

From dashboard to decision

Martijn van Grieken

September 2026

CONTENTS

Executive summary	
1. The misunderstanding: data-driven working as a data problem	
2. The decision chain: five links	
Where the chain does not break	
3. Where the chain does break: two places	
Break 1: between meaning and decision	
Break 2: between execution and feedback	
Three situations in which you will recognise it	
How we ran into it ourselves	
4. Four stages of decision maturity	
5. Five measures for the chain	
6. Why more data does not repair the chain	
7. What AI changes here, and what it does not	
8. What you can do tomorrow	
9. What we do and do not cover	
About the author	
About Twentynext	
Accountability	
References	

Executive summary

Ask yourself one question, and take the time to answer it honestly. Which decision in your organisation was demonstrably taken differently this past month because data was available? Not: which decision was justified with figures afterwards. Not: which report was read. Which choice came out differently from the way it would have gone without that data.

In the conversations we have with boards and management teams, that question is usually followed by a pause. Not because there is no data; there is plenty. Not because there are no reports; there are too many. The pause falls because the question assumes a connection that is often missing in practice: the one between the figure and the choice.

That is what this whitepaper is about. In most organisations, data-driven working is approached as a data problem: a platform, a governance model, a dashboard environment, a team of specialists. Those things are necessary. They are simply not the problem. The problem sits one link further along, at the point where information should become a decision and does not. An organisation that leaves that alone can multiply its data foundation tenfold and see little of it in the speed and the quality of the decisions it takes.

We map that point with a chain of five links: signal, meaning, decision, execution and feedback. The chain rarely breaks at the front. It breaks in two places further along: between meaning and decision, where the information reaches someone who is not allowed to choose, and between execution and feedback, where nobody measures whether the choice worked. The first break makes your data powerless. The second makes your organisation forgetful: every decision is a first decision.

We then translate that into something you can use. Four stages of decision maturity, described in observable behaviour rather than a score, so that you can recognise your organisation without a questionnaire. Five measures for the chain, four of which most organisations never count. And five things you can do tomorrow, starting with the question above and with a test that costs nothing: switch off a report and see who complains.

One warning up front, because it runs through the whole piece. More data does not usually repair a broken chain, and often it makes matters worse. Every report added raises the price of choosing: more to read, more to weigh, more ways to postpone a decision until the figures are complete. They are never complete.

1. The misunderstanding: data-driven working as a data problem

Ask any organisation what it is doing to become data-driven, and the answer is about infrastructure. A data platform arrives. Sources are unlocked. A definition list is drawn up so that everyone means the same thing by “active customer”. A dashboard environment follows, and then a governance model to settle who may change which figure.

Each of those steps is defensible. Together they form a programme that can run for years and whose progress is excellent to report: number of sources unlocked, number of reports published, coverage of the definition list. It is work with a clear beginning and a clear end, and it produces artefacts you can show.

That is precisely where the trap lies. A programme that judges itself on artefacts reaches its goal the moment the artefacts exist. Whether the organisation started taking different decisions falls outside the scope, and so the question is not asked. That produces the situation we regularly find: a technically successful data programme inside an organisation that steers exactly as it did five years ago.

The misunderstanding is understandable. It is baked into the way we talk about information. The familiar ladder from data to information to knowledge to wisdom, described by Ackoff (1989), suggests that the steps follow one another by themselves: lay out enough data and insight appears, and from insight a better decision follows. That last arrow is the weakest in the whole sequence, and in practice it is the only one that really matters.

Why is that arrow so weak? Because decisions are not taken by organisations but by people, in meetings, under time pressure, with interests that are not all on the table. Simon (1947) described long ago that administrators do not optimise but settle for the first solution that will do. March (1994) showed that decision-making in organisations is often a distribution of attention rather than a weighing of alternatives. And Kahneman

(2011) mapped how systematically intuition departs from what the figures say, and how rarely people notice it in themselves. Kahneman et al. (2021) added a second source of error that hits at least as hard here: alongside bias there is noise, the spread between judgements that should have been identical. Two managers who assess the same file differently do not each have an opinion; together they have a measurement problem.

None of this is an argument against data. It explains why data alone is not enough. A figure changes a decision only when it reaches the right person at the right moment, in a form that forces a choice, and when that person has the mandate to choose differently from yesterday. Any of those four conditions can fail, and usually more than one does.

2. The decision chain: five links

To see where things go wrong, it helps to cut the route from figure to choice into pieces. We use five links.

Link	What happens	Who does it
1. Signal	Something changes in reality and is recorded	The system, the measurement, the registration
2. Meaning	The signal is interpreted: is this a deviation, and from what	The analyst, the report, the norm
3. Decision	Someone with a mandate chooses between courses of action	The decision-maker
4. Execution	The choice is brought into reality	The team, the process, the system
5. Feedback	The effect of the choice is measured and assessed	The organisation, if it remembers

The chain is a circle, not a line. Link five feeds link one: what you measure about the effect of a decision is itself a signal. Organisations that close this circle get a little better at choosing with every turn. Organisations that do not repeat their decisions without knowing whether they worked.

The value of this division lies in the precision it forces. “We want to work data-driven” is not an assignment anyone can act on. “We receive the signal about stock deviations eleven days after it arises, and only then does it reach someone allowed to decide” is.

WHERE THE CHAIN DOES NOT BREAK

It is worth establishing first where the problem usually is not, because that is where most of the money goes.

Link one is the easiest to delimit in many organisations. A great deal is recorded, often more than needed, and where it is missing it is usually a technical question with a predictable solution: a connection, a sensor, an adjusted form. That does not always make it light work: recording behaviour and data quality in old systems can be stubborn. But it is work whose nature is known.

Link two is a craft, and organisations that invest in it see results. Analysts get better, definitions get sharper, reports get more readable. Here too the progress is visible and the investment explicable.

Precisely because the first two links are manageable, they attract the attention and the budget. The result is an organisation that records and interprets excellently, and then does nothing differently from before.

3. Where the chain does break: two places

BREAK 1: BETWEEN MEANING AND DECISION

The first break is the best known and at the same time the most underestimated. The information is there, it is correct, it is comprehensible, and it reaches someone who may do nothing with it.

This rarely happens on purpose. It is the sum of small, defensible choices. The dashboard was built at the request of the department that knows the data best, so it sits with the analysts. The report goes to the management team, which meets monthly, while the signal changes weekly. The employee who sees the deviation first may not adjust the order without approval from someone who does not open the dashboard.

We call this the last metre. The data has travelled kilometres, through connections, transformations and visualisations, and strands in the last metre between screen and decision. Three forms occur most often.

- **The mandate gap.** Whoever sees it may not choose; whoever may choose does not see it. This is the purest form and can usually be established in a single conversation.
- **The rhythm gap.** The signal moves faster than the decision moment. A weekly deviation discussed monthly is no longer a deviation by the time it is discussed, but history.
- **The form gap.** The information is presented as a picture of reality rather than as a choice. A chart that falls raises the question of why it falls. A chart with a threshold line and two courses of action raises the question of what you are going to do.

That third point deserves emphasis, because it is the cheapest to repair and the most often overlooked. Most reports are built descriptively: they show what is. A report that feeds decisions is built normatively: it shows what is against what should have been, and makes visible which action belongs to that. That difference is not a matter of technology but of design, and it is usually faster to achieve than a new data project, provided the owner, the norm and the courses of action are clear.

BREAK 2: BETWEEN EXECUTION AND FEEDBACK

The second break is quieter and more damaging in the long run. A decision is taken, it is carried out, and after that nobody looks back.

This is not negligence. It is the natural outcome of how organisations divide their attention. A decision taken is off the agenda; the next matter is already waiting. Looking back also demands something uncomfortable: the willingness to establish that a choice was wrong, in an environment where decisions attach to people.

The consequences are twofold. First, the organisation does not learn. Without a measured effect, every decision remains a first decision, and experience remains a collection of anecdotes rather than a growing capability. Second, and this is subtler, the data itself loses its edge. Models that predict without ever being held to what actually happened drift away from reality unnoticed. Nobody notices, because nobody looks.

The test for this break is simple. Name a decision your organisation reversed in the past year because the measurement showed it was not working. If you cannot name one, then one of two things is true: all your decisions were good, or you do not measure them.

THREE SITUATIONS IN WHICH YOU WILL RECOGNISE IT

Described in the abstract, this pattern is the easiest to deny. Three situations we keep seeing in various guises.

The stock that turns too late. An organisation measures its stock levels daily and has a well-working forecasting model. The deviation becomes visible on Monday. The purchasing decision falls in the weekly meeting on Thursday, and the order goes out on Friday. Nobody does anything wrong, and yet the time from signal to action is four days in a market that turns in two. The solution does not lie in a better model but in a decision mandate with the planner, with a threshold above which they may act themselves.

The dashboard that raises questions instead of choices. A management team receives a monthly overview with twelve charts. Every month the discussion turns on the same two questions: is this figure right, and how come that line is falling. A request for further analysis follows. The next month the analysis is there, the situation has changed, and it starts again. This is the form gap in its purest form: the overview describes reality excellently and forces no choice at all.

The project nobody dares to stop. An initiative has been running for eighteen months. It was never established what it would be measured against, so there is no moment at which it demonstrably fails. The conversation is about effort and confidence rather than result, and the decision to stop would be a personal judgement rather than an agreed outcome. This is the absence of stated grounds for refutation, and it costs more money than all the dashboards put together.

HOW WE RAN INTO IT OURSELVES

Finally, an example from our own house, because the pattern affects us as much as our clients. Every week we measure how our own website is found and what comes of it, in a fixed set-up with recorded definitions, so that progress cannot be confused with a different delimitation.

That set-up initially held only a count over twenty-eight days. Such a window hides a standstill. A period without a single enquiry looks in that count like a slightly lower average, not like a stop, and a lower average raises at most the question of whether the season is playing a part. The figure was there, the meaning was not, and so no decision followed.

The repair cost no new data and no new dashboard: next to the count over twenty-eight days there is now the date of the last actual enquiry. One field, and the question of whether it is still running can be answered at a glance rather than after an analysis. That is the form gap from the previous section, in our own house, and the lesson is the same: the problem was not in the measurement but in what it showed a reader.

4. Four stages of decision maturity

Maturity models have a bad name, and rightly so. They produce a score that feels comfortable and changes nothing. We therefore describe the four stages in observable behaviour: things you can watch happen in an organisation, without a questionnaire.

The best-known model in this field is that of Davenport and Harris (2007), who described five stages of analytical maturity, rising from organisations with scattered individual analyses to organisations that build their competitive position on them. That model looks at the analytical capability of an organisation. Ours looks at something else, and therefore counts differently: not how well you analyse, but whether the outcome of that analysis changes a decision. An organisation can score highly in their model and remain stuck in stage two of ours.

Stage	How you recognise it	What is needed next
1. Collecting	Figures are requested after the decision, to support what has already been decided. Discussions are about whether the figure is correct, not about the choice.	One process in which the figure is on the table before the decision moment.
2. Reporting	There are many reports and they are looked at. Switch one off and nobody complains. Nobody can name a decision that followed from it.	An owner per report who may also decide, and a threshold that prompts an action.
3. Deciding	People can name concrete choices that fell differently because of data. The justification can still be found later.	Recording in advance what would reverse the decision, and testing that afterwards.
4. Learning	Decisions are measured afterwards, sometimes reversed, and the outcome changes the next round. Looking back is routine and not personal.	Guarding that the rhythm survives when things get busy.

Two remarks about this model, because without them it works against you.

First, it applies per process, not per organisation. Almost every organisation we see sits in several stages at once: financial steering in stage three, stock management in stage one, customer service somewhere in between. An organisation-wide score averages that away and hides precisely the place where there is ground to gain.

Second, stage four is not an end point but a state that needs maintenance. Looking back is the first thing to go once the agenda fills up. Organisations that reach this stage and do nothing to hold it slip back to three within a year.

5. Five measures for the chain

Anyone who wants to improve the chain has to be able to measure it. The five measures below are deliberately taken not from the data world but from the decision world. Four of them are counted nowhere in most organisations, and that is exactly what makes them worth having.

- **Time from signal to decision.** The number of days between the moment reality changed and the moment someone with a mandate chose in response. Measure this on one process, not on all of them. The number is usually a multiple of what those involved expect.
- **Share of decisions with a traceable justification.** Of the decisions of the past quarter: for how many can you still find which figure or which argument was decisive? Not whether a justification existed, but whether it can still be found now.
- **Share of reports actually in use.** Of all the reports you maintain: how many are actually opened, and by how many different people? The tail is almost always longer than expected, and that tail costs maintenance, attention and trust.
- **Feedback frequency.** Of the decisions with a measurable effect: for how many was that effect actually measured, and within what period? This is the measure that makes the second break visible.
- **Grounds for refutation, stated in advance.** For how many decisions was it recorded beforehand at which outcome you would conclude that the decision was wrong? This is the sharpest of the five and the least common. A decision without a refutation criterion cannot turn out to be wrong, and therefore cannot teach anything either.

That last measure deserves explanation, because it often meets resistance. Recording in advance what would make you turn back feels like inviting criticism. In practice it works the other way round: it lowers the stakes of the decision. Whoever says beforehand “we will evaluate this in the second quarter and stop below this result” does not have to win in order to stay credible. That makes it easier to take a decision, and much easier to reverse one.

6. Why more data does not repair the chain

The most common response to faltering decision-making is a request for more information. An additional analysis arrives, an extra dashboard, a deeper study. That feels like progress and rarely is.

There are three reasons why more data tends to aggravate a broken chain rather than heal it.

The price of choosing rises. Every extra report enlarges the amount a decision-maker is assumed to have weighed. That raises the threshold for choosing and supplies an excellent argument for waiting: the figures are not complete yet. They never are.

Attention fragments. Attention is scarce in every organisation, scarcer than money. A report nobody opens is not free: it demands maintenance, it appears in overviews, and at every disruption someone has to establish whether it matters. The twentieth report makes the other nineteen less visible.

It displaces the question of blame. As long as the answer to a faltering process is “we need more insight”, nobody has to establish that the mandate is in the wrong place or that the decision rhythm does not match the problem. That is comfortable, and it is why this pattern can repeat itself for years.

The way out is not less data but a different order. Start with the decision you want to see fall differently, work back to the signal that requires, and build only that. Whatever remains of the existing reports may prove that it is missed.

7. What AI changes here, and what it does not

This whitepaper is about decision-making, and the most frequently asked question of the past year is what AI changes about it. The answer is sharper than expected: AI improves the link that already worked best, and leaves both breaks untouched.

AI is usually deployed first in link two: interpreting, summarising, connecting, forecasting. That goes faster and often better than a person with the same time can manage, and it is real gain; we build those solutions with conviction. AI can help elsewhere in the chain too, in preparing a decision, in execution and in watching the effect. What it nowhere solves by itself is who may choose and whether anyone looks back.

There is also an effect running the other way. AI sharply lowers the cost of a first analysis. Where an analyst used to need a week, a first version is there within an hour. The costs of validation, context and interpretation remain; those of producing all but disappear. That means the supply of signals grows faster than the capacity to decide on

them. Precisely the mechanism from the previous chapter, now accelerating: the price of choosing rises, attention fragments further, and the argument that one more analysis is needed becomes cheaper than ever.

A second point is often missed in the boardroom. A model that forecasts without ever being held to the outcome drifts. With a human analyst you notice that: they lose credibility. With a model it happens silently, because it keeps producing figures with the same confidence. The second break, the absence of feedback, is therefore no longer a learning problem with AI but a quality problem.

The practical conclusion is not to be cautious with AI. It is that the order matters. Repair the mandate and the rhythm on one process first, and deploy AI on that same process afterwards. Then every hour gained in link two is genuinely an hour gained towards the decision. The other way round it is an investment in speed on a road that runs out further along.

8. What you can do tomorrow

Five steps, rising in effort. The first two cost nothing and are done this week.

- **Put the monthly question to your management team.** Which decision fell differently this past month because of data? Ask for one concrete example, with the name of the decision-maker. If it stays quiet, you have your baseline, and it is sharper than any maturity score.
- **Switch off one report and see who complains.** Pick one you suspect is no longer used, turn it off for two weeks and measure the reactions. This is the cheapest measurement of usage that exists, and the outcome is almost always instructive.
- **Give every dashboard an owner who may also decide.** Not the builder, not the analyst: the person who may take an action on the basis of this screen. If you cannot find that person, you know this screen feeds no decision.
- **Record the refutation in advance for your next decision.** One sentence will do: at which outcome, at which moment, will we conclude that this was wrong. Put that sentence in the decision document and schedule the moment straight away.
- **Measure the time from signal to decision on one process.** Pick a process with a clear owner and a measurable lead time. Record the baseline before you change anything, or in six months you will have an improvement you cannot demonstrate.

That last clause is not a formality. The most common reason improvement efforts around decision-making peter out is that nobody can show anything changed. A baseline measurement of half a day makes the difference between an effort that still exists next year and one that gets cut.

9. What we do and do not cover

This whitepaper describes a pattern we recognise from our work on data and AI solutions for organisations in healthcare, industry and the public sector. We build and maintain those solutions, and that is exactly why we often see what happens after delivery, and what does not.

What we do not do here is diagnose your decision-making from a distance. The four stages and the five measures are meant as language and as a yardstick, not as a verdict. Every organisation has its own reasons why the chain runs as it runs, and those reasons are rarely foolish.

Nor do we claim that technology does not matter. A decision waiting eleven days for a manual export is a technical problem with a technical solution. Our point is narrower than that: technology is necessary and not sufficient, and the organisations that get stuck almost never get stuck on the part they master best.

If you recognise the pattern and want to know where your own chain breaks, our free Data Maturity Scan will give you a first picture. If you would rather have the sharp version straight away, pick one process and put the question from the executive summary. We are glad to think along about what the answer means.

About the author

Martijn van Grieken is the founder of Twentynext. Since 2014 he has worked on data and AI solutions for organisations in healthcare, industry and the public sector, from the first question to working software that runs in production and is maintained.

About Twentynext

Twentynext is a data and AI consultancy in Eindhoven. We advise, build and maintain: from business question to working software, and then the maintenance needed to keep that software working. That combination is deliberate: advice, delivery and maintenance require coherence in our experience. Advice without delivery stays a report, and delivery without maintenance becomes a problem in time.

Accountability

The chain division and the four stages in this whitepaper are our own framework, formed in the work. They build on a tradition in the decision-making literature that reaches further back than the field of data and analytics: on Simon (1947), who described that administrators do not optimise but settle; on March (1994), who analysed decision-making in organisations chiefly as a distribution of attention; and on Kahneman (2011) and Kahneman et al. (2021) on the systematic ways judgements depart, through both bias and noise. For the distinction from existing maturity models we refer to Davenport and Harris (2007); the DIKW ladder that opens chapter 1 goes back to Ackoff (1989).

The references follow APA 7. We verified every source in this list on title, year, publisher and place of publication; where a publisher differs by market, we name the publisher of the first edition.

The examples in this piece are patterns, not cases. Where we write about client situations we do so only with permission and in the form agreed. Figures about our own practice we give only when they can be traced to a recorded measurement.

References

- Ackoff, R. L. (1989). From data to wisdom. *Journal of Applied Systems Analysis*, 16, 3–9.
- Davenport, T. H., & Harris, J. G. (2007). *Competing on analytics: The new science of winning*. Harvard Business School Press.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown Spark.
- March, J. G. (1994). *A primer on decision making: How decisions happen*. Free Press.
- Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organization*. Macmillan.

Colophon

How we make this. This whitepaper was produced with the help of AI. Twentynext is responsible for the content, but cannot guarantee that it is accurate or complete.

Disclaimer. This whitepaper is general in nature and not advice tailored to your situation. Read our full disclaimer at [twentynext.nl/en/website-disclaimer](https://www.twentynext.nl/en/website-disclaimer).

Publisher and editorially responsible: Twentynext B.V., Eindhoven, the Netherlands · info@twentynext.nl · [twentynext.nl](https://www.twentynext.nl)