

LIVE WEBINAR SERIES

Claude × Smartsheet

Session 3: Where it goes wrong

The failure modes and how experts design around them.

1 Intro to Claude

2 What good looks like


3 Where it goes wrong

4 Doing it at scale

5 Running it unattended

Where we are in the journey

You've had the overview and seen what good looks like. Today: where it goes wrong and how experts design around it.



Three live failure modes coming up, each with an expert fix.

A quick refresher

First, the three ideas from last week and then we go find where they break.

1

What good looks like

Context, specificity, and a few follow-up turns separate expert use from amateur.

2

Write back safely

Parse the notes, review the table, approve, then let it write to the sheet.

3

Adopt it well

Right-size usage, start with one workflow, and name an owner.

The takeaway: *“Better questions get better answers.”*



● LIVE

DEMO 1

The confident hallucination

The most dangerous answer is the one that's wrong but sounds completely sure of itself.

Asking what the data can't answer

When the sheet can't support an answer, the prompt decides whether you get an honest one.



Amateur prompt

"Which projects are going to finish on time?"

It wants a verdict but sets no rules, no scope, and nothing telling it to flag what it can't answer.

→ A confident verdict, even when the sheet can't support one.



Expert prompt

"Using only the columns present in the "Q3 Delivery Plan" sheet, tell me which projects you CAN and CANNOT assess for on-time completion. For each one you can assess, state the specific column values you used. For each one you cannot, say exactly what data is missing. Do not estimate or fill gaps. If the sheet has no progress or forecast field at all, tell me that first."

→ Tells you what it can and can't assess, and shows the evidence.

Why this one is dangerous

1

Confidence $\neq$ correctness

A sure-sounding answer is the one nobody thinks to double-check.

2

Demand the evidence

Make it cite the exact values behind every single claim.

3

Let it say “I can't tell”

An AI that never admits uncertainty is a risk, not a feature.

“The dangerous output isn't the wrong one. It's the confident one.”

A confident answer from a partial view

The permission model is a security win and a quiet trap.



The guarantee

Claude inherits your exact permissions. It can never see more of a sheet than you can.



The hidden gotcha

So if your view is restricted, the answer still looks complete and nothing flags the rows, or the systems, you can't see.

Before you trust it, ask: am I looking at the whole dataset, or just my slice?



● LIVE

DEMO 2

The silent bad write

An update that quietly hits the wrong rows is worse than no update at all when no one notices.

Watch it succeed and quietly get it wrong

The danger isn't a refusal. It's a confident “done” that hides a bad write.

- 1** A confident “done” before anyone checks the work.
- 2** The difference reveals rows matched to the wrong task.
- 3** You revert the bad ones; the unsure ones get quarantined, before a wrong value reaches the systems downstream.

The review-the-difference pattern

1

Match, don't guess

Tie each change to exactly one row by name and flag the ambiguous ones.

2

Show the difference

Current → proposed, with the source note, before any write happens.

3

Quarantine the doubtful

Low-confidence matches go to a separate list for you to handle.

The difference is the line between automation you trust and silent corruption.

● LIVE

DEMO 3

Design the sheet for reliability

The biggest improvement isn't always a cleverer prompt. It's also about sheets the AI can actually trust.

Same prompt, two very different sheets

We didn't touch the prompt. We changed the data the AI had to work with.



Messy sheet

- Owners often blank
- Free-text status: “wip”, “ok?”, “done!”
- Inconsistent date formats
- No progress or forecast field

→ The AI has to guess. Output is shaky.



AI-ready sheet

- Owner column, always filled
- Status from a controlled list
- One consistent date format
- % Complete + Forecast Date columns

→ The AI cites real values. Output is reliable.

The biggest lever is the sheet

1

Controlled lists

Status from a fixed set, not free text the model must interpret.

2

Real fields

Owner, % complete, forecast date, not buried in notes.

3

One source of truth

Consistent formats the AI can read the same way every time.

We didn't make Claude smarter. We made the data legible.

THE TAKEAWAY

Reliability is designed, not hoped for.

Trustworthy AI on your data depends on structure, evidence, and checkpoints, and that expertise is exactly what we bring.

Still to come

SESSION 4

Doing it at scale

A full governed weekly workflow running on live data and what it takes to run it for real.

SESSION 5

Running it unattended

Portfolio roll-ups across many sheets, running on a schedule with your checkpoint.

YOUR TURN

Questions and a takeaway

“How would you make this reliable on my sheet?” drop it in the chat.

Get our Reliability Checklist for AI-ready Smartsheet. Leave your email in the chat and we'll send it.