

**Dashboard:** a screen that shows what already happened. **Decision:** a choice someone has to make, and own.

We have built a lot of the first kind. Screens everywhere, in every tool, on every wall. Somewhere along the way, the screen became the point.

So here is a test. If your analytics stack went dark tonight, who would notice by Friday? Not who would miss the charts. Who would be unable to make a call they actually make. For most teams the honest answer is almost no one, and that is the tell. The work was never waiting on the dashboards. They ran alongside it, watched and admired and rarely load-bearing.

A decade of spend has not changed this. In 2025, 98.4% of organizations increased their data and AI investment, and only about 46% reported real value from it (Wavestone, 2025). Spend is not the constraint. Something else is.

This paper names it. The gap between insight and action is mostly an adaptive problem wearing a technical disguise. People do not fail to act because a chart was ugly or the data was late. They fail because no one owns the decision, because acting is riskier than sitting still, and because the call is genuinely hard. That matters for the fix, because most of what we sell as a fix, a better dashboard, cleaner data, now AI, is a technical answer to a problem that is not technical. The only way design helps is if it forces the one thing design can force: a name on the decision, and a person who can be held to it.

***The one move.** Stop building dashboards forward from your data. Build them backward from your decisions, and put a name on each one. Every metric on a screen should answer five questions: the decision it informs, who owns that decision, the early signal that should move it, the threshold that should trip it, and the action to take when it does. Four of those*

*are scaffolding. The owner is the load-bearing one. A metric that cannot name its decision and its owner is decoration. Cut it.*

## The report graveyard

Open any company's BI tool and you will find the graveyard. Dashboards built, presented once, never opened again. In my experience most of an estate goes stale within a quarter, and you should not take my word for it. Pull your own access logs and count what has not been opened in 90 days. Someone is still paying to keep all of it alive.

The numbers around it are stubborn. About a quarter of large firms call themselves data-driven, about a fifth say they have built a data culture, and those figures have barely moved across a decade of investment (Davenport and Bean, MIT Sloan, 2023). Ask leaders why, and roughly 80% point to people, culture, and process, not technology. Yet under 2% of investment goes to data literacy (same source). We diagnosed the disease and kept funding a different cure.

Here is what that looks like up close. Earlier in my career I built an analysis matching recruiters to customers. I scored each recruiter by their submittal-to-hire ratio, broke it down by the type of role coming in, and produced a clean recommendation: send these jobs to these recruiters and you fill more, faster. It pointed at a real, recurring decision. It was never used. Not because anyone disagreed with it. It died for two reasons that had nothing to do with the analysis. No one's actual job was to act on it, and rerouting work across the floor was hard to make happen. The design was fine. The decision had no owner and no easy path to action. That is most of analytics, in one chart.

# Why a better dashboard won't fix it

So why does good analysis stall? The research is consistent, and it does not flatter the tools.

Pfeffer and Sutton named the oldest version the knowing-doing gap. Organizations know what to do and fail to do it, and the gap is structural, not informational. More information rarely closes it, because information was never the binding constraint.

Then there is who pays for a wrong call. Acting on data exposes you. A decision carries your name, and if it goes wrong the blame is easy to trace. Inaction is quiet and shared. Tetlock showed that accountability deepens the pull toward the status quo exactly when action could create a visible loser. People under scrutiny stop reaching for the best outcome and start avoiding the worst one they could be blamed for (Tetlock, 1994). A dashboard does not touch that.

It is worse on the high-stakes calls, which is where analytics is supposed to earn its keep. People drop a model that makes a visible mistake faster than they drop a colleague who makes the same one. Researchers call it algorithm aversion, and it runs strongest on the biggest decisions (MIS Quarterly, 2024). The more automated your analytics, the more your best people quietly route around them.

And the finding that should stop the “just fix the culture” crowd too: even when companies build an analytical culture, it does not reliably produce data-driven decisions. A study of 305 firms found that leadership support and data quality build the culture, and the culture still does not consistently turn into the decisions (Szukits and Moricz, 2024). The last step resists from both sides.

Ron Heifetz drew the line that explains it. Technical problems have a known solution and a clear owner. Adaptive problems are ambiguous, involve competing values, and require someone to change their behavior. Most decisions analytics is meant to inform

are adaptive. Most analytics investment is technical: better charts, cleaner pipelines, faster queries. Heifetz put it plainly: a common cause of failure is treating an adaptive challenge as if it were a technical one.

This is the part most of these papers skip, so I will say it plainly too. A better dashboard discipline is also a technical fix. If “build better reports” is the whole answer, it lands in the same graveyard, just tidier. Design earns its keep in exactly one way. It can force the thing the adaptive problem is missing: a named owner, on the hook for a specific decision. Everything that follows is built to make that one thing unavoidable.

## The framework: Decision-Backward Design

Build reports backward from the decision, and make the decision have a person. That is the whole idea. It has a few interlocking parts, and they all serve the owner.

None of these parts are new on their own. Decision intelligence said start from the decision (Kozyrkov). Metric and driver trees mapped the levers. RACI and its cousins named owners. What is missing in practice is a single, ruthless test applied to every metric on every screen, and a hard gate on the one word everyone fudges, “leading.” That is the contribution.

**The Decision Line.** Every metric you put in front of someone carries one line that answers five questions:

- **Decision:** what recurring decision does this inform?
- **Owner:** who makes it, by name, and can be held to it?
- **Leading indicator:** what early signal should move it? Two gates, both required. Predictive, meaning it runs reliably ahead of the outcome. And actionable, meaning intervening on it changes that outcome. Fail the first and it is noise. Fail

the second and it is a lagging number with a head start, which you can watch but not change.

- **Trigger:** at what threshold does the decision actually change?
- **Move:** what is the recommended action when the trigger trips?

Four of those are scaffolding. The owner is the beam they hang from. A metric that cannot name its decision and its owner does not belong on the screen, and most cannot.

Here is the honest part. Writing a name in a box does not create accountability. It makes the absence of it impossible to ignore. When you run a real decision through the Line, reach “owner,” and the room goes quiet, you have found the actual problem, and no dashboard was ever going to solve it. Real accountability comes from incentives, decision rights, and a mandate from the top, and the Decision Line supplies none of those. What it does is turn a vague cultural complaint into a fact on the page. That is worth more than it sounds, because most organizations never make the vacuum visible enough to fix.

Take a software company deciding which customers to save. Churn is the lagging indicator, and by the time it lands on a dashboard the customer is gone. The Line forces something better. Decision: do we open a save play on this account? Owner: the named success manager for that book. Leading indicator: usage down more than 30% for two straight weeks, paired with a support escalation. Predictive, because it runs ahead of churn. Actionable, because a save play can still change it. Trigger: both true. Move: open a retention play inside 48 hours. That metric makes a decision. A churn dashboard reports a funeral.

**The Decision Register.** The Line works on one metric. The Register works on the portfolio. It is a short catalog of the handful of recurring decisions that actually matter, each with its owner and the one or two signals that drive it:

Decision	Owner	Leading signal	Move
Save this account?	Named success manager	Usage down >30% for 2 weeks, plus a support escalation	Open a retention play in 48 hours
Re-run this campaign?	Named demand-gen lead	Cost per acquisition up >20% week over week	Pause and reallocate the budget

It is what you build instead of two hundred dashboards. Most organizations have never written their important recurring decisions down in one place, which is why the sprawl has no center of gravity.

**Technical or adaptive.** Before you build, ask what kind of decision you are serving. Technical, operational, repeatable, bounded, like the save play? The Line can close the loop. Adaptive, a strategic pivot, a reorganization, a real trade between values? Then data informs and will not decide, and pretending a dashboard can is the original mistake. Design well for the technical ones. Structure the debate for the adaptive ones. Do not confuse them.

One caution on the Line itself. The predictive-and-actionable bar is high, and not every decision that matters has a signal that clears it. The risk is that you drift toward the decisions that are easy to instrument and away from the ones that count. A decision on the Register with no good leading indicator does not get cut. It gets a human owner, a review cadence, and honest judgment, which is the older technology and still the right one when the data cannot lead.

## What it's worth

Start with what you can count, and be honest about what kind of number it is.

Take a mid-market company with eight analysts and a BI estate of about 220 dashboards. Fully loaded, those analysts cost roughly \$1 million a year. If they spend 40% of their time building and maintaining reports, that is about \$400,000 a year going into report production. If a large share of those reports sit unopened, and your access logs will tell you the real share, a good chunk of that \$400,000 is producing analysis that changes no decision.

Be careful what you claim here. That money is mostly already spent, sunk in the build. Deleting a dead dashboard does not send cash back to the P&L, and your eight analysts do not get cheaper. What you reclaim is their capacity, and the value depends entirely on what they do with it. The honest move is to point that reclaimed time straight at the work that shifts decisions: writing the Register, and finding the leading indicators that clear the bar. The waste number gets you the meeting. The reallocation is the point.

For the upside, price one decision instead of waving at the portfolio. Say a save play, opened two weeks earlier because the leading indicator fired, retains an account worth \$40,000 a year, and across a book you catch three such accounts a quarter you would otherwise have lost. That is about \$120,000 a quarter in retained revenue, from one Decision Line. The figure is illustrative. The method is not. Pick your highest-value recurring decision, put a number on a single good call, and multiply by how often it recurs.

We are not claiming the third thing, and you should distrust anyone who does. Decision-Backward Design produces more decisions actually made per cycle. It does not make each one correct, and a dollar figure on decision quality is the same inflated math that built the graveyard.

## Where to start

You do not need a platform or a reorganization. You need a week and the willingness to delete things.

1. **Audit the estate.** Pull access logs. Find every dashboard untouched in 90 days. That list is your graveyard, and its length is your opening case.
2. **Write the register.** List the ten or so recurring decisions that actually move the business. Most teams have never done this. It takes an afternoon, and it is clarifying.
3. **Run three through the Line.** Take three of those decisions and answer the five questions. Say the owner's name out loud. If you cannot, you have found your real problem, and you found it in week one.
4. **Kill the zombies.** Retire the dashboards that serve no decision. This frees the analyst time, and it is the part everyone resists.
5. **Then automate, carefully.** Once a decision has an owner, a validated leading indicator, and a trigger, it earns an automated brief that lands in front of the owner the moment the trigger trips.

Start with the technical decisions, where the loop closes. Earn trust there before you point any of this at the adaptive ones.

## What to watch

A few honest cautions, because the framework has failure modes of its own.

**Keep an exploration budget.** “Cut every metric without a decision” has a trap: the open scan that finds a decision you did not know you had. Protect a small, explicitly labeled exploration budget with an owner and a cadence, and cap it so it cannot grow

back into a graveyard. Most genuinely strategic questions are themselves slow recurring decisions, where to put capital, where to compete, and they belong on the Register at the top. Kill the metrics that are neither a known decision nor deliberate exploration, and only those. The old name for the balance is exploration versus exploitation (March, 1991).

**You still keep the data layer.** This is not an argument against dashboards as infrastructure. The save-play trigger needs the underlying usage data to compute. What you cut is the presentation sprawl and the unowned reports, not the pipelines. An automated brief is just a dashboard that earned its place by passing the Line.

**Automation runs straight into algorithm aversion.** The same research that says people abandon erring models fastest on the biggest decisions is a warning about the endgame, where a system writes the Move. The mitigation is structural: the human owns the decision, the machine only surfaces the trigger, and every number traces back to its source. The moment you let AI draft the recommendation, that traceability stops being optional, or you have swapped a slow honest process for a fast plausible one. That discipline is its own subject, and the next paper in this set.

**The owner is the real barrier, and the register can become shelfware.** I have watched a register fill with names that could not actually say no, owners with the title and not the authority, and it quietly died. Naming an owner is uncomfortable because it removes the safety of shared blame. If your organization will not give an owner real decision rights, no framework saves you. You should at least know that is the thing in your way.

## **The point was never the dashboard**

A dashboard answers a question nobody asked and waits to be admired. A decision has a person behind it, a moment it has to happen, and a cost when it does not. The whole

job of analytics is to get from the first to the second, and nearly all the difficulty lives in that last step, which is the step most of our spending skips.

So build fewer things. Put a name on each one. Track what you can still change, not just what already happened. The dashboard was never the point. The decision was, and the person willing to own it.

---

## Sources

- Davenport, T. and Bean, R. “Action and Inaction on Data, Analytics, and AI.” MIT Sloan Management Review, 2023. <https://sloanreview.mit.edu/article/action-and-inaction-on-data-analytics-and-ai/>
- Wavestone. 2025 AI and Data Leadership Executive Benchmark (via DataIQ). <https://www.dataiq.global/articles/2025-ai-and-data-leadership/>
- Pfeffer, J. and Sutton, R. *The Knowing-Doing Gap*. Harvard Business School Press, 2000.
- Tetlock, P. “Accountability amplifies the status quo effect when change creates victims.” *Journal of Behavioral Decision Making*, 1994. <https://onlinelibrary.wiley.com/doi/abs/10.1002/bdm.3960070102>
- “An Integrative Perspective on Algorithm Aversion and Appreciation in Decision-Making.” *MIS Quarterly*, 2024. <https://misq.umn.edu/misq/article/48/4/1575/2300/An-Integrative-Perspective-on-Algorithm-Aversion>
- Szukits, A. and Moricz, P. “Towards data-driven decision making: the role of analytical culture and centralization efforts.” *Review of Managerial Science*, 2024. <https://link.springer.com/article/10.1007/s11846-023-00694-1>
- Heifetz, R. *Leadership Without Easy Answers*. Harvard University Press, 1994.

- Kozyrkov, C. On decision intelligence (Google's first Chief Decision Scientist). <https://www.datacamp.com/podcast/making-better-decisions-using-data-and-ai-with-cassie-kozyrkov-googles-first-chief-decision-scientist>
- March, J. G. "Exploration and Exploitation in Organizational Learning." *Organization Science*, 2(1), 71-87, 1991. <https://www.jstor.org/stable/2634940>