combination of hard work, considered strategy, and things beyond your control - market conditions, the behaviour of others, timing, luck. In a difficult market, losses can feel like personal failures. They usually aren't. The discipline is knowing the difference and being kind to yourself and your colleagues when the outcome doesn't match the effort.

Get that right, and a life in preconstruction can be a genuinely enjoyable one.

Back to clickable contents page


AKERLOF
www.akerlof.co.uk