


BEYOND THE DASHBOARD

The Data You Don't Have

Why large-scale change keeps stalling in the space between the data it has and the data it wishes it had



Giovanni Leonardi

July 2007

Contents

1. Executive Summary	3
2. The Meeting That Waits for a Number	3
3. Risk Is Not the Same as Uncertainty	4
4. Why Organisations Wait: The Structural Forces	5
4.1 The Compliance Reflex	5
4.2 The Warehouse Promise	5
4.3 The Asymmetry of Blame	5
5. The Seduction of the Single Version of the Truth	6
6. The Honest Objection	6
7. What Judgement Under Uncertainty Actually Looks Like	7
8. The Gap Between Intent and Machinery	8
9. A Different Temperament	9

The Data You Don't Have

Transformation · No. 144

Originally written July 2007 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

*“A precise answer to an unanswerable question is not knowledge.
It is decoration.”*

1. Executive Summary

There is a moment, familiar to anyone who has sat on the steering committee of a large change programme, when the decision on the table is plainly bigger than the evidence available to make it. The instinct, almost universally, is to send the analysts away for more. More modelling, more benchmarking, a firmer set of numbers. It feels responsible. It is often the opposite.

This essay is about the decisions transformation cannot make well because it insists on treating them as problems of missing data — when they are in fact problems of irreducible uncertainty. The distinction, drawn eighty years ago by an economist and largely forgotten by practitioners, is between risk, which can be measured, and uncertainty, which cannot. Much of the difficulty in large-scale change lives in the second category and is quietly, expensively, managed as though it belonged to the first.

The pattern is not a failure of individual judgement. It is structural. Three forces — the compliance reflex that now governs how large organisations account for themselves, the promise implied by years of investment in data warehouses and management information, and the deep asymmetry between the blame attached to a considered wrong call and the blame attached to a wrong call made without full analysis — combine to make waiting for data the rational choice for the individual and the wrong choice for the enterprise.

The argument here is not against data. It is against the confusion of two different things: reducing the uncertainty you can, and pretending you can reduce the uncertainty you cannot. Transformation stalls in the gap between them. Closing it is less a matter of method than of temperament — and temperament, unlike method, cannot be bought in.

2. The Meeting That Waits for a Number

Consider a scene that repeats itself, with minor variations, in boardrooms and programme offices everywhere. A consolidation programme — several business units, incompatible systems, a single target platform — has reached its investment decision. The business case runs to sixty million pounds over five years. The benefits, as benefits always do, depend on things that have not happened yet: how quickly people adopt the new processes, how much of the promised standardisation survives contact with local exception, how many of the old systems can actually be switched off rather than quietly kept alive.

The steering committee reads the paper. Someone — usually the most senior person in the room, and often the most careful — observes that the adoption assumptions are “soft.” They are. No one can know, in advance, how a workforce of nine thousand people will respond to a change none of them asked for. The decision is deferred. The analysts are

sent to firm up the numbers. A quarter passes. The revised paper returns with a more elaborate model, a wider set of benchmarks, a sensitivity analysis. The assumptions are no firmer, because they cannot be; they have simply been dressed more warmly. The decision is deferred again.

By the time the programme is finally approved, nine months have gone, several million pounds of run-rate have been spent keeping the condemned systems breathing, and the market the programme was meant to address has moved. The delay was not caused by disagreement about the answer. Everyone in the room, if pressed privately, would have given roughly the same one at the first meeting. The delay was caused by the belief that a number existed which, once found, would make the judgement unnecessary.

It did not exist. It was never going to exist. And the most revealing part of the story is what the eventual sensitivity analysis showed: across the entire plausible range of the adoption assumptions — from pessimistic to optimistic — the decision did not change. The programme was worth doing at the low end and worth doing at the high end. The data everyone waited for would not have altered the outcome. It would only have altered how defensible the outcome felt.

3. Risk Is Not the Same as Uncertainty

In 1921 the economist Frank Knight drew a line that our management vocabulary has never properly absorbed. Risk, he said, is a situation in which the outcome is unknown but the distribution of outcomes is knowable — the roll of a die, the failure rate of a component, the default rate of a large book of similar loans. You cannot know the next result, but you can know the odds, and knowing the odds you can price, hedge, and plan. Uncertainty, by contrast, is the situation in which the distribution itself is unknown or unknowable — where the event is genuinely novel, where the sample size is one, where the future does not resemble the past closely enough for the past to instruct it.

The whole apparatus of modern decision support — the models, the dashboards, the fact-based rigour that Six Sigma and its descendants have trained a generation to demand — is built for risk. It is superb at risk. Give it a stable process and enough history and it will tell you, with real authority, what to expect and where to intervene. The trouble is that it presents the same confident face when handed a question of uncertainty, and the face is a lie. A precise answer to an unanswerable question is not knowledge. It is decoration.

The failure at the heart of so much stalled transformation is not that organisations lack data. It is that they apply the machinery of risk — measurement, distribution, expected value — to conditions of uncertainty, where that machinery cannot work, and then mistake the resulting precision for insight.

The data you don't have, then, comes in two very different kinds. There is the data that exists but has not yet been gathered — last year's actuals still sitting in a subsidiary's ledgers, the process timings no one has measured, the customer information scattered across systems that do not speak to each other. That data is worth chasing, and chasing it is ordinary diligence. And there is the data that does not exist and cannot, because it concerns a future that has not been written: how a market will respond, whether a competitor will move, how people will behave when their working lives are rearranged. Waiting for the first kind is prudent. Waiting for the second is a category error dressed as prudence, and it is the more dangerous precisely because it wears the same clothes.

4. Why Organisations Wait: The Structural Forces

If this were merely a matter of individuals failing to see a distinction, it would be easily fixed by explaining the distinction. It is not easily fixed, which tells us the causes are structural. Three forces, in particular, make the counsel of delay almost irresistible inside a large organisation.

4.1 The Compliance Reflex

The last five years have reshaped how large enterprises account for their own decisions. In the wake of the corporate scandals at the start of the decade, the regulatory settlement — most visibly the requirements now placed on financial reporting and internal control, and, for those in financial services, the capital and risk-modelling regime being bedded in this year — has installed a reflex at the top of every substantial organisation: the decision must be documented, traceable, and defensible after the fact. This is, in its own domain, entirely proper. Money should be handled by people who can show their working.

But a reflex trained on one domain does not stay in it. The habit of defensibility migrates from the accounts to every consequential choice, and in the domain of genuine uncertainty it becomes corrosive. A judgement, honestly made under uncertainty, cannot be fully documented, because its essential ingredient — the practitioner's read of a situation the data cannot settle — leaves no audit trail. So the organisation reaches for the thing that can be documented: more analysis. The analysis is not sought because it will improve the decision. It is sought because it will improve the file.

4.2 The Warehouse Promise

For the better part of a decade, large organisations have poured money into the infrastructure of knowing. Data warehouses, business intelligence layers, the long and painful pursuit of a "single version of the truth," the executive dashboard that promises the whole enterprise on one screen. These are real achievements and they have made much that was once invisible visible.

They have also made an implicit promise that no infrastructure can keep: that if the picture is only complete enough, the decision will become obvious. When an organisation has spent years and fortunes on the premise that everything important can be measured, the admission that some important things cannot be measured feels like a betrayal of the investment. The dashboard glows green; surely, somewhere in all that instrumentation, the answer is waiting to be found. The more comprehensive the reporting, the harder it becomes to say the honest thing — that the number which would settle this particular question is not in the warehouse, and never will be, because it concerns something that has not happened.

4.3 The Asymmetry of Blame

The third force is the most human and the most powerful. Consider the position of the individual executive facing an uncertain call. If they decide on judgement and are wrong, the post-mortem is brutal: you had a gap in the analysis and you proceeded anyway. If they wait, commission more work, and are wrong, the post-mortem is gentle: you did everything reasonable; the information simply wasn't available. The outcomes may be identical — the same failed programme, the same wasted money — but the personal consequences are not remotely symmetric.

Delay is rarely the organisation's rational choice, but it is almost always the individual's. As long as a wrong decision made carefully is punished more lightly than a wrong decision made boldly, careful people will keep choosing to wait – and the sum of their careful waiting is an enterprise that cannot move.

This asymmetry is not written into any policy. It lives in the culture, in the tone of the review meeting, in whose career survives a failure and whose does not. And because it is unwritten, it is almost never addressed directly. Organisations exhort their leaders to be bold while maintaining, in every mechanism that actually distributes reward and blame, a standing incentive to be cautious. The exhortation is a poster. The asymmetry is the weather.

5. The Seduction of the Single Version of the Truth

There is a particular danger in the models that do get built for genuinely uncertain questions, and it deserves naming, because it is the point at which waiting-for-data turns from wasteful into actively harmful.

A model of an uncertain future must assume something about how that future behaves, and the convenient, tractable, universally used assumption is that it behaves like the recent past – that the distribution which held yesterday will hold tomorrow. For long stretches this assumption is invisible, because for long stretches it is roughly true. The danger is that the periods when it fails are precisely the periods that matter most: the discontinuities, the regime changes, the events for which there is no useful precedent. A model built entirely on the past is, by construction, blind to exactly the events that would most damage the decision it informs.

An argument much discussed this year has put a memorable name to the phenomenon – the improbable, high-consequence event that our history-based models are structurally unable to see coming, and that we rationalise, afterwards, as though it had been predictable all along. Whatever one makes of the wider thesis, the warning for the transformation practitioner is exact. The elaborate model produced to justify a delayed decision is not neutral. It does not merely fail to reduce the uncertainty; it disguises it. It converts a range of genuine possibilities into a single confident line, and a single confident line is far easier to act on badly than an honest admission of doubt. The organisation that waited nine months for a firmer number often ends up more exposed than the one that decided in the first week, because the first-week decision knew it was a judgement and hedged accordingly, while the nine-month decision believed it had been proven.

“Precision is not the same as accuracy, and a model's confidence is a property of the model, not of the world it claims to describe.”

6. The Honest Objection

It would be easy, and wrong, to read all this as an argument against analysis — a licence for the confident executive to trust their gut and dismiss the analysts as obstacles. The strongest case against the position taken here must be met squarely, because it is a good case.

More information, the objection runs, is genuinely and almost always better. Much of what looks like irreducible uncertainty is nothing of the kind; it is reducible uncertainty that the impatient decision-maker cannot be bothered to reduce. The adoption rate no one can "know" in advance can in fact be estimated far better by running a genuine pilot than by guessing. The competitor's move can be anticipated by better market intelligence. To wave the flag of "genuine uncertainty" over every hard question is to give laziness and bias a respectable name, and the history of decisions made boldly on insufficient evidence is not a happy one.

All of this is true, and it sharpens rather than defeats the argument. The discipline being urged here is not decide without data. It is distinguish, before you spend a quarter chasing it, which kind of uncertainty you face. The reducible kind should be reduced, and vigorously — a pilot that costs a month and genuinely narrows the range is worth ten steering committees. The test is not whether more analysis is possible; more analysis is always possible. The test is whether the analysis, once done, would actually change the decision. Where it would, commission it without hesitation. Where it would not — where, as in the consolidation programme, the answer is stable across the whole plausible range of the missing numbers — further analysis is not diligence. It is displacement activity, and expensive displacement at that.

7. What Judgement Under Uncertainty Actually Looks Like

If the answer is not to wait, and not to guess recklessly, what is the practitioner actually meant to do? The tradition of thinking about decisions under uncertainty is older and richer than our dashboards, and it offers several disciplines that the data-driven era has let fall into disuse.

The first is to ask the question the elaborate model obscures: what would I have to believe for this to be the wrong decision? This inverts the whole exercise. Instead of assembling evidence toward a confident point estimate, it names the specific conditions under which the choice fails, and asks how plausible those conditions are and whether they can be watched for. A decision whose failure requires three unlikely things to happen at once is a different animal from one that fails if a single ordinary thing goes wrong, and no amount of point-estimate modelling reveals the difference as clearly as this single question.

The second is to prize reversibility over confidence. The reason the search for certainty is so costly is that it treats every decision as though it were irreversible and therefore had to be right first time. Most are not. A choice that can be unwound, staged, or run as a genuine trial before full commitment can be made under far more uncertainty than one that cannot, because the cost of being wrong is bounded. The practitioner's craft lies in structuring the decision so that it becomes reversible — sequencing the programme so the uncertain, expensive, hard-to-undo commitments come after the cheap, informative, early ones, not before. This is not a new idea; those who think in terms of options have argued it for years. It remains widely ignored, because staging a decision to preserve the freedom to change your mind reads, in a defensibility culture, as indecision.

The third is the oldest of all, and the most resisted: the acceptance that a good-enough decision made in time will usually beat a perfect decision made too late. The economist

Herbert Simon called this satisficing — choosing the first option that clears the bar of good enough rather than searching exhaustively for the optimum — and demonstrated, decades ago, that it is not a lazy compromise but the only rational strategy available to any decision-maker with finite time and attention facing a world too complex to optimise. In transformation, where the ground shifts while you analyse it, the optimum is a moving target, and the exhaustive search for it guarantees you will aim at where the target was.

Reducible uncertainty	Irreducible uncertainty
The data exists but has not been gathered	The data concerns a future that has not happened
Closed by diligence: gather, measure, pilot	Cannot be closed by any amount of analysis
More work genuinely improves the decision	More work improves only the file, not the decision
Correct response: reduce it, vigorously	Correct response: decide under it, and stay reversible

None of these disciplines is technical. Not one requires a system, a licence, or a consultant. What they require is a tolerance for acting on incomplete information without pretending the information is complete — and that tolerance is a matter of character before it is a matter of technique.

8. The Gap Between Intent and Machinery

Here we reach the structural story the topic asks for. Every organisation of any size now declares that it wants to be bold, adaptive, decisive — that it will not be left behind by the pace of change. The intent is genuine. And yet the same organisations have spent the decade building machinery whose every setting rewards the opposite behaviour: a compliance apparatus that prizes the defensible file over the sound judgement, an information infrastructure that has promised omniscience and made the admission of ignorance feel like failure, and a reward system that punishes the bold wrong call far more harshly than the cautious one.

The gap between transformation intent and transformation reality is not, at root, a gap in capability or in data. It is the gap between what the organisation says it values and what its mechanisms actually reward. Leaders who exhort their people to embrace uncertainty while running review meetings that flay anyone who acted without complete analysis are not being hypocritical, exactly; they are simply unaware that the meeting is a stronger signal than the speech. The machinery teaches, every day, in a hundred small verdicts, that the safe move is to wait for a number. And so people wait, and transformation — which is nothing but a long chain of consequential decisions made under uncertainty — slows to the speed of its most cautious gate.

To close the gap, the intervention has to be aimed at the machinery, not the exhortation. It means making it legitimate, in writing and in the tone of the review, to record that a decision was a judgement under irreducible uncertainty, and to be judged afterwards on the quality of the reasoning rather than the accuracy of the outcome — because under genuine uncertainty a good decision can have a bad result and a bad decision a good one, and an organisation that cannot tell these apart will learn exactly the wrong lessons from both. It means an executive prepared to say, on the record, we do not have this number, we will not have it, and here is the judgement we are making anyway — and to be visibly backed when they do. One such act, publicly supported, teaches more than a year of

posters about agility.

9. A Different Temperament

We are, as a profession, extraordinarily fluent in method. We have never been better equipped to measure, model, report, and analyse. What we are far less fluent in is temperament — the settled capacity to act well when the method runs out, when the data you don't have is the data that cannot exist, and the decision will not wait.

The paradox of the data-rich organisation is that its very richness can make it worse at exactly this. The more it can measure, the less tolerant it becomes of what it cannot, and the more it mistakes the comfort of a full dashboard for the confidence of a sound judgement. The discipline this moment asks for is almost countercultural: to invest seriously in reducing the uncertainty that can be reduced, and then to have the nerve to decide, openly and reversibly, in the presence of the uncertainty that cannot. The first half is a matter of tools, and we are good at tools. The second is a matter of character, and character is the one thing no warehouse holds and no analyst can be sent away to fetch.

The organisations that will move — and in a changing market, moving is survival — are not the ones with the most data. They are the ones that have learned the difference between the questions their data can answer and the questions it cannot, and have built the temperament, the reward structures, and the plain institutional courage to answer the second kind anyway. That, and not another quarter of analysis, is what the deferred decision has been waiting for all along.