Decision Reconstruction Risk

The failure mode that turns sound decisions into indefensible records, and why AI is accelerating it

By Phillip Wikes

I have spent much of my career reading records after the fact. In civil-rights investigations, in administrative reviews, in the moment when a decision made months or years earlier is finally examined by someone who was not in the room. And I kept seeing the same thing.

The decision was often fine. The record was not.

Not sloppy, not obviously incomplete. It read as professional. It was organized. It reached a conclusion. But when you asked it the only question that matters under scrutiny, *how was this conclusion actually reached?*, the record could not answer. The reasoning had lived in the decision-maker's head, in a hallway conversation, in context that was obvious at the time and invisible later.

I had not encountered a term that precisely described this pattern, so I began referring to it as Decision Reconstruction Risk.

Decision Reconstruction Risk (DRR) is the condition in which a record cannot explain, on its own terms, why a consequential decision was made.

The definition is deliberately narrow, and the narrowness is the whole point.

DRR is not "bad documentation"

This is the distinction most people miss, so it is worth stating plainly: **decision quality and record defensibility are different things.**

A decision can be excellent, reached by an expert, correct on the merits, and the record documenting it can still carry high Decision Reconstruction Risk, because the basis for it was never committed to the page. A different decision can be mediocre and yet fully reconstructable, because every step was anchored.

DRR does not measure whether a decision was right. It measures whether an independent reviewer, arriving with no prior knowledge, could rebuild the reasoning from the record alone. That is a separate property, and it is the property that determines whether an organization can defend itself when the decision is later questioned in litigation, audit, investigation, or a public-records request.

Why DRR has always existed, and why AI now accelerates it

Context is perishable. The people who made a decision understood why; the evidence was fresh; the reasoning was present. All of that decays. Personnel change, memories fade, informal explanations disappear. What remains is the record. DRR is the gap between what the record says and what it can prove.

That gap is not new. What is new is that artificial intelligence has made it both far more common and nearly invisible.

AI-assisted drafting produces records that are fluent, confident, and professionally worded at scale. And fluency is not evidence. The old warning signs of a weak record, prose that read as rushed or vague, are exactly what a language model smooths away. The result is a record that looks *more* finished while being *less* reconstructable. AI raises the volume of records and the polish of records without guaranteeing the one thing that matters under scrutiny: that the basis for the conclusion is present. DRR now spreads while wearing the appearance of its opposite.

How to recognize Decision Reconstruction Risk

A record carries DRR when one or more of these are true:

- The basis for a conclusion is asserted but not identifiable in the record.
- The chronology cannot be rebuilt from the record's own dates.
- The path from evidence to conclusion, what I call **Decision-Process Traceability**, is not visible on the page.
- Evaluative characterizations appear without a documented anchor.
- The record depends on its author remaining available to explain it.

Each is a symptom. Together they describe a record that looks complete and cannot be defended.

The property DRR destroys: documentation defensibility

For a long time, organizations talked about documentation *quality*: completeness, formatting, policy compliance. Those are real, and they are not the same as **documentation defensibility**.

Defensible documentation contains identifiable evidence, a traceable line of reasoning, an accountable decision path, and a chronology a stranger can follow. It holds up when the author is gone and the context is lost. Decision Reconstruction Risk is precisely the erosion of that property. You can have high quality and high DRR at the same time, which is why quality checklists miss it.

Naming the problem is how you begin to manage it

Every discipline has a moment when a diffuse, chronic frustration is finally given a name and becomes a manageable thing. "Technical debt" changed how engineering teams talk about shortcuts. "Alert fatigue" changed how clinicians and analysts design warnings. Naming does not solve the problem, but it makes the problem visible, discussable, and measurable. You cannot manage what you cannot name.

Decision Reconstruction Risk is diagnosable before a record is finalized. It can be detected, reduced, and designed out of a workflow. I have built a review discipline for exactly that, the Justification Review Standard, but I want to be clear about the order of importance: the tool matters less than the recognition. The first and most valuable step is to stop treating "the record didn't hold up" as bad luck or bad writing, and to start treating it as a specific, nameable condition with specific causes.

The point

The organizations that do well in the AI era will not be the ones that used AI the most. They will be the ones whose records can still explain themselves after the people and the context are gone. That is a solvable problem, but only once you can see it clearly enough to name it.

I call it Decision Reconstruction Risk. Now that it has a name, you can look for it, measure it, and reduce it, before the record becomes the only thing that remains.

Phillip Wikes introduced the term Decision Reconstruction Risk and developed the Justification Review Standard (JRS), a record-level review method for detecting and reducing it. He is a former Lead Civil Rights Officer at the Maryland Commission on Civil Rights. More at jrsstandard.com.