



Signals in the Noise

Aurora Market wants to become "data-driven," but the journey doesn't start with clarity. It starts with **contradictions**.



Act I – Confusion, Not Vision

The Complaints

Different departments complain about opposite problems:

- Marketing claims customer churn is completely random and unpredictable
- Operations insists forecasting is fundamentally broken — yet each store blames the model rather than their dodgy inputs
- Finance insists fraud is rising sharply, yet can't agree on which transactions actually look suspicious

The Reality

Everyone wants answers. Everyone disagrees on the **questions.**

In this fog of confusion and misalignment, Aurora forms a small ML team and tells them vaguely: "Fix the data."

The team has no idea where to start. No clear objectives. No prioritisation. Just noise.

The Work Begins

With conflicting demands and unclear priorities, the newly formed ML team starts digging through years of messy data, trying to make sense of the chaos.



Act II – First Models, First Failures

Early Success

The ML team jumps into churn prediction. Early results look brilliant – almost **92% accuracy**.

Celebration. Congratulatory emails. A polished slide deck for leadership.

1

2

3

The Truth

The model was trained on highly imbalanced data, mostly predicting "no churn." The impressive accuracy metric masked complete uselessness.

Marketing loses trust immediately.

The Question

Then someone from Marketing asks a simple question: "Which customers should we actually target to prevent churn?"

Silence. The model can't answer.

📌 **Lesson 1:** Business metrics \neq model metrics. High accuracy means nothing if the model can't drive actual decisions.

Act III — Forecasting Chaos

Operations insists demand forecasting must be improved urgently, but the data feeds are an absolute mess:

Format Inconsistencies

One store uses decimal commas; another uses decimal points. Simple numbers become parsing nightmares.

Missing Context

Promotions aren't labelled anywhere in the system. Spikes in demand look random.

False Signals

Historical stockouts appear as "zero demand" — teaching the model that nobody wants these products.

The first forecasting model goes into production — and fails spectacularly. One store receives **triple stock** of cereal during a low-demand week. Shelves overflow; staff joke about building a "Cereal Castle."

Operations blames ML. ML blames data quality. Leadership blames everyone.

📄 **Lesson 2:** Data problems surface only when a model amplifies them. What looks like minor inconsistencies become catastrophic failures at scale.

Act IV – A Breakthrough from the Wrong Direction



The Unexpected Discovery

Whilst chaos continues elsewhere, Finance quietly tests an anomaly-detection model for fraud detection.

It works surprisingly well.

But the real discovery isn't fraud at all. The model flags weekly patterns of suspicious returns that turn out to be **staff errors**, not theft. Aurora realises they have a process problem, not a criminal one.

This unexpected insight shifts the entire tone: Even imperfect models can surface **useful truths the company wasn't looking for.**

The breakthrough comes from solving a different problem than intended. Sometimes the most valuable insights are the ones you stumble upon.



Act V – Business Friction

1

The Pitch

The ML team attempts dynamic pricing next. Leadership gets excited – too excited about the possibilities.

2

The Vision

They want algorithmic prices updated every hour, responding to demand in real-time like tech giants.

3

The Reality

Legal intervenes forcefully: "You can't price products this erratically; customers will revolt and regulators will investigate."

The project stalls completely because the **business side changed the objective mid-way**, hoping ML could somehow bend regulations and customer psychology.



Lesson 3: Solving the wrong business problem faster isn't progress. Technical feasibility doesn't equal business viability.



Act VI – Finally: Convergence

After failures, misunderstandings, painful lessons, and unexpected wins, the company finally starts aligning around shared understanding:



Marketing

Provides clearer churn definitions and actively collects new behavioural signals that matter.



Operations

Cleans data pipelines, properly labels promotions, and reports anomalies systematically.



Finance

Integrates ML insights directly into auditing workflows rather than treating them as separate.



Leadership

Stops demanding "AI magic" and starts asking structured, answerable questions instead.



Final Message

Aurora's transformation isn't a straight path.

It's a cycle of **false starts, misaligned incentives, poor data quality, over-ambitious goals, unexpected insights, and gradual alignment.**

Заставило взглянуть на себя трезво

Machine learning didn't save the company. It **forced the company to see itself clearly** — exposing hidden problems, surface assumptions, and reveal uncomfortable truths.

And that brutal clarity is where genuine impact truly began.



The Real Lesson

0

Perfect Models

Success isn't about flawless algorithms

100%

Alignment Required

It's about organisational readiness

∞

Iterations Needed

Progress comes through persistent learning

The companies that succeed with ML aren't the ones with the best data scientists. They're the ones willing to confront uncomfortable truths about themselves — and change accordingly.