

**DIRTY DATA EXPERT**

*Structure Before Intelligence™*

---

# WHEN THE AGENT ACTS

**Executive Brief**

Gary Segal  
[www.dirtydataexpert.com](http://www.dirtydataexpert.com)  
[gary@dirtydataexpert.com](mailto:gary@dirtydataexpert.com)

Chatbots describe your data. Agents act on it. When the record is wrong, you don't get a wrong answer. You get a wrong action.

Chatbots describe your data. Agents act on it. For three years, most AI in business has been conversational. A chatbot reads your data and gives you an answer. If the data is wrong, you get a wrong answer, and a person decides what to do with it.

That is changing fast. Agentic AI does not just answer. It raises purchase orders, approves invoices, moves stock and adjusts prices. Gartner predicts that 40% of enterprise applications will include task-specific AI agents by the end of 2026, up from less than 5% in 2025.

The risk changes with it. When an agent reads a dirty record, you don't get a wrong answer. You get a wrong action, at machine speed, with nobody in the loop to stop it.

This guide explains how that happens, what it costs, and the three questions every leadership team should answer before giving an AI agent the authority to act.

## The yoghurt that broke Christmas

---

*An illustrative scenario, built from data failures common in inventory-carrying businesses. It hasn't happened to your company. Yet.*

A national FMCG distributor goes live with an AI replenishment agent in October, just in time for festive-season volumes. The pilot went well: buyer workload fell by 60%, and management is pleased.

Three years earlier, during an ERP migration, one yoghurt SKU was loaded with its unit of measure as "each" instead of "case of 12". The buyers knew and quietly worked around it. Nobody fixed the record, because everybody "knew".

The agent didn't know. It read stock at a twelfth of reality, calculated a severe shortfall ahead of December, and raised a purchase order twelve times too large. It flagged the order as urgent and chose the premium supplier for speed.

A second agent in accounts payable received the invoice. The three-way match between purchase order, goods received note and invoice passed perfectly, because all three carried the same wrong data. Payment was released.

The dairy arrived to full cold rooms, so stock was rerouted to third-party cold storage at emergency rates. Because the cold rooms now showed as full, the agent cancelled orders for three other chilled lines. Those lines ran out in the week before Christmas.

By January there were expired write-offs, premium freight bills, two unhappy retail chains, and an audit committee asking who had approved a purchase order that no person had approved. The board suspended the AI programme.

The yoghurt record is still wrong.

## How it happens: five data failures agents can't see

---

Every one of these failures already exists in most inventory-carrying businesses. People catch them today because they know the data is wrong. An agent doesn't know, so it acts on the record as written.

| Data failure                                                                                | What your people do today                 | What an agent does                                                       |
|---------------------------------------------------------------------------------------------|-------------------------------------------|--------------------------------------------------------------------------|
| Wrong unit of measure — "each" instead of "case of 12"                                      | Mentally convert, every time              | Orders twelve times too much, or too little                              |
| Duplicate records — one supplier, SKU or customer under two codes                           | Know which code is the "real" one         | Splits stock across both codes, sees two shortfalls, orders twice        |
| Stale master data — old prices, lead times, discontinued lines still active                 | Check with the supplier before ordering   | Uses the old lead time and price, or orders a line that no longer exists |
| Overloaded fields — a notes field that says "DO NOT ORDER", a flag reused for a new purpose | Read the note and act on it               | Reads only the structured fields and ignores the note                    |
| Systems that disagree — ERP, warehouse and finance hold different versions of the truth     | Know which system to trust for which fact | Trusts whichever system it is connected to                               |

Cross-checks rarely catch any of these. When the purchase order, goods received note and invoice are all built from the same wrong record, they agree with each other. That is why the three-way match in the yoghurt story passed.

## The control nobody documented

The most important control in the yoghurt story was never written down. It was a buyer who knew the unit of measure was wrong and corrected it in their head.

Every workaround in your business is a control like that. It works, it has worked for years, and it exists only in someone's memory. Auditors can't test it, because nobody knows it is there.

When an agent takes over the task, that person leaves the loop, and the control leaves with them. Nothing records that it ever existed. On paper, every documented control still passes.

Before any agent goes live, ask the people whose work it will replace one question: "What do you check that the system doesn't?" Every answer is either a record to fix or a rule to write down.

## Seven implications for leaders

Agentic AI turns data quality from a reporting issue into an operational risk. That makes it a leadership decision, not an IT task.

- 1. Dirty data now causes actions, not just wrong reports..** A bad record used to mislead a report. Now it raises an order, releases a payment or cancels a delivery.
- 2. Speed multiplies every error..** A buyer makes one wrong order and spots it. An agent can make hundreds before anyone looks.
- 3. Your existing controls may not see it..** Controls that compare documents fail when every document comes from the same wrong record.

**4. Accountability must be named before go-live..** When an agent approves something, a person must still own the decision. Decide who that is before the audit committee asks.

**5. A good pilot doesn't prove readiness..** Pilots run on a clean sample, in a quiet period, with people watching closely. Live operations have none of that.

**6. Agents feed each other..** One agent's output is the next agent's input. In the yoghurt story, one wrong record moved through purchasing, payments and warehousing.

**7. Fixing the record is cheaper than suspending the programme..** The yoghurt error was a single field. The board suspended the whole AI programme instead.

## This has happened before: Target Canada and Knight Capital

---

The yoghurt story is illustrative. The pattern behind it is not. Automated systems have already acted on bad data, and on bad instructions, with enormous costs.

### Target Canada: automation on dirty master data

Target opened its first Canadian stores in March 2013. Merchandisers had to load about 75,000 products into a new SAP system, largely by hand, under severe time pressure.

According to Canadian Business, the information in the system was accurate about 30% of the time. In the US, it was between 98% and 99%. Dimensions were entered in inches instead of centimetres, or in the wrong order. Sometimes the wrong currency was used.

The replenishment system acted on that data. Stores had empty shelves while distribution centres overflowed with stock.

In January 2015, less than two years after opening, Target announced it would close all 133 stores, affecting about 17,600 employees. CBC News reported a write-down of about \$5.4 billion.

Target Canada didn't use AI agents. It didn't need to. Automation acting on dirty data was enough.

### Knight Capital: automation with no brake

Knight Capital shows the second risk: speed with nothing to stop it. On 1 August 2012, a faulty software deployment at the trading firm sent more than 4 million orders into the market in 45 minutes while trying to fill 212 customer orders.

The SEC found that Knight lost more than \$460 million. Its systems had sent 97 automated warning emails before the market opened, and nobody acted on them.

The cause was software, not dirty data. The lesson still applies: an automated system that can act must also be able to be stopped.

## Structure Before Intelligence™: three questions before giving an agent authority

---

AI agents fail on structure, not intelligence. The model in the yoghurt story worked exactly as designed. The data and the decision rights around it did not.

Before any agent is allowed to act, the leadership team should be able to answer three questions.

### **1. Is the data the agent will act on fit to act on?.**

Data that is good enough for a report is not always good enough for an action. List the fields the agent will rely on, such as unit of measure, supplier, price, lead time and capacity. For each one, name an owner and state when it was last checked.

### **2. What do our people know that our systems don't?.**

Find the workarounds before the agent removes the people who carry them. Each one must become a corrected record or a written rule. Anything left in people's heads is a control you are about to lose.

### **3. What stops the agent when it is wrong?.**

Set limits on value, volume and deviation from normal patterns. Name the person who owns each agent's decisions. Make sure someone can pause it immediately, and that every action is logged with the data and reason behind it.

If the answer to any of these is "we're not sure", the agent is not ready. The organisation isn't either.

## **Agent readiness checklist**

---

Use this before any agent is given authority to buy, pay, move stock or change prices. Every unticked box is a risk someone should own.

### **DATA**

- Every field the agent acts on has a named owner
- Units of measure and pack sizes are checked for every item the agent can order
- Duplicate supplier, customer and product records are found and merged
- Discontinued items and inactive suppliers are blocked, not just labelled
- Key systems (ERP, warehouse, finance) agree on the facts the agent uses

### **PEOPLE AND KNOWLEDGE**

- The people whose work the agent replaces have been asked what they check that the system doesn't
- Every known workaround is now a corrected record or a written rule

### **CONTROLS AND ACCOUNTABILITY**

- Limits are set on order value, quantity and deviation from normal patterns
- A named person owns each agent's decisions
- Someone can pause the agent immediately
- Every action is logged with the data and reason behind it
- Agents that feed each other have been tested together, not only on their own
- The pilot included live conditions, including a peak period

### **Find your yoghurt records before your agents do**

The AI Risk Diagnostic shows where your data, controls and decision rights are not yet ready for AI that acts.

## About Gary Segal

---

Gary Segal is a consultant, author and executive speaker, known as the Dirty Data Expert. He created the Structure Before Intelligence™ framework, which treats AI failure as a leadership and data governance problem, not a technology problem.

He works with boards and executive teams in organisations that carry inventory, across FMCG, retail, manufacturing, logistics and financial services. He works in South Africa and internationally, as a keynote speaker and as a consultant.

[gary@dirtydataexpert.com](mailto:gary@dirtydataexpert.com) · +27 68 645 1673 · [www.dirtydataexpert.com](http://www.dirtydataexpert.com)

---

## Sources

- The Last Days of Target Canada, Canadian Business
- Target closes all 133 stores in Canada, CBC News
- SEC charges Knight Capital with violations of market access rule, US Securities and Exchange Commission