GENAI INVESTMENT · DECISIONAL ARCHITECTURE

# The Cheapest Line in Your AI Budget Is the Intelligence

Jensen Huang has spent four years calling intelligence a commodity. Last week he declared it arrived. For a copper concentrator, that is a procurement instruction, not a prophecy.

*A diagnostic paper for boards, COOs, and operations directors funding GenAI initiatives in copper mining and adjacent heavy industry across LATAM*

**AUDIENCE**
Boards · COOs · Ops Directors

**FOCUS**
Decision-space design for GenAI

**REGION**
LATAM Mining · Heavy Industry

**Jaime Rebolledo** · Director Value Proposition Mining · NTT DATA Chile

# Contents


01 Executive Summary 3

02 Four Years of the Same Sentence, and the Week It Came Due 4

03 What “a Decision” Actually Weighs at 150,000 Tonnes a Day 5

04 The Committee Pattern, and the Pilots That Actually Paid 6

05 Five Capabilities, Read as a Build Order 7

06 When the Model Costs Nothing, the Margin Moves to the Design 9

07 The Bug That Never Throws an Error 10

08 What to Ask Before You Fund the Next Pilot 11

09 When Automation Is Really Just Decision Laundering 14

10 Intelligent Debt: the Liability That Accrues Without an Invoice 15

11 The Test You Can Run on One Pilot This Week 16

12 Key Insights 17

13 Sources 18

# 01

## EXECUTIVE SUMMARY

The intelligence is now the cheapest line item in your GenAI[1] budget, and getting cheaper every quarter. What stays scarce, and what no vendor can sell you, is the design of the space where that intelligence decides: which decisions in your operation it owns, at what level of autonomy, under what single objective function, with what override telemetry[2] a human can read. That design is the work NTT DATA does with mining operators, and this paper is its argument: answer those four questions before you procure the intelligence, not after, because every documented AI success in copper concentration answered them first, and the initiatives that stall are the ones that bought the intelligence and skipped the design.

The claim about cheap intelligence is not ours. It belongs to the man who sells the world its compute. Jensen Huang has held the same functional definition of intelligence since at least 2022, told Cambridge students in November 2025 that intelligence was "about to be a commodity," and on June 1, 2026, at GTC Taipei, declared the thesis fulfilled: *"agentic AI has arrived… useful AI has arrived"* — two weeks after telling Dell Technologies World that demand was going "utterly parabolic." When the supplier of intelligence spends four years telling his customers the intelligence is the commodity, the differentiation has already moved to what you point it at.

The empirical backdrop is uncomfortable and worth stating plainly. The MIT NANDA *GenAI Divide* study reported in 2025 that roughly ninety-five percent of enterprise GenAI initiatives produced no measurable return. Closer to home, the NTT DATA and MIT Technology Review study of autonomy in mining found that seventy-two percent of surveyed mining organizations admitted their projects delivered only part of the benefits they expected. Both are directional signals rather than precise censuses, and the argument that follows holds across their error bars.

In my own conversations with operations leadership across the region, what strikes me is never the absence of ambition. It is that the question "which decision does this AI own" so often has no written answer. The pages that follow give that question structure: what a decision actually weighs at concentrator scale, what the funded-pilot pattern looks like from inside a technology committee, how Huang's five capabilities read as a build order for an orchestration layer designed for the global optimum, and the four questions that turn the next AI proposal on your desk into something a CFO can audit.

[1] GenAI — Generative Artificial Intelligence.
[2] Override telemetry — the record that lets a human see what the model did and reverse its action.

**KEY INSIGHT —** Intelligence became the cheapest line in the AI budget. The scarce asset is the design of the decision space: which decisions exist, at what autonomy, under what objective function, with what override telemetry. That design is not for sale, because it is a description of your own operation.

# 02

## FOUR YEARS OF THE SAME SENTENCE, AND THE WEEK IT CAME DUE

Most boards met Huang's thesis as a 2026 headline. It is older than that, and its age is the point.

In a 2022 interview with the technology analyst Ben Thompson, Huang defined intelligence in the most deflationary terms the chief executive of NVIDIA could have chosen: *"Intelligence is the ability to recognize patterns, recognize relationships, reason about it and make a prediction or plan an action. That's what intelligence is. It has nothing to do with general intelligence. Intelligence is just solving problems."* Five capabilities, no mystique. In November 2025, at the Cambridge Union, he sharpened the economic half: intelligence is "about to be a commodity." In March 2026, on the Lex Fridman podcast, he went as far as saying AGI had, by functional definitions, already been achieved, while drawing the line that matters for this paper: commodity intelligence is not the same thing as humanity, and was never going to be.

Then, on June 1, 2026, at GTC Taipei, he closed the arc: *"agentic AI has arrived... useful AI has arrived."* Tokens, he told the audience, "are now profitable units of revenues." Two weeks before, at Dell Technologies World, he had described demand as going "utterly parabolic." His company put its next-generation Vera Rubin platform into full production with a roughly tenfold reduction in inference cost per token against the prior generation. Read those statements together and the conclusion is not philosophical. It is pricing. The five capabilities scale with compute and data, both of which his company sells in volume, and the unit cost of exercising them is falling on a published curve.

Boards heard the first clause of the 2022 definition and concluded that machine intelligence has matured. They funded pilots, hired a head of AI, and asked vendors how good the model is. The load-bearing clause is the second one. If the intelligence is a commodity, the scarce asset moved. It moved to the design of the space the intelligence operates in. Which decisions in the operation has the intelligence been pointed at. At what level of autonomy. Under what objective function. With what override telemetry. Those four design choices are not the vendor's to make. They are the board's, and a board can no more subcontract

the design of its own decision space to an AI vendor than it could subcontract its operating model to the ERP[3] vendor in 1998.

**KEY INSIGHT —** Huang has held the same definition since 2022 and declared it fulfilled in June 2026. When the supplier of intelligence tells you for four years that the intelligence is the commodity, the differentiation has already moved to what you point it at.

# 03

## WHAT "A DECISION" ACTUALLY WEIGHS AT 150,000 TONNES A DAY

The AI literature a board reads in 2026 is saturated with verbs that describe what the technology does: it orchestrates, it augments, it automates, it copilots. None of those verbs name the thing the technology touches. The thing it touches is a decision. Before the rest of this argument can land, fix one decision in mind at the scale of a world-class concentrator.

A large copper concentrator processes on the order of one hundred and fifty thousand tonnes of ore per day. At the rougher flotation head, an operator resets the reagent dosing setpoint roughly once an hour during stable operation, and far more often when feed grade shifts. The decision sounds clerical. It is not.

The setpoint depends on the upstream ore blend, on the pyrite content fixed by a mine plan executed twelve hours earlier, on a recovery curve the on-stream analyzer reports with a lag, on a reagent lot that arrived three days ago and has not been validated against the historical baseline, and on the operator's private read of which pH probes in the cell row are currently telling the truth. One setpoint, resolved against three competing pressures at once:

1. Recovery against reagent cost.
2. Recovery against the tailings load sent downstream.
3. The instrument reading against the operator's trust in that instrument.

A senior metallurgist resolves it in about thirty seconds. The resolution rests on two decades of pattern memory, two decades of cross-checked cause and effect, and a value system that ranks the recovery curve above any single shift's reagent budget. The thirty seconds look like instinct. They are a compressed architecture.

A major concentrator runs roughly twenty-five decisions of comparable density in parallel at any moment. Each routes through a human or a small team today. The gap between that operation and any defensible level of autonomy is not whether the company has an AI. The gap is whether each of those twenty-five decisions has been catalogued, mapped against

[3] ERP — Enterprise Resource Planning, the corporate system of record for finance, procurement, and operations.

the five capabilities Huang named, declared at a level of autonomy the organization can defend in front of a regulator, attached to an objective function the board has endorsed, and wired with override telemetry the next plant manager can read. That is the work. The model is the easy part.

**KEY INSIGHT —** A model's quality is meaningless until you name the decision it serves. The decision comes first; the model comes second.

# 04

## THE COMMITTEE PATTERN, AND THE PILOTS THAT ACTUALLY PAID

There is a scene that repeats, with local variations, across the technology committees of the region's producers. I have sat through versions of it. A portfolio of AI pilots comes up for renewal after eighteen or twenty-four months. Each pilot hit the target it was given: the maintenance model lifted equipment availability, the dispatch model pushed fleet utilization to a record, the energy model cut consumption against the tariff. And then someone, usually from finance, asks the only question that survives contact with a balance sheet: which of these moved the margin. The room goes quiet, because availability bought with deferred overhauls, utilization bought with longer hauls, and energy saved by throttling the mills are three ways of moving cost from one ledger to another inside the same operation. Each pilot optimized the thing it was told to optimize. Nobody told them to optimize the same thing.

That scene is not an anecdote dressed as data; the data corroborates it. In the NTT DATA and MIT Technology Review study *Autonomy in the Mining Business* (2025), seventy-two percent of surveyed mining organizations admitted their autonomy and digitalization projects delivered only part of the benefits they expected. Only twenty-eight percent reported full delivery. The pattern is partial fulfilment: activity everywhere, reconciliation nowhere.

Now look at the initiatives that cleared the bar, because they are public and they share one design choice. At Freeport-McMoRan's Bagdad operation in Arizona, a machine-learning model was pointed at the concentrator itself: it classified mill feed into ore types from sensor data and recommended operating setpoints. The published results: throughput up roughly ten percent in a quarter, and metallurgical recovery up about one percentage point, an outcome the company described as adding thousands of tonnes of copper with very little capital investment, at a time when the alternative on the table was a capital expansion of the plant. Closer to home, Antofagasta Minerals reports in its 2025 Annual Report that its own operational-recommendation system, deployed across its plants, lifted recovery at Centinela by three quarters of a percentage point and at Antucoya by a similar margin, alongside measured savings in acid and water. Different companies, different toolchains, same design

signature: the AI was scoped against the metallurgical variable that reconciles directly to margin, not against a department's local KPI[4].

Why recovery is the variable worth fighting for is a matter of public arithmetic. A world-scale operation like Escondida produces over 1.2 million tonnes of copper a year at sulphide-flotation recoveries that typically sit in the mid-eighties. Work the numbers and one additional percentage point of recovery at that scale is on the order of ten to fifteen thousand tonnes of copper a year; at prevailing prices, roughly one hundred million dollars annually, every year it holds. That is the order of magnitude sitting behind a design choice, and it is why the question "which variable is this pilot scoped against" is worth more than any question about the model.

The committee pattern and the public successes are two halves of one observation. The pilots that stall are not failures of intelligence; the models mostly work. They are failures of scope: four intelligences bought, zero decision spaces designed. The ones that paid answered, sometimes almost by accident, the four questions before writing the scope: which decision, at what autonomy, under what objective function, with what telemetry. The rest of this paper is about removing the "by accident."

**KEY INSIGHT —** The pilots that pay are scoped against the variable that reconciles to margin. At world-concentrator scale, one percentage point of metallurgical recovery is worth on the order of US$100 million a year. The model that finds it costs a fraction of that; the design that lets it be found costs discipline.

# 05

## FIVE CAPABILITIES, READ AS A BUILD ORDER

Return to Huang's 2022 sentence and read it the way a systems architect reads a specification rather than the way an audience reads a keynote. He named five capabilities in sequence: recognize patterns, recognize relationships, reason, predict, plan an action. Read as a build order, those five capabilities are five layers of an autonomous operating architecture. The mapping is not a metaphor. It is one to one, and the correspondence is the load-bearing observation of this paper.

In the autonomous operating model NTT DATA designs for world-scale copper concentrators, the component that coordinates the whole chain is an orchestration layer designed for the global optimum: a layer that sits above the domain systems of mine, plant, logistics, and shared services, and resolves their conflicts in favor of the operation as a whole. That layer decomposes exactly along Huang's five verbs:

[4] KPI — Key Performance Indicator.

- **Recognize patterns** becomes detection across the live DCS[5] and historian data. The sensing layer reads what the plant is doing now.
- **Recognize relationships** becomes a causal model with explicit anti-hallucination validation: before any relationship is acted on, the layer confirms the relationship is causal and not a coincidence in the data.
- **Reason** becomes the resolution of the roughly ten inter-domain trade-offs the orchestration layer carries, the same class of trade-off the single setpoint in section 03 exposed, now resolved across mine, concentrator, transport, and shared services at once.
- **Predict** becomes a scenario engine that reports outcomes as confidence bands, the tenth, fiftieth, and ninetieth percentile (P10/P50/P90) of contribution margin, not a single deceptive point estimate.
- **Plan an action** becomes a set of operating scenarios projected across readiness, ramp-up, and sustaining horizons.

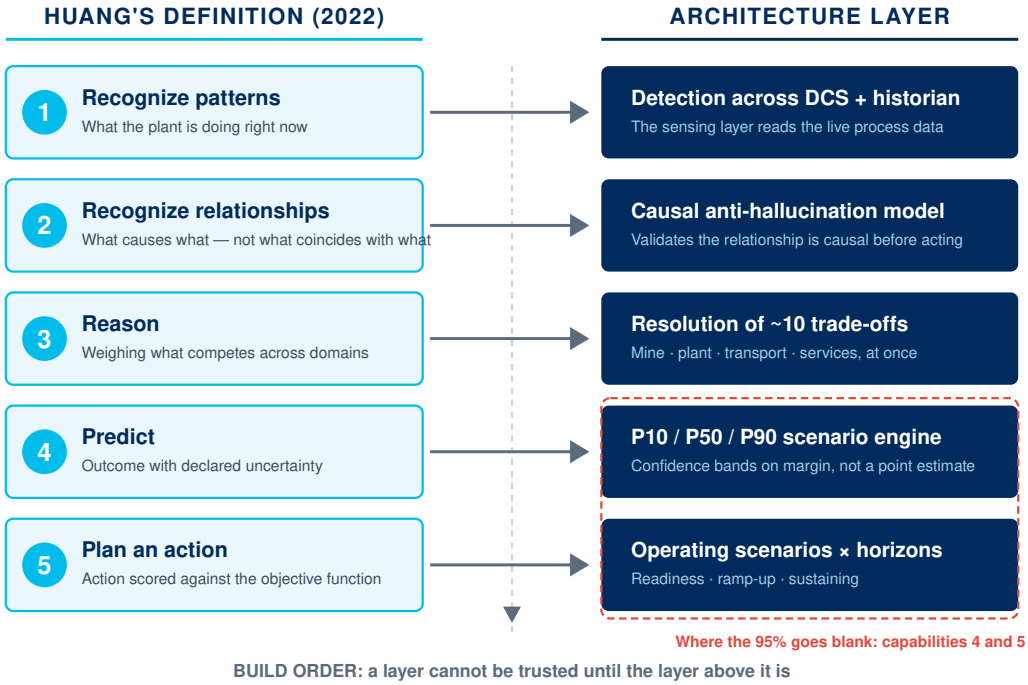


**Figure 1** Huang's five capabilities of intelligence mapped one-to-one to five layers of a global-optimum orchestration architecture in a copper concentrator

The reason the mapping matters to a board is not its elegance. It matters because it converts a vague question (is our AI good enough) into a verifiable one (which of these five layers exists, for which decision). A program that has built pattern recognition and stopped there has built one fifth of an intelligence and called it done. That is what most of the ninety-five percent look like from the inside: excellent at the first capability, blank at the fourth and fifth.

[5]DCS — Distributed Control System, the real-time process-control layer of a plant.

# 06

## WHEN THE MODEL COSTS NOTHING, THE MARGIN MOVES TO THE DESIGN

Take Huang's commodity claim seriously for a moment, then test it against its own limits.

If the five capabilities scale with compute and data, and compute and data are purchasable at falling prices, the intelligence converges toward a commodity input. Vera Rubin's published tenfold reduction in inference cost per token is that convergence expressed as a price list. The same model weights that read your concentrator can read your competitor's. The differentiation cannot live in the weights, because the weights are for sale to both of you. It lives in what you point the intelligence at, and how you constrain it. That is the design of the decision space, and it is not for sale, because it is a function of your operation's specific decisions, your specific risk tolerance, and your specific objective function.

Honesty about the evidence sharpens the argument rather than softening it. Huang's statements are a CEO's position articulated across keynotes and interviews, not peer-reviewed research; he is the most interested possible party, and his thesis should be weighed as such. The MIT NANDA ninety-five percent is a directional signal from interview data, not a precise measurement of the global enterprise population. The seventy-two percent partial-fulfilment figure from the NTT DATA and MIT Technology Review mining study corroborates the direction inside a vertical cohort without nailing the magnitude. And the strongest single counterweight comes from an operator's own lab: Boliden's engineers published that their reinforcement-learning controller for the Aitik mill outperformed the existing control strategy in simulation, and then said plainly that more study was needed before it could approach a viable production deployment. None of these is a precise instrument. All of them point the same way, and the architectural response is invariant across their error bars: design the decision space first, procure the intelligence into it second.

There is a deeper distinction the board should hold. The ninety-five percent is two different problems wearing one number. One part is a comprehension gap: the organization never wrote down which decision the AI owns, so the question of return was never answerable. That part is corrigible. It dissolves the moment the four questions are answered, because the answers are information the organization can produce on its own. The other part is structural: each domain has its own incentive to optimize its own KPI, and no amount of better dashboards reconciles a maintenance team graded on availability with a metallurgy team graded on recovery. That part does not dissolve with information. It requires an act of governance, the imposition of a single objective function above the silos, which is the subject of the next section. Confusing the two leads to the standard mistake: buying more intelligence to fix a problem that better intelligence cannot touch.

**KEY INSIGHT —** The honest caveat is also the frontier: no operator has yet published a validated, full cross-domain orchestration running on margin in copper. The leaders publish single-plant wins. The global optimum is still being engineered, and that is precisely why the design discipline, not the model, is where the race is.

# 07

## THE BUG THAT NEVER THROWS AN ERROR

The most expensive defect in the ninety-five percent has no stack trace, no alert, and no error message. It is the absent objective function.

An organization installs intelligence across its domains without declaring, in a single equation the whole operation is graded against, what the operation is trying to maximize. In that vacuum, every agent does the rational thing: it optimizes the local objective it was handed. The maintenance agent maximizes availability. The dispatch agent maximizes utilization. The energy agent minimizes the tariff. Each is locally correct. The sum is sub-optimal, because the local optima conflict, and there is no arbiter to resolve the conflict in favor of the whole. The committee pattern in section 04 is a forensic photograph of exactly this state.

The orchestration layer resolves it with one equation that sits above every domain and serves as the single north star:

```
contribution margin = tonnage × grade × recovery × price −
                      operating cost
```

Every agent's local recommendation is scored against its effect on that number before it is allowed to act. The maintenance recommendation that lifts availability but defers an overhaul into a high-grade campaign is rejected, because the equation can see the recovery it would cost. This is not a more sophisticated KPI. It is the precondition that makes any KPI meaningful, because it gives the orchestration layer a way to rank conflicting recommendations from domains that would otherwise each declare victory.

The board's role here is specific and non-delegable. The objective function is a statement of what the company values, expressed in a unit that reconciles across domains. A vendor cannot write it, because a vendor does not own the trade-off between this quarter's cost and next year's recovery. The CFO and the COO own it. And in our experience, the hardest part is not the algebra; it is the governance moment where two executives who have never shared a scoreboard agree, in writing, on who owns the recovery term. If that equation does not exist on paper, the AI program does not have an aiming problem. It has a target problem.

**KEY INSIGHT —** An objective function is not a KPI. It is the precondition that makes KPIs meaningful. And no model, at any price, solves a target problem.

# 08

## WHAT TO ASK BEFORE YOU FUND THE NEXT PILOT

Once the objective function exists and the decisions are catalogued, the conversation with a vendor changes shape. The legacy question (how good is your model) becomes unanswerable in a useful way, because a model's quality is meaningless until you specify the decision it serves. Four questions replace it, and each is operational, each translates into a line in a scope of work, and each is defensible in front of a CFO who does not speak the AI vocabulary.

For any decision the intelligence is proposed to touch, the board asks: Which specific decision in the chain does this own. At what level of autonomy does it own it. Against what objective function is its recommendation scored. And what override telemetry lets a human see and reverse its action. The level of autonomy is not a slider the vendor sets; it is governed by the confidence the causal layer can demonstrate, on the ladder Parasuraman, Sheridan, and Wickens formalized in 2000: decision support, decision augmentation, and decision automation, with the automation level reserved for relationships the causal layer can validate at high confidence, and the override-telemetry requirement rising with every step up. A decision running at full automation with thin telemetry is the single most dangerous configuration a board can fund, because it is autonomous and unaccountable at the same time.

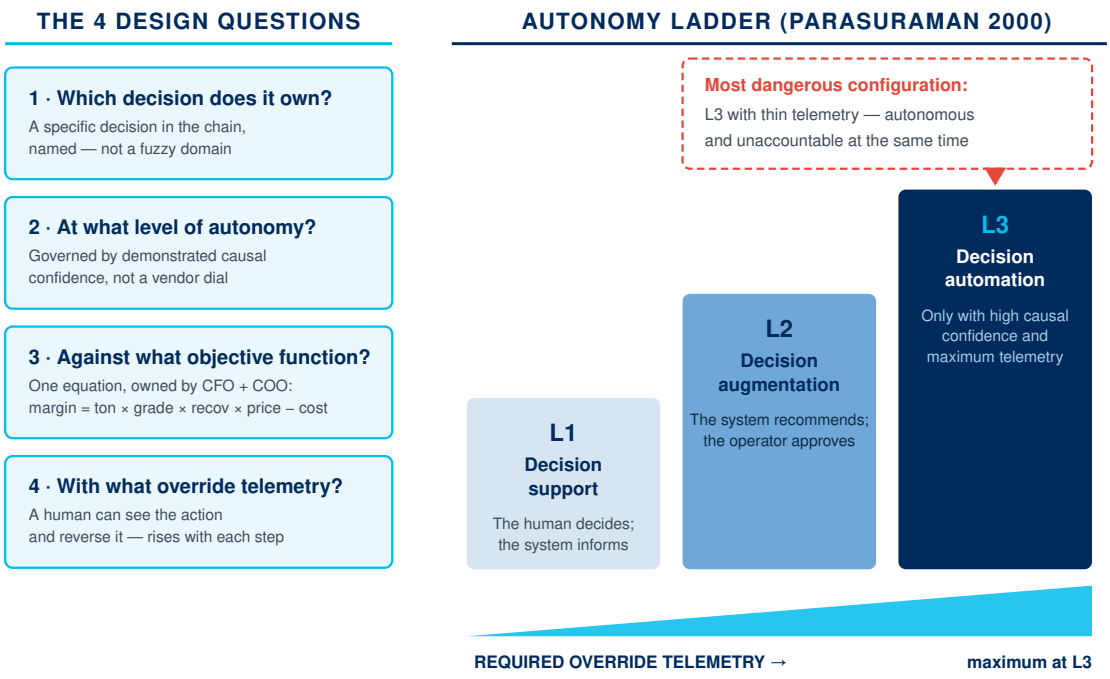


**Figure 2** The four design questions cross-tabulated against the three-level autonomy ladder, showing override-telemetry intensity rising with autonomy

**DIAGNOSTIC QUESTION** *For the AI proposal currently on your desk: can the sponsor name the single decision it owns, the autonomy level, the objective function it is scored against, and the override telemetry, in four sentences, without slides? If not, you are funding an intelligence without a decision space.*

**Table 1** Decision space catalogue · four representative examples

| Decision name | Cadence | Current decider | Trade-offs | Margin sensitivity | Information substrate | Autonomy level |
|---|---|---|---|---|---|---|
| Rougher reagent dosing setpoint | Hourly / shift | Process operator | Recovery vs. reagent cost; recovery vs. tailings load; instrument trust vs. historian reading | USD 2–5M/yr at top-quartile execution | DCS, historian, LIMS, on-stream analyzer | L1 (manual with advisory screen) |
| SAG mill speed / load setpoint | 15-min cycle | Control-room operator | Throughput vs. liner wear; throughput vs. downstream slurry density | USD 4–8M/yr | DCS, historian, vibration sensors, ore-hardness model | L2 (augmented — operator approves) |
| Weekly maintenance work-package sequencing | Weekly | Maintenance planning lead | Equipment availability vs. production campaign timing; overhaul deferral vs. failure risk | USD 3–6M/yr (via recovery loss from unplanned stops) | CMMS, historian, equipment-criticality registry, mine plan | L1 (planning assistant output reviewed and approved) |
| Concentrate shipment lot commitment | Monthly | Commercial + operations joint | Grade variability vs. contractual penalty; port slot vs. stockpile cost; blend flexibility vs. customer spec | USD 5–15M/yr | LIMS, MES, ERP, port-logistics system, price-forward curve | L1 (scenario engine with human sign-off) |

*A full concentrator catalogue runs 20–30 rows. The four above span operations, maintenance, and commercial — the three domains where silo optimization most reliably conflicts when the objective function is absent.*

The friction nobody puts in the brochure: building this catalogue is not a data exercise, it is a series of interviews with people who have never been asked to write down what they decide. The senior operators who hold the twenty-five decisions in their heads are often the last to be consulted and the first to be automated around. Cataloguing with them, rather than about them, is the difference between a decision space that survives its first night shift and one that gets quietly overridden by the people who know better.

# 09

## WHEN AUTOMATION IS REALLY JUST DECISION LAUNDERING

A board that accepts this argument will start seeing programs that look like designed decision spaces and are not. Three look-alikes recur, and each fails the same way under one diagnostic question.

The first is decision theater: a model that produces predictions on a dashboard while every actual decision still routes through the same humans it always did. It looks like augmentation. The detection question: what decision changed owner when this went live? If the answer is none, you bought a viewer, not a decision system, and its return is the value of the screen, which is roughly zero.

The second is silo optimization wearing the language of orchestration. Each domain has a model, each model has an objective, and there is no equation above them. It looks like an autonomous operation. The detection question: is there a single number every agent is scored against, or does each domain keep its own scoreboard? If each keeps its own, you have funded the committee scene of section 04 in advance.

The third is decision laundering: a decision declared at full automation with no telemetry that lets a human see what the model did or reverse it. It looks like maturity and confidence. The detection question: when the model acted last Tuesday, who saw it, and could they have stopped it? If the action is invisible until it shows up in a monthly report, the autonomy is not a capability. It is an accountability gap with a confident name.

Each look-alike passes a casual review because each produces activity, and activity reads as progress. None survives the question of which decision moved under what objective function with what visibility. That question is the whole audit.

# 10

## INTELLIGENT DEBT: THE LIABILITY THAT ACCRUES WITHOUT AN INVOICE

There is a balance-sheet metaphor a CFO will recognize immediately, and it is the most useful frame this paper offers. Installing intelligence without designing the decision space creates a liability that accrues silently, the way debt accrues, until something forces a refinancing.

The mechanism is concrete. Each undesigned pilot creates a dependency: a model in production, an integration into the historian, a team that now reads its output, an expectation in the operating routine. None of it is tied to the objective function, so none of it produces margin, but all of it produces cost and obligation. The interest on this debt is the management attention spent maintaining pilots that cannot be defended, the credibility spent each time the board asks for a return that does not exist, and the opportunity cost of the recovery points left on the floor because the one pilot that could find them was never scoped. The principal comes due at the portfolio review, when a CTO has to stand in front of the committee scene of section 04 and write off most of the portfolio.

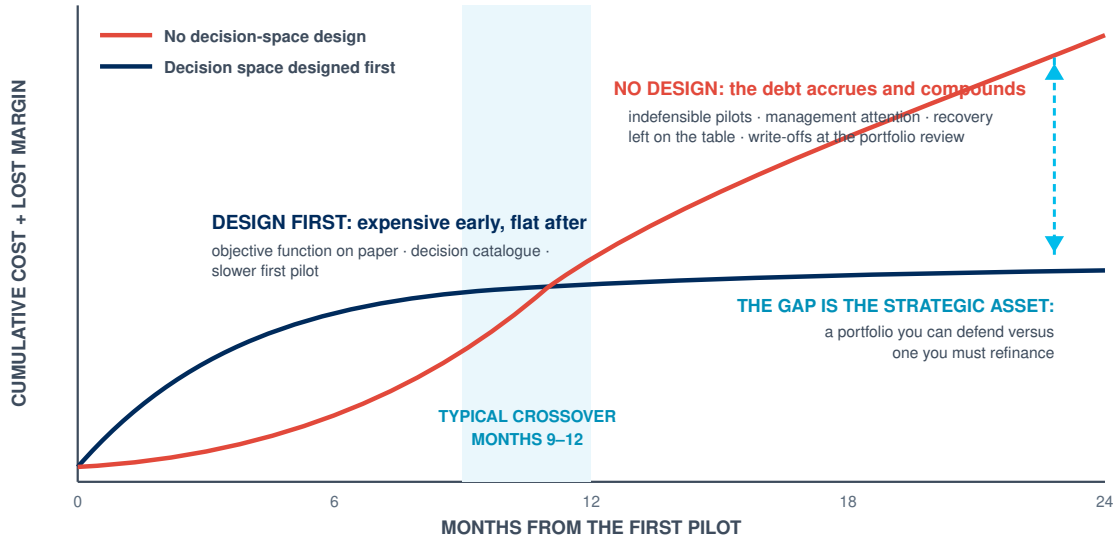


**Figure 3** The asymmetric cost over time of not designing the decision space versus designing it, with the crossover and the residual divergence marked

The asymmetry is the point. Designing the decision space first is expensive early: it forces the objective function out of people's heads and onto paper, it forces the decision catalogue, and it slows the first pilot. Skipping that design is cheap early and ruinous late. The choice a

board faces in 2026 is not whether to spend on AI. It is whether to pay the design cost now, in cash and discipline, or pay the debt later, in write-offs and lost recovery. Doing nothing is not the cheap option. It is the option that accrues interest you cannot see until the principal is due.

# 11

## THE TEST YOU CAN RUN ON ONE PILOT THIS WEEK

This argument is falsifiable in your own operation, and the test takes one pilot and an afternoon. Pick any AI initiative currently funded. For the single decision it claims to own, write Huang's five capabilities applied to that exact decision:

1. Recognize patterns: name the live data the pilot reads and the pattern it detects.
2. Recognize relationships: name the causal relationship it acts on, and how that relationship was validated as causal rather than coincidental.
3. Reason: name the trade-off it resolves and what it weighs against what.
4. Predict: name the outcome it forecasts and whether it reports a confidence band or a single number.
5. Plan an action: name the action it takes or recommends, and at what autonomy level.

The outcome is binary. If you can write all five for that decision without ambiguity, the pilot has an architecture, and its return is a question of execution. If any of the five comes back blank, that blank is the bug, and in our experience the blank falls almost every time on capability four or five: the pilot recognizes patterns beautifully and neither predicts in confidence bands nor plans a scored action. That blank is the architectural footprint the ninety-five percent leaves behind, visible on a single decision in a single afternoon.

Run it on three pilots before the next board meeting. The ones that come back fully written are the ones to fund again. The ones with a blank in row four or five are not failures of intelligence to be fixed with a better model. They are decisions you have not finished designing. Finish the design, or wind the pilot down. Both are defensible. Funding it again unchanged is not.

One last collapse to avoid. This argument is not a case for a better model: a better model improves the first three capabilities and leaves the blank in the fourth or fifth exactly where it was. It is not a case for better data, which is the same mistake one layer down. And it is not a case for an AI governance committee, because a committee produces a posture while a decision space is a specification you can audit. Every adjacent claim improves the intelligence. This paper's claim is that the intelligence is already the commodity; the work that remains is a description of your own operation's decisions, written in your own objective function, owned by your own board.

Huang gave the industry a definition in 2022 that sounded like awe, and in June 2026 he declared it delivered. Read on the second clause, it was always an instruction. The intelligence is cheap and getting cheaper. What stays scarce is knowing which decisions exist in your operation, at what level of autonomy each should run, under what single objective function, with what telemetry a human can read. Buy the commodity. Design the space.

**KEY INSIGHT —** Buy the commodity. Design the space.

# 12

## KEY INSIGHTS

1 Intelligence became the cheapest line in the AI budget. The scarce asset is the design of the decision space: which decisions exist, at what autonomy, under what objective function, with what override telemetry. That design is not for sale, because it is a description of your own operation.

2 Huang has held the same functional definition of intelligence since 2022 and declared it fulfilled at GTC Taipei in June 2026. When the supplier of intelligence spends four years telling you the intelligence is the commodity, the differentiation has already moved to what you point it at.

3 A model's quality is meaningless until you name the decision it serves. The decision comes first; the model comes second.

4 The pilots that pay are scoped against the variable that reconciles to margin. At world-concentrator scale, one percentage point of metallurgical recovery is worth on the order of US$100 million a year. The model that finds it costs a fraction of that; the design that lets it be found costs discipline.

5 The honest caveat is also the frontier: no operator has yet published a validated, full cross-domain orchestration running on margin in copper. The leaders publish single-plant wins. The global optimum is still being engineered, and that is why the design discipline, not the model, is where the race is.

6 An objective function is not a KPI. It is the precondition that makes KPIs meaningful. And no model, at any price, solves a target problem.

# 13

## SOURCES


1. Thompson, B. — *An Interview with Nvidia CEO Jensen Huang about Manufacturing Intelligence* — Stratechery, 2022. Primary source of the functional definition of intelligence quoted in §02. stratechery.com/2022/an-interview-with-nvidia-ceo-jensen-huang-about-manufacturing-intelligence/
2. Huang, J. — Address at the Cambridge Union, November 4, 2025 — "Intelligence is about to be a commodity." Coverage: news.futunn.com/en/post/65135996/
3. Fridman, L. — *Lex Fridman Podcast #494: Jensen Huang*, March 23, 2026 — AGI-achieved claim (functional definition, qualified) and the commodity-versus-humanity distinction. lexfridman.com/jensen-huang/ · Coverage: Fortune, March 30, 2026.
4. NVIDIA — Jensen Huang keynote, GTC Taipei, June 1, 2026 — "agentic AI has arrived… useful AI has arrived"; tokens as profitable units of revenue; Vera Rubin to full production. Coverage: SiliconANGLE, June 1, 2026. The "utterly parabolic" demand remark is from Huang at Dell Technologies World, May 18, 2026 — blogs.nvidia.com/blog/dell-technologies-agent-enterprise-ai/.
5. MIT Project NANDA — *The GenAI Divide: State of AI in Business 2025*, July 2025 — ~95% of enterprise GenAI initiatives without measurable return; directional signal from interview data. github.com/aidecentralized/nandapapers · Coverage: Fortune, August 18, 2025.
6. NTT DATA × MIT Technology Review en Español — *Autonomy in the Mining Business: Beyond Technology, the Challenge for Talent and Culture*, late 2025 — 72% of surveyed mining organizations reported their projects delivered only part of the expected benefits. Coverage: Mexico Business News.
7. Parasuraman, R., Sheridan, T.B., Wickens, C.D. — *A Model for Types and Levels of Human Interaction with Automation* — IEEE Transactions on Systems, Man, and Cybernetics, Part A, 30(3), 2000. DOI 10.1109/3468.844354.
8. International Copper Association — *Freeport-McMoRan Looks to the Future with Artificial Intelligence*, July 2020 — Bagdad ML program: ore-type classification (seven ore types), ~10% throughput gain, +9,000 tonnes of copper, capital expansion avoided. internationalcopper.org. The +1 percentage-point recovery gain and the ~85,000 t/day quarterly throughput are documented in independent published analyses of the same Bagdad deployment.
9. Antofagasta plc — *Annual Report 2025* — SIRO/operational-recommendation results: Centinela recovery +0.75 pp (and +1% throughput), Antucoya recovery +0.8 pp, sulphuric-acid and water savings. antofagasta.co.uk/media/4905/49913-antofagasta-ar25-web-ready.pdf · Coverage: International Mining, March 2026.

10. Boliden — *Can AI Optimize Our Processes?* — reinforcement-learning mill control at Aitik outperformed existing control in simulation; further study required before production deployment. boliden.com/news/can-ai-optimize-our-processes/
11. BHP — *BHP and Microsoft use AI to lift Escondida copper recovery*, May 30, 2023, and *AI is improving performance across global mining operations*, January 2026 — program scope and reported water and energy savings; production and recovery baselines used in the §04 arithmetic. bhp.com

**NTT DATA · Natural Resources Practice · 2026** · biztalksnttdata.com

*For information about designing the decision space behind GenAI investments in mining operations, contact: Jaime Rebolledo, NTT DATA Chile · jaime.rebolledocampos@nttdata.com*

*For more on this topic, contact:* ***Jaime Rebolledo****, Director Value Proposition Mining, NTT DATA ·* jaime.rebolledocampos@nttdata.com.