THE DECISION GAP Part I

The Confidence Trap

Why Your Best IC Memos Are Hiding the Information That Matters Most

By Scott Sellers

July 2026

I have spent nearly three decades in Special Operations, making consequential decisions under uncertainty. Sometimes with minutes to decide, sometimes with months. But always with the same structural problem at every scale: the teams doing the analysis know far more than the information that reaches the person making the final call. This is true in a tactical operations center. It is true in an investment committee room.

A $100+ million irreversible decision is up for approval. Four months of diligence has led to this day. A respected partner is sponsoring the deal. A partner and firm known for their discipline. The decision memo runs thirty-five pages and is, by any reasonable standard, rigorous.

Everyone in the room nods. The narrative is clean. The numbers work. The recommendation is clear. And yet there is a moment, brief and easy to dismiss, where something does not quite settle. Not a specific objection. More like a frequency the room can almost hear but nobody names. Two years later, in the post-mortem, someone will say: *We **actually knew** that at the time*. And they did. It just never made it into the memo.

Now ask the senior partner sponsoring that deal how much of what his team actually knew (the 200 meetings, the 50 reference calls, the dozens of side conversations, the years of watching the industry) made it into that memo. If he's honest, the answer is something like twenty percent.

Anchor on that number – twenty percent; twenty percent of the knowledge is in the memo, but the other eighty percent is in the team. The partner's track record.

The associates' diligence and judgment. The firm's institutional experience. The IC is not really approving the deal based on a document; it is approving a recommendation backed by people whose judgment has been earned over years.

This is the architecture institutional investors and decision-makers operate under. Document + human judgment + track record = decision. It has been the standard for decades. It is considered sound.

It is not sound. Across studies of deal performance, across every fund size and vintage, a consistent pattern emerges. Even at top-quartile firms with experienced partners and disciplined processes, a meaningful fraction of approved deals fail to meet the expectations held at IC. Not because the teams were inexperienced. Not because the diligence was thin. Because the compensating layer, human judgment carrying the eighty percent the memo cannot, does not close the gap as reliably as the architecture assumes.

This is the confidence trap. It is the single largest source of preventable failure and underperformance in institutional investing. And until now, it has had no name.

The trap is built on a context gap. The team carrying out the analysis knows far more than the artifact the analysis becomes, and the gap between the two is invisible to the decision-maker reading the artifact. What they see is a coherent narrative, a credible team, and a polished recommendation. What they cannot see is everything the artifact failed to carry. The confidence is real but the foundation supporting the decision is more fragile than the room believes.

That is part of what makes it so persistent. You cannot solve a problem you cannot articulate.

The trap has a structure. It has causes that are accelerating. And the era of tolerating it is ending — not because firms have suddenly developed more discipline, but because the tools to close it have finally arrived. The firms that adopt them first will extract an unreasonable advantage from those that do not.

> *The Confidence Trap: How complex reality compresses into polished artifacts, creating perceived confidence that exceeds actual support.*

. . .

What the Gap Actually Is

In any institutional decision of consequence, the team making the recommendation carries a vast amount of knowledge about the situation. A PE take-private. A consulting recommendation to a Fortune 500 board. A strategic pivot at a public company. An M&A advisory brief on a generational acquisition. In all of these, the knowledge exists in four layers.

Layer 1: Explicit Documented Information

The CIM. The financial model. The market research report. The management presentation. The reference calls written up as formal notes.

This is the 10 to 20 percent of total context that ends up in formal artifacts and gets passed up the chain. It is also the layer that IC memos are almost entirely built from.

Layer 2: Explicit Undocumented Information

The side conversations. The impressions from meeting the CFO in person versus on video. The thing the founder said at a dinner that revealed the real reason he wants to sell. The competitive dynamic you learned about from a former employee who spoke off the record. The management team dynamic you observed during a plant tour.

This layer is known to the individual analyst or partner who experienced it. It is almost never transmitted cleanly into the formal artifact. It gets flattened into phrases like “management demonstrates strong alignment” or “competitive position appears defensible.”

Layer 3: Tacit Pattern Recognition

The uneasy feeling that the margin structure looks similar to a deal you watched fail in 2019. The recognition that this CEO’s career arc resembles a pattern that tends to end with founder burnout. The suspicion that the industry’s reported growth rates will not survive an interest rate shift because you have seen it happen twice before.

This layer lives in the heads of senior partners who have earned their pattern recognition through decades of deals. It is extremely valuable. It is almost entirely unarticulated.

Layer 4: Load-Bearing Assumptions

Every deal rests on assumptions that, if wrong, invalidate the thesis. Most of these are never surfaced explicitly. They live in the unstructured space between the analysis. An assumption about customer renewal rates buried in the model. An assumption about regulatory stability baked into the multiple. An assumption about talent retention implicitly driving the value creation plan.

These assumptions are load-bearing. The entire investment case depends on them being correct. They are almost never stated, calibrated for confidence, or stress-tested.

. . .

The IC memo, as currently practiced at even the most sophisticated firms, captures Layer 1 in detail. It summarizes Layer 2 through soft qualitative language. It barely touches Layer 3. And it almost completely fails to surface Layer 4.

This is how the trap is built.

> *The Four Layers: What the team knows versus what the IC actually reads. The most decision-critical knowledge lives in the layers least likely to survive translation.*

When I say an IC memo contains twenty percent of the team's knowledge, I mean this. It contains the documented layer with high fidelity. It contains the other three layers in progressively diminished form.

The IC then makes its decision using this compressed artifact as its primary input.

The partners know they are working with incomplete information. But the compression is so standardized across the industry that it has become invisible. Normalized risk. The compression has become invisible — mistaken for how decisions get made rather than recognized as a known weakness in how decisions get made.

Why This Gap Is Getting Worse, Not Better

Twenty years of investment in technology, data, and analytical sophistication should have closed this gap. The opposite has happened. Three forces are actively widening it.

Force 1: Deal Complexity Has Outpaced Memo Format

A PE deal in 2005 might have had 30 material considerations. A PE deal in 2025 has 150. Regulatory complexity has multiplied. Portfolio construction questions have grown more sophisticated. ESG and sustainability factors now matter. Geopolitical exposures require analysis they didn't ten years ago.

The number of load-bearing assumptions per deal has roughly tripled in two decades.

The IC memo format, meanwhile, has barely changed. It is still a 15 to 40 page document produced in Word or PowerPoint. It is still reviewed in a 90-minute meeting. It is still approved or rejected on the basis of its narrative coherence plus the track record of the sponsoring partner.

The capacity of the format to hold context has not grown. The demand on that capacity has tripled. The compression ratio, the gap between knowledge acquired and knowledge transmitted, has gotten dramatically worse.

Force 2: Deal Teams Have Fragmented

Twenty years ago, a single senior partner led a deal and carried most of the context in their own head, with a small team of analysts supporting.

Today's deals involve larger teams, multiple external advisors, specialized diligence providers, and consultants who each hold a slice of the total context. The sector lead carries one set of patterns. The financial diligence firm carries another. The commercial diligence firm a third. The legal team a fourth. Management consultants a fifth.

When the IC memo is produced, it is typically synthesized by a VP or Principal aggregating inputs from all these sources.

Each handoff loses context. The synthesizer cannot fully represent what each contributor actually knows. They can only represent what each contributor wrote down. The Layer 2, 3, and 4 knowledge held by each specialist evaporates at each translation.

Force 3: Better Writing Tools, Worse Decision Tools

The rise of AI writing assistants has made IC memos smoother, more fluent, more professional-sounding. It has done almost nothing to improve the underlying decision quality.

It is actively degrading it. Research published in Harvard Business Review in March 2026 makes the case sharply. Researchers at Esade, the University of Sydney, and NYU ran more than 15,000 strategic advice simulations across six leading foundation models, spanning seven core strategic tensions. What they found is that even rich, industry-specific context shifted recommendations by only about 11 percent. Simply flipping the order of options in a prompt moved outputs by roughly 19 percent, nearly double the effect of adding detailed context (Romasanta, Thomas & Levina, HBR, March 2026). The models were not synthesizing context to produce bespoke judgment. They were producing what the researchers termed "trendslop," advice that sounds tailored while quietly steering every leader toward the same small cluster of managerial trends.

The same pattern applies when these models are used to draft or polish IC memos. The context-insensitivity the researchers documented does not disappear when the output format changes from a strategy recommendation to a deal memorandum. The models generating the polished prose are not evaluating the specifics of the deal. They are pattern-matching to whatever recommendation sounds most defensible.

And it gets worse. A beautifully written memo is harder to interrogate for gaps than a clunky one. Narrative coherence hides missing information.

An experienced IC member used to be able to tell when an analyst was uncertain by the quality of the prose. Uncertainty showed up as hedging, awkward transitions, and places where the argument was thin.

Today's AI-assisted memos have a uniform confident fluency regardless of whether the underlying analysis is rigorous or thin. The IC's most important quality-detection signal, narrative friction, has been sanded away. This is not an accident of style. OpenAI's own researchers documented in 2025 that models are trained and scored to guess confidently rather than admit uncertainty, and that a grounded answer and a fabricated one read identically on the page (Kalai et al., OpenAI, 2025). The confident fluency is the product working as designed.

What the Trap Costs

The confidence trap does not usually cause catastrophic failure. That is why it persists.

The damage is slower and harder to see: a persistent degradation of decision quality visible only in aggregate, over years, across portfolios. The typical underperforming deal was not flagrantly wrong at IC. The thesis rested on three

or four assumptions nobody explicitly surfaced. Each seemed reasonable. Together, they had maybe a 70 percent probability of holding. The failure mode is in the stacking, the compound probability across multiple unsurfaced assumptions, and it is invisible because the assumptions themselves are invisible.

The IC process that partners consider best-in-class is, on close inspection, selecting against surfacing the very information most predictive of deal outcome. The format rewards narrative confidence, not assumption transparency.

McKinsey has documented that roughly two-thirds of buyout returns between 2010 and 2021 came from multiple expansion and leverage rather than operational improvement (McKinsey Global Private Markets Review). The underlying theses were rarely the primary driver of returns and were never truly tested. In the era Bain has now named ("12 is the new 5"), where 10 to 12 percent EBITDA growth is required to clear hurdles that 5 percent used to clear (Bain Global Private Equity Report, February 2026), that rescue layer is gone.

Institutions that recognize when the character of their environment has shifted compound advantage over those still optimizing for conditions they were built under.

What It Takes to Escape

No one has escaped because escaping requires two capabilities that have never coexisted: extracting tacit knowledge through structured interaction (what the best consultants do in a boardroom, expensive and unscalable), and synthesizing that knowledge into rigorous analytical artifacts at scale (what templates attempt but fail to do, because they accept whatever the preparer puts in regardless of whether the critical layers have been included).

The irony is sharp. The same large language models currently making IC memos *worse*, sanding away narrative friction, producing trendslop, giving thin analysis the same confident sheen as rigorous work, are the first technology capable of making the underlying decision architecture *better*. The capability that enables polished writing is the same capability that enables structured knowledge extraction. The question is which direction you point it. Even the labs concede the point at their own layer: research from Thinking Machines Lab showed in 2025 that identical prompts return different answers under production load, and that determinism is not a model property but an engineering discipline that must be built deliberately (He et al., 2025). Between them, the industry's own researchers have now documented all three failures: judgment that homogenizes, confidence that fabricates, and outputs that do not reproduce.

Escaping requires three shifts. **First, every load-bearing assumption is surfaced and calibrated**, stated explicitly with a confidence level and supporting evidence, not buried in a model. An assumption stated with 70 percent confidence is fundamentally different from an assumption that was never stated at all. **Second, fragility is mapped before the decision, not discovered after**, the specific conditions under which the thesis fails are identified before commitment, creating monitoring capacity that catches underperformance 6 to 18 months earlier. **Third, the reasoning is preserved**: the recommendation, the dissent, the assumptions with their confidence levels, and the conditions that would have changed the call, written once and kept, so the firm can reconcile every thesis against how reality unfolds. A decision process without a record is a process that cannot learn.

This discipline needs a name, because you cannot demand what you cannot name. Call it **decision governance**: structure for how the decision itself gets made, not the evidence beneath it. Your firm already governs its evidence. Diligence has process, models have review, data has provenance. The decision, the step where all of it converges into an irreversible commitment, is the last ungoverned function in the building. Decision governance closes that gap, and it has a simple membership test: assumptions surfaced and calibrated, fragility mapped before commitment, and the reasoning preserved in a permanent record. A tool or a process that does all three is decision governance. Anything less is a better memo.

These are not theoretical capabilities. The technology to deliver them exists now. The firms that understand this first will escape the trap while competitors still debate whether it exists.

. . .

The compression was tolerable in an era when sourcing, multiple expansion, or cheap debt carried returns. That era is over.

Every technological discontinuity in investing has followed the same pattern: spreadsheets in the 1980s, electronic market data in the 1990s, alternative data in the 2010s. A small number of firms adopt early. They extract outsized advantage for three to five years. Then the capability becomes table stakes. The window is never as wide as the laggards assume.

The confidence trap now has a name. So does the discipline that escapes it. Decision governance exists today; I have spent the past year building Pronoa to be exactly this, and I have too much conviction in the category to pretend

otherwise. The firms that adopt it first will capture an advantage that compounds: every governed decision makes the next one better, and the record never stops accruing. The firms that do not will spend the next cycle producing ever more polished memos and wondering why their decisions keep getting harder to defend.

The only remaining question is whether your firm names it before or after the firms competing for the same deals do.

. . .

Scott Sellers is the founder of Pronoa, a decision governance platform for institutional investors, executives, and advisors. Scott is a career U.S. Navy SEAL officer with 30 years of service and holds an MBA from Thunderbird School of Global Management and a Master's degree from the National War College.

This is Part I of The Decision Gap, a series on decision governance: the structural failures in how institutions make their most consequential decisions, and the discipline that fixes them.

Part II asks a question most firms have never asked: what happened to the deals you passed on? The data in your declined opportunities may be more valuable than the data in your portfolio.

Join the launch list at pronoa.io for future installments and first access to the platform.
