DATA & ANALYTICS

From Dashboards to Decisions

Why the next generation of business analytics will be judged by the decisions it changes, not the reports it produces.

Berry Ostrovsky · 2026-08-15 · 8 min read

EXECUTIVE SUMMARY

- Reporting capacity has grown far faster than decision capacity. More dashboards rarely produce more decisions.
- A dashboard answers "what happened". A decision needs a recommended action, an owner and an expected effect.
- AI is most useful upstream of the chart: consolidating messy commercial data, classifying situations and ranking where to look first.
- The practical unit of analytics output is a short, ranked list of actions, not a page of visuals.
- The analyst's role shifts from producing numbers to defining the decision, the logic behind it and its limits.

Almost every commercial organization has invested in reporting. Data warehouses were built, BI tools were rolled out, and teams can now see sales by product, customer, channel and week within seconds. Yet the familiar complaint remains: leadership sees more than ever and still argues about what to do next.

The gap is not a data gap. It is a decision gap. Reporting capacity has grown far faster than the organization's ability to convert what it sees into a small number of clear, owned actions. This article looks at why that happens, what a decision-oriented analytics output actually looks like, and where AI genuinely helps rather than adding another layer of automated description.

The reporting layer solved the wrong bottleneck

When BI tooling matured, the assumption was straightforward: if people could see the numbers faster, they would act faster. That assumption held while access to data was the constraint. Once access became cheap, the constraint moved.

In most commercial teams the real bottleneck is attention. A category manager responsible for thousands of product-store combinations does not need another view of the same data. They need to know which twenty of those combinations deserve a conversation this month, and why.

A dashboard is optimised for exploration by someone who already knows what they are looking for. It is a poor instrument for prioritisation, because it treats every cell of the dataset as equally worth inspecting.

What separates a report from a decision

A useful test: read the output and ask what changes tomorrow morning. If nothing changes, the output described the business rather than steering it.

A decision-oriented output tends to carry four elements that a standard report does not.

- A specific situation, described at the level someone can act on — this product, in this chain, at this price point.
- A recommended action, stated plainly, with the alternative that was rejected.
- An expected effect, expressed as a direction and rough magnitude rather than false precision.
- An owner, because an insight without an owner quietly becomes background noise.

None of this requires more sophisticated modelling. It requires the analysis to end one step later than it usually does.

Where AI actually adds value

The most visible use of AI in analytics — asking a chatbot to summarise a chart — is also the least valuable. Summarising a chart that a competent manager can already read compresses a task that was never expensive.

The expensive work sits earlier in the chain. Commercial data arrives inconsistent: the same product appears under different codes, store hierarchies drift, historical data is incomplete, and definitions differ between systems. Reconciling that has traditionally consumed a large share of an analyst's week.

Language models and structured automation are well suited to this messy middle: matching entities that are not identical but refer to the same thing, classifying situations against a defined rule set, and flagging records that do not behave as expected. The output is not a narrative — it is a cleaner, comparable dataset on which ordinary business logic can run.

Ranking beats describing

Once data is comparable, the analytical question becomes ordering rather than description. Which gaps are worth closing first?

A simple structure works well in commercial contexts. Classify each product-customer combination into a small number of states: actively selling, previously selling and now dormant, or never listed at all. Each state implies a different conversation. Dormant combinations point to an execution or availability problem. Never-listed combinations point to a range or negotiation opportunity.

Estimating potential for each state does not require a forecasting model to be accurate to the shekel. A directional estimate based on comparable stores or comparable customers is usually enough to rank opportunities, and ranking is what the team actually needs. The discipline is to be explicit that the number is directional and to say what it assumes.

The cost of false precision

Decision-support output loses credibility quickly when it implies more confidence than the underlying data supports. A pipeline that presents an opportunity as being worth an exact amount invites a debate about the number instead of a debate about the action.

Presenting a range, naming the assumption behind it, and separating what is measured from what is estimated tends to survive scrutiny better. It also makes the analysis reusable: when an assumption turns out to be wrong, only that assumption needs revisiting.

This is a communication choice as much as a technical one. Analysts who state limits clearly are trusted with larger questions.

What this changes for the analyst

If AI absorbs a meaningful share of consolidation, formatting and first-pass description, the remaining value of the role concentrates in areas that are harder to automate: framing the question, choosing what to compare, deciding which differences are meaningful, and understanding the commercial context well enough to know when a number is misleading.

That is a shift in emphasis rather than a replacement. The analysts who benefit are the ones who move upstream into problem definition and downstream into the decision itself, rather than defending the middle ground of report production.

It also raises the bar on judgement. When output is cheap to generate, the scarce skill is deciding what deserves to exist at all.

A practical way to start

Improving decision quality does not require a platform rebuild. A narrower approach usually moves faster.

- Pick one recurring commercial decision, such as monthly range reviews or promotional prioritisation.
- Write down how that decision is made today, including the inputs people actually trust.
- Replace the report with a ranked list of no more than twenty items, each with a recommended action and an owner.
- Use automation for consolidation and classification, and keep the ranking logic transparent enough to be challenged.
- Review after one cycle: how many items were acted on, and what stopped the rest?

The last question is the important one. It usually reveals whether the obstacle was analytical, organizational, or simply that nobody was accountable for the follow-up.

The value of analytics is not in the volume of what it explains but in the quality of what it changes. Dashboards remain useful as an interface for exploration, and they will not disappear. But the next meaningful improvement in commercial analytics is unlikely to come from another visualisation layer.

It will come from shortening the distance between a dataset and a decision — using automation for the reconciliation work, being disciplined about ranking, and being honest about what the numbers can and cannot support. The measure of success is simple enough to state: how many decisions were made differently because the analysis existed.

KEY TAKEAWAYS

- Measure analytics by decisions changed, not reports delivered.
- Use AI for consolidation and classification before using it for narrative.
- Rank a short list of actions instead of describing the whole dataset.
- Prefer directional estimates with stated assumptions over false precision.
- Every recommendation needs an owner, or it will not be executed.