EXECUTIVE BRIEF · DIRTY DATA EXPERT

Executive Assumptions

The greatest risk in every AI initiative isn't the model. It's the belief, never tested, that the business underneath it is ready.

By **Gary Segal** · Founder, Dirty Data Expert

August 2026

Most AI failures begin as assumptions, not technology failures. Technology rarely creates the original problem — more often, AI simply exposes assumptions that have existed inside the business for years, unexamined because no one was ever forced to examine them.

Executives do not usually experience this as a data problem. They experience it as a delivery problem: the pilot stalls, the model underperforms, the dashboard nobody trusts, the automation that has to be quietly re-checked by a human anyway. By the time it reaches that stage, the postmortem tends to focus on the technology. The actual failure happened months earlier, in a planning meeting, when a reasonable-sounding assumption went unchallenged.

> The greatest AI risks often begin as executive assumptions: that clean data is decision-ready, that ERP structure equals business understanding, and that AI will fix weaknesses the business has never governed.

The argument in brief

- **AI does not create business understanding — it amplifies what is already there.** Structure, definitions and governance that are weak before AI stay weak after it; they simply get executed faster and at greater scale.
- **Ten specific assumptions carry most of the risk,** and they cluster into three categories: what the business believes about its data, what it believes about its own capability, and what it believes AI itself can do.
- **None of the ten are exotic.** Each one sounds reasonable in a steering committee. That is precisely why they survive unchallenged into production.
- **The fix is not more technology.** It is a disciplined, recurring practice of surfacing and testing assumptions before they are built into a roadmap.

Where the assumptions hide

Assumptions are rarely stated out loud. They live inside data structures, system configurations, legacy reports and operational habits — inherited, not decided, and therefore never tested. The ten below are the ones we see most often, grouped by the kind of confidence they misplace.

EXHIBIT — THE TEN ASSUMPTIONS, IN ORDER

01

BELIEF ABOUT THE DATA

Clean Data = Decision-Ready Data

Data can be accurate and still be structurally unsuitable for forecasting, segmentation, governance and AI.

> **EXAMPLE —** Inventory records are 99% accurate at the SKU level, but the product categories were built for accounting, not demand planning. A forecasting model trained on that structure produces confident, wrong forecasts for entire categories.

02

BELIEF ABOUT ITSELF

ERP Structure = Business Understanding

ERP systems manage transactions. They do not automatically create business intelligence.

> **EXAMPLE —** The ERP correctly records every invoice and shipment, but two divisions hold the same customer under different account IDs. "Customer lifetime value," reported to the board as one number, is quietly undercounted for every multi-division client.

03

BELIEF ABOUT AI

AI Will Fix Governance

AI depends on governance. It cannot create accountability where none exists.

> **EXAMPLE —** An AI-driven approval workflow is deployed to speed up procurement, but no one owns the vendor master data. The system now approves duplicate vendor records faster than a human ever could.

04

BELIEF ABOUT ITSELF

Dashboards Create Clarity

Dashboards only visualize what already exists. They cannot resolve conflicting definitions underneath.

> **EXAMPLE —** Sales and Finance each have a "revenue" dashboard, built on different definitions of when a sale counts. The two numbers have never matched, and no amount of dashboard redesign will fix that.

05

BELIEF ABOUT AI

Automation Improves Decisions

Automation accelerates decisions. It does not guarantee better ones.

> **EXAMPLE —** An automated reorder system cuts replenishment decisions from days to minutes — and keeps reordering a discontinued SKU, because no one updated the exception rules. Automation just executes the bad decision faster.

06

BELIEF ABOUT THE DATA

Segmentation Is a Reporting Exercise

Segmentation is the operating system beneath forecasting, pricing, replenishment and AI — not a chart produced after the fact.

> **EXAMPLE —** Marketing segments customers one way for campaigns; supply chain segments them a different way for replenishment. A model trained on one segmentation quietly fails when applied to decisions made on the other.

07

BELIEF ABOUT ITSELF

Technology Ownership = Business Ownership

Owning the platform does not mean owning the business definition, accountability or outcome.

> **EXAMPLE —** IT keeps the CRM running flawlessly, but no business owner has ever decided what counts as a "qualified lead." The field means ten different things to ten different sales reps.

08

BELIEF ABOUT THE DATA

More Data = Better Decisions

More data often creates more noise. Better decisions require better structure, not simply more information.

> **EXAMPLE —** Three new data sources are added to a forecasting model expecting a lift in accuracy. Performance drops instead, because two of the sources use conflicting product hierarchies no one reconciled first.

09

BELIEF ABOUT AI

AI Understands Context

AI understands patterns in data. It does not automatically understand business intent, strategy or human priorities.

> **EXAMPLE —** A support chatbot correctly spots complaint patterns in ticket language, but recommends the same fix for a churn-risk enterprise account as for a one-time retail buyer — it has no sense of account strategy.

10

BELIEF ABOUT ITSELF

Clean Transactions = Clean Strategy

A business can execute transactions perfectly while still making poor strategic decisions, because the underlying structure is weak.

> **EXAMPLE —** Every order, invoice and payment reconciles perfectly, yet the business is discounting its highest-margin product line, because no one connected transaction data back to the pricing strategy it was meant to serve.

What this looks like in practice

These assumptions rarely announce themselves. They surface as symptoms that get treated individually, rather than traced back to a shared root: product categories that don’t match between systems, customer segments that meant something different two reorganizations ago, ownership of a KPI that three departments each believe belongs to someone else, master data that changes without anyone downstream being told. Each looks like a local, fixable glitch. Together, they are the reason the same class of AI project keeps stalling at the same stage.

The executive response

Before scaling AI further, the assumption itself — not just the system built on it — needs to be put in front of the people who can validate it.

1 **List the assumptions the roadmap is standing on**

Every AI initiative already has a working list of assumptions; it is just implicit. Make it explicit before the next investment decision, not after the pilot underperforms.

2 **Validate them with the people who own the data, not the platform**

The system owner can confirm the pipeline runs. Only the data owner can confirm the definition underneath it still means what everyone assumes it means.

3 **Prioritise the assumptions that actually drive decisions**

Not every assumption is worth testing immediately. Start with the ones sitting directly under a decision the business is about to make at scale.

4 **Govern the assumptions that recur**

Some assumptions resurface in every project — segmentation, ownership, definitions. Those belong in standing governance, not in a one-off workshop.

> The most dangerous assumptions are the ones executives never realise they are making. AI does not remove assumptions — it scales their consequences.
>
> Next step: use this brief to frame the conversation before the next investment in AI, analytics or automation — while the assumptions are still cheap to test.