

Decisions shouldn't become obsolete silently

OCTOBER 2026 · CONCEPT STAGE · CLIMATESOLVE.ORG

An institution commits capital on the best evidence it has. The evidence then moves. Nothing connects that change back to the decision it undermines, and the decision stays on the books until a scheduled review, an audit, or something going visibly wrong. This paper sets out why that gap exists, what would close it, and the smallest honest test of whether closing it is worth anything.

SUMMARY

In one page

Organisations are good at recording **why** a decision was made. Almost none can answer the reverse question: *something we believed has changed — which of our live decisions should we now look at again?*

Answering it requires keeping the reasoning behind a decision in a form that can be queried backwards. Record the evidence a decision rests on, the claims drawn from that evidence, and the assumptions the decision depends on. Watch the evidence. When something material changes, trace it forward to the handful of live decisions it actually reaches — and, just as importantly, past the many it does not.

Software has done a version of this for decades. Change one component and the build system knows exactly what depends on it and what has to be rebuilt. **Nothing does that for the decisions an institution has already made.**

We are starting with climate capital because the evidence moves fastest and the sums are largest: over **\$2 trillion a year** now flows into climate solutions, against an estimated need of around \$7.8 trillion a year to 2030.¹ But the mechanism is not climate-specific, and the paper says so plainly rather than pretending otherwise.

Nothing is built. The proposal is a single institution, a single decision already taken, and a defined point at which we stop.

THE PROBLEM

Everybody does their job properly, and it still goes wrong

Someone gathers the evidence. They write it up. A committee reviews it and signs it off. The reasoning is filed in minutes. That is the process working exactly as intended, and most institutions do it well.

Then the world moves on. A cost forecast is revised. A policy shifts. A pilot fails quietly. The new information is published, read, perhaps even discussed — and it never finds its way back to the decision that depended on it.

A worked example

A fund commits £40m to a technology, partly because it is forecast to halve in cost by 2030.

Two years later, that forecast is revised. It won't.

The revision is public. It is in a report that people in the sector have read. **But nobody connects it to the £40m.** The money stays where it is until the next scheduled review — or until something goes visibly wrong and somebody asks why.

The example is invented; the shape is not. It is the ordinary condition of every institution that makes decisions faster than it revisits them, which is all of them.

Why it persists

This is not carelessness, and treating it as such is the reason it has gone unaddressed. **No person can hold the reasoning behind every live decision in their head and simultaneously watch the entire world for changes to it.** The failure is structural, not personal.

Three things make it durable:

- **Reasoning is stored as prose.** Minutes, memos and board papers record conclusions in a form that cannot be queried. You can read them; you cannot ask them a question.
- **Evidence arrives without addressing.** A revised forecast does not know which decisions were built on the old one. It has no way of knowing.
- **Review is calendar-driven, not evidence-driven.** Decisions are revisited when the diary says so, which is uncorrelated with when the ground beneath them actually shifts.


The result is that the interval between a fact changing and a decision noticing is measured in quarters or years, and is set by the review cycle rather than by the change.

THE MECHANISM

Keep the decision joined to what it rests on

A decision rests on an assumption. The assumption rests on a claim. The claim rests on evidence. That chain exists in every considered decision; it is simply never written down in a form that survives the meeting.

Record it once, and the chain can be read in both directions. Forwards, it explains why the decision was made — which is what minutes already do. Backwards, it answers the question nothing currently answers.



The chain is constructed in one direction and has to be interrogated in the other. That asymmetry is the whole problem, and the whole product.

What the system actually does

1. **Records the reasoning at the point of decision.** Not a summary — the actual evidence cited, the claims drawn from it, and the assumptions the decision depends on, each as a separate object that can be referred to later.
2. **Watches the evidence those claims came from.** Continuously, rather than at review time.
3. **Tests whether a change is material.** Most changes are not. A system that flags everything is a system nobody reads, and this is the point at which most comparable efforts fail.
4. **Traces a material change forward** to the live decisions that depend on it, and presents those — and only those — to a person.
5. **Records what the person decided.** Changed, reaffirmed, or still open. The record is the asset: it accumulates into an account of why an institution believes what it believes.

evidence → claim → assumption → **decision**

evidence moves → which decisions does that reach?

The point is not finding new research. Plenty of tools do that, and the result is a longer reading list. The point is knowing that this particular change matters to that particular live decision — and to almost none of the others.

BOUNDARIES

What this is not

These are constraints we have chosen, not features we have yet to build. They are stated here because the category is crowded with products that do the opposite.

- **Not an AI that decides where climate capital should go.** The system has no view on the merits of any allocation.
- **Not a ranking of climate solutions.** It does not score, rank or compare options.
- **Not another feed of new research.** The problem is relevance, not volume.

- **Not a replacement for anyone’s judgement.** People decide and keep the authority to decide. The register only ensures they are deciding on evidence that still holds.

The last is the one that matters commercially. A system that recommends decisions inherits responsibility for them, and no institution will accept that from a third party. A system that surfaces *which decisions deserve a second look* asks for no authority at all.

THE PROPOSAL

One institution. One decision already made.

Not a pilot programme, not a platform rollout. One real allocation that is already on the books, and permission to map the evidence, claims and assumptions it actually rests on.

What that involves, concretely:

FROM THE INSTITUTION

One decision already taken, with the papers behind it.
Access to the person who can confirm what the decision actually assumed — usually a few hours, not a workstream.

A decision, at the end, on whether anything we surfaced was worth knowing.

FROM US

The dependency map, built by hand. The monitoring. The materiality judgements, each one shown with its reasoning so it can be disagreed with.

A written record of why that decision was made and what has changed since — which the institution keeps either way.

If something material changes and it genuinely reaches that decision, the institution knows before its next scheduled review. If nothing does, they have spent a few conversations and gained a documented account of their own reasoning.

How we would know it worked — and when we would stop

The question this tests is narrow and answerable:

Does tracing evidence changes back to live decisions surface something an institution would not otherwise have known, in time to act on it, often enough to be worth maintaining?

THE FALSIFICATION QUESTION

Three outcomes, and we have decided in advance what each one means:

- **Nothing material reaches the decision.** Either the evidence genuinely did not move, or our materiality test is too strict. Distinguishable, and worth a second attempt on a different decision.
- **Things reach it, but the institution does not care.** The mechanism works and the value does not exist. This is the result that ends the project.
- **Things reach it and the institution acts.** Then the question becomes whether the dependency map can be built cheaply enough to scale — a different and more tractable problem.

The stopping rule

If the first institution finds nothing worth having, **we stop and say so**. That is the point of starting with one decision rather than a programme: it is small enough to fail cleanly, and failing cleanly is the only way the answer means anything.

H O N E S T Y

The risks we think are real

Stated here rather than discovered later.

Building the map may cost more than it returns

Constructing an accurate chain of evidence, claims and assumptions is laborious, and it is the reason nobody has done this. Our position is that the cost should be paid by hand first, on one decision, to find out whether anybody values the result — *before* any effort goes into automating the construction. Automating the creation of something nobody wants is the most expensive possible mistake here.

Nobody may want to own it

This is the risk we take most seriously, and the one that is least discussed. A register that reopens decisions creates work, and occasionally creates blame. **An institution may be structurally better off not knowing that a decision has gone stale**. If that is true, no amount of engineering fixes it, and we would rather discover it in the first quarter than the third year.

Materiality is a judgement, not a calculation

Deciding whether a change matters enough to raise is the hard part, and it cannot be fully automated without reintroducing exactly the judgement we say we are not making. Our approach is to show the reasoning behind every materiality call so it can be argued with, and to accept that the threshold will be set differently by different institutions.

S C O P E

Climate is where this starts. It is not where it ends.

We begin with climate because the evidence moves fastest and the decisions are largest. But the same question sits under every field where consequential decisions outlive the evidence behind them — investment theses, clinical guidance, infrastructure business cases, public policy.

And increasingly, decisions made by machines. **An AI acts on what is true at the moment it acts**. Six months later some of that is wrong, and nothing currently asks which of its outstanding decisions should be reconsidered. It is the same dependency graph, and nobody owns it yet.

None of that is a reason to build anything now. It is a reason the narrow test is worth running.

Where this actually is

Concept stage. Nothing is built. There is no product, no pricing and no customer. This paper describes a mechanism and a test, and the only thing being sought is the first institution willing to put one real decision through it.

If this is a problem you recognise, the next step is a conversation rather than a demonstration — steve@cbmediagroup.com.

¹ Climate Policy Initiative, *Global Landscape of Climate Finance 2026*. Global climate finance crossed \$2 trillion for the first time in 2024, with roughly \$2.1 trillion estimated for 2025, against an estimated need of about \$7.8 trillion a year from 2025 to 2030.