

BEYOND DELIVERY

Measuring the Success of Data Products: Adoption, Trust and Decisions, Not Just Delivery

A Practitioner's White Paper on Portfolio Measurement for Data Products

Giovanni Leonardi

February 2023



Contents

1. Executive Summary

2. Necessary, But Not Sufficient

3. Six Dimensions of Data Product Success

3.1 Usage and Adoption

3.2 Data Quality

3.3 Reliability

3.4 Decision Impact

3.5 Business Value

3.6 Reduced Operational Friction

4. The Data Product Scorecard

5. Reading the Scorecard Without Gaming It

6. Embedding Measurement in the Lifecycle and Governance

7. Recommendation

3

4

4

4

5

5

5

5

6

6

7

8

8

Measuring the Success of Data Products: Adoption, Trust and Decisions, Not Just Delivery

Portfolio · No. 103

Originally written February 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“A data product can be delivered flawlessly against plan and be a complete failure in use.”

1. Executive Summary

For a generation, we have measured the success of data initiatives the way we measure any project: did it land on time, on budget, and in scope? Those three questions were always imperfect for data, and the shift from data projects to data products has made their inadequacy impossible to ignore. A data product is not delivered and finished; it is released and then lives, serving decisions day after day. Judging it by whether it was delivered on plan is like judging a restaurant by whether it opened on schedule and never asking whether anyone eats there.

This paper argues that milestones, budget, and scope are necessary but wholly insufficient measures of a data product, and that they actively mislead when used alone: a data product can be delivered flawlessly against plan and be a complete failure in use. In their place — or rather, around them — it proposes a balanced scorecard of six dimensions that together describe whether a data product is actually succeeding: usage and adoption, data quality, reliability, decision impact, business value, and reduced operational friction.

The paper sets out each dimension, provides a worked scorecard with example metrics and accountable owners, explains how to read the scorecard without gaming it, and concludes with a specific recommendation: that every data product carry this scorecard from inception, that measurement be built into the product lifecycle rather than bolted on at review time, and that portfolio governance shift its primary question from whether the product was delivered to whether it is being used, trusted, and turned into decisions and value.

The stakes are portfolio-wide. An organisation that funds data products on delivery metrics alone will systematically over-invest in building and under-invest in adoption, and will accumulate a portfolio of technically complete products that change nothing. The measure you choose is the behaviour you get.

2. Necessary, But Not Sufficient

Delivery metrics are not wrong; they are incomplete. On time, on budget, and in scope tell you whether a team executed a plan. For a bridge or an office fit-out, execution and success are nearly the same thing, because the asset delivers its value simply by existing. A data product is different. Its value is created not when it is built but when it is used, repeatedly, by people who trust it enough to change what they do. Delivery is the precondition for value, never the proof of it.

This is why delivery metrics mislead when used alone. They reach their maximum reading — one hundred per cent delivered — at the exact moment the product has produced no value whatsoever: the day it goes live. Everything that determines whether the investment pays off happens afterwards, in a period the delivery metrics do not measure at all.

“A portfolio steered by delivery status is steering by the one indicator guaranteed to look best when nothing has yet been proven.”

There is a second, subtler failure. Because delivery is the only thing measured, it becomes the only thing managed. Teams are rewarded for shipping, not for adoption; sponsors declare victory at go-live and move their attention elsewhere; and the unglamorous work that actually creates value — driving usage, earning trust, embedding the product into how decisions are made — has no owner and no metric, so it does not happen. This is the mechanism by which organisations accumulate shelfware with a data warehouse behind it: products that were delivered competently, reported as successes, and used by no one. The failure is invisible precisely because the only instrument on the dashboard is showing green.

A data product delivered on time, on budget, and in scope has merely cleared the starting line. Every measure that matters describes the race — and the race begins the day it goes live.

3. Six Dimensions of Data Product Success

What, then, should be measured? Success for a data product is not one number but a balance across six dimensions. No single dimension is sufficient, and a product can score well on some while failing on others in ways that matter. The point of a balanced scorecard is precisely to make those trade-offs visible rather than to collapse them into a single reassuring figure.

3.1 Usage and Adoption

The first and most basic question is whether anyone is actually using the product. A data product with no consumers has failed regardless of how well it was built. Adoption measures both depth — how intensively users engage — and breadth — how much of the intended audience has taken it up — and it distinguishes one-time curiosity from repeated, embedded use. The failure mode this dimension catches is the most common of all: a product that launched to fanfare, drew a wave of first visits, and then flatlined as the intended users quietly returned to the spreadsheet they trusted. Adoption that is not growing, or not sticking, is the earliest signal that value will never materialise.

3.2 Data Quality

Usage without trust is fragile and short-lived. Quality measures whether the product is accurate, complete, and current enough for the decisions it serves, against explicitly defined standards for its critical data elements rather than a generic completeness percentage. It is the dimension that most directly underwrites trust, and the first to destroy it when it slips. Trust in a data product is asymmetric: it is earned slowly over many correct answers and lost instantly on a single visible error, after which users revert to their own sources and rarely return. Measuring quality against the elements that actually matter, not the average across all fields, is what protects the trust the whole product depends on.

3.3 Reliability

A product that is trusted but unavailable, or late, cannot be relied upon in a live decision process. Reliability measures whether the product is there when it is needed, at the freshness it promises, and how quickly service is restored when it is not. Adoption depends on reliability in a compounding way: people will not build a product into a recurring decision or an automated process if they cannot count on it being current and available at the moment the decision is made. Every missed refresh teaches the user to keep a manual fallback — and a manual fallback is adoption quietly leaking away.

3.4 Decision Impact

This is the dimension delivery-based measurement omits entirely, and the one that most defines a data product. It asks whether the product actually changes decisions: whether it is embedded in named business processes, cited in real choices, and shortening the time it takes to make them. A product that is used, trusted, and reliable but changes no decision is a well-run library that no one acts on. Decision impact is harder to measure than usage, because it requires tracing the product into the processes it serves, but it is the dimension that separates a data product that informs the business from one that merely describes it.

3.5 Business Value

Ultimately a data product must move an outcome the organisation cares about, whether revenue, cost, risk, or service quality. Value measures the benefit realised against the case that justified the investment, and the value returned per unit of ongoing run cost. It is the hardest dimension to attribute cleanly, because business outcomes have many causes and a data product is rarely the only one, and it is the most important not to abandon for that reason. A disciplined, honestly caveated estimate of value that keeps the question alive is

worth more than a precise measurement of delivery that answers a question no one should be asking.

3.6 Reduced Operational Friction

Many data products earn their keep less by enabling new decisions than by making the organisation's existing work easier: retiring duplicated and shadow data, eliminating manual reconciliation, and replacing bespoke data requests with self-service. This dimension captures the efficiency dividend, which is often the most immediate and cleanly measurable return a data product produces, and frequently the one that funds the harder-to-attribute value elsewhere. A product that quietly removes thousands of hours of manual data wrangling has succeeded, even before its decision impact is proven.

4. The Data Product Scorecard

The six dimensions become usable when they are expressed as a scorecard, with an explicit question, example metrics, and a named owner for each. Ownership is not incidental. A dimension without an accountable owner is a dimension that will not be managed, and the single most common reason value measurement fails is that everyone assumed value was someone else's job. The scorecard below is a template, not a mandate; the specific metrics will vary by product, but the dimensions and the discipline of owning each one should not.

Dimension	The question it answers	Example metrics	Accountable owner
Usage & Adoption	Is anyone actually using it?	Active consumers; access volume; share of intended audience onboarded; repeat vs one-time use	Product Owner
Data Quality	Can it be trusted?	Accuracy & completeness of critical data elements; conformance to quality SLAs; consumer-facing quality incidents; freshness	Data Steward
Reliability	Is it there when needed?	Availability against SLA; pipeline success rate; timeliness of refresh; incidents and mean time to recover	Engineering / Platform Lead
Decision Impact	Does it change decisions?	Named processes it is embedded in; decisions citing the product; decision cycle-time reduction	Business Owner
Business Value	Does it move outcomes that matter?	Benefit realised vs business case; attributable revenue, cost or risk effect; value per unit of run cost	Business Sponsor
Reduced Operational Friction	Does it make work easier and cheaper?	Manual hours eliminated; duplicate/shadow sources retired; reconciliation effort removed; self-service rate	Operations Owner

5. Reading the Scorecard Without Gaming It

The scorecard is a management instrument, and like any instrument it can be gamed and misread. Three principles keep it honest. Balance beats any single number: a product excelling on adoption while failing on quality is heading for a trust collapse, and the scorecard should force that tension into view rather than average it away into a comfortable composite. Leading indicators earn their place alongside lagging ones: adoption and quality move early and predict value, whereas realised business value confirms late, so a scorecard of only lagging measures tells you the outcome long after you could have changed it, while a scorecard of only leading measures flatters activity that may never convert. And every metric must be owned, because an unowned metric is a number that is reported, not managed, and reporting without ownership is how a scorecard becomes theatre.

A word on gaming. Any metric that determines funding will be optimised, and some of the dimensions are easier to inflate than to satisfy honestly — adoption most of all, which can be manufactured by mandating use or counting incidental access. The defence is the balance itself: mandated adoption that produces no decision impact and no value shows up as a divergence across the scorecard, and a product whose adoption climbs while its decision impact stays flat is telling you something true, however uncomfortable. The scorecard is hardest to game precisely when all six dimensions are read together.

The attribution problem deserves an honest answer rather than avoidance, because it is the reason business value is so often dropped from measurement altogether. No data product is the sole cause of a business outcome; the market, the people, and a dozen other initiatives all contributed. But the choice is not between perfect attribution and none. A defensible approach names the decisions the product informed, estimates the value of those decisions being made better or faster, and states the assumptions openly so that others can challenge them. This is the same discipline any credible business case already uses. A transparent estimate that invites scrutiny is worth incomparably more than silence, because silence is what allows a product that delivers no value to survive indefinitely on the strength of having been delivered on time.

The scorecard's most valuable output is not the celebration of products that are succeeding but the identification of those that are not, and what to do about them. A product scoring badly is not automatically a failure to be cancelled; the scorecard tells you which kind of intervention it needs. Low adoption with high quality is usually a change and enablement problem, not a data problem, and the answer is investment in adoption rather than in the product. High adoption with falling quality is an urgent reliability and stewardship problem to fix before trust collapses. Strong usage and trust but no measurable decision impact suggests the product answers a question no important decision actually turns on, and should be refocused or retired. Read this way, the scorecard becomes a portfolio management instrument: it directs remedial investment to where it will pay, and it gives the organisation the evidence and the permission to stop funding products that will never earn their keep. The willingness to retire a delivered product is the clearest sign that an organisation has genuinely stopped measuring success by delivery.

6. Embedding Measurement in the Lifecycle and Governance

Measurement of this kind cannot be an annual event or a slide prepared for a review. To work, it must be built into the product from inception and into governance as a standing discipline. In practice this means three things. First, define the scorecard before build, at the point the product is conceived, so that adoption, quality, and decision-impact targets are commitments made alongside the delivery plan, not questions asked awkwardly after go-live when it is too late to instrument for them. Second, instrument the product to measure itself, so that usage, quality, and reliability are captured automatically as the product runs, rather than reconstructed by hand in time for a review — because measurement that depends on manual effort is measurement that will lapse the moment attention moves on. Third, make the scorecard the unit of portfolio governance, so that the question asked of every data product, at every review, is not only whether it was delivered but whether it is being used, trusted, and converted into decisions and value, with continued investment following the products that are earning it and being withdrawn from those that are not.

That last move is the one with teeth. Measurement changes nothing unless it changes where the money goes. A portfolio that measures the six dimensions but still funds on delivery has merely added reporting; a portfolio that reallocates towards the products demonstrating adoption, trust, and impact has changed its own behaviour, and through it the behaviour of every team that wants to be funded next year.

7. Recommendation

The recommendation of this paper is therefore specific, and can be adopted without reorganising anything.

1. Retain delivery metrics, but demote them. On time, on budget, and in scope remain necessary hygiene, not evidence of success.
2. Adopt the six-dimension scorecard as the definition of data product success, and require every data product to carry one from inception.
3. Assign each dimension a named, accountable owner — value that is everyone's job is no one's.
4. Instrument products to measure themselves, rather than relying on periodic manual assessment that will lapse.
5. Shift the primary question of portfolio governance from delivery to value in use — funding the products that demonstrate adoption, trust, and decision impact, and challenging those that were delivered but change nothing.

None of this is expensive relative to what a data portfolio already costs to build. What it changes is not the spend but the aim. An organisation that measures only delivery will build data products efficiently and discover, too late and too often, that efficient building was never the problem. An organisation that measures adoption, trust, and decisions will build fewer products that no one uses, and more that change what the business does. The scorecard is not additional overhead. It is the difference between a data portfolio that is merely busy and one that is worth the money.