QUADRATIC RESEARCH BRIEF 01

Claude in Excel: what works, what breaks, *and who it's for.*

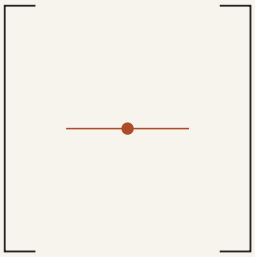

A workflow-maturity framework for operators, analysts, and technical buyers evaluating AI inside the modern spreadsheet.

AUTHOR
Dr. M. J. Lindeman
Editorial contributor, Quadratic Research

PUBLISHED
April 7, 2026
Findings current as of March 2026

READING TIME
~18 minutes
10 sections · 1 framework

00 CONTENTS

In this brief.

01 Executive summary 02
The headline finding in three paragraphs.

02 Why this matters now 03
The decision risk underneath the noise.

03 What "Claude in Excel" is 03
Add-in vs. native AI architectures.

04 Workflow maturity framework 04
Three tiers: convenience, builder, power.

05 What CiE does well 05
Formulas, cleanup, document extraction.

06 What breaks (and why) 06
Persistence, automation surface, reproducibility.

07 Choose Quadratic if... 08
Seven decision criteria.

08 Appendix 09
Two formula-equivalence examples.

01 EXECUTIVE SUMMARY

The headline finding.

Claude in Excel (CiE) brings Claude's assistance into Excel as an add-in, and it's strongest for day-to-day spreadsheet tasks where speed matters: cleaning messy data, generating and explaining formulas, tracing dependencies across complex workbooks, and producing quick summaries.

For repeatable reporting workflows—dashboards, models, or recurring reports that need to be updated, revisited, and maintained over time—CiE can work, but with an important constraint: its chat reasoning does not persist between sessions, creating auditability and handoff gaps unless you build workarounds.

For automation, governance, or integration-heavy workflows at scale (VBA / Power Query-driven processes, scheduled runs, and auditable pipelines), CiE's architectural limits are often significant enough to make it a poor fit for core parts of the workflow.

THE BOTTOM LINE

CiE works well for users who work in self-contained tasks, value speed over traceability, and are not building processes that others will maintain or audit. It is less well matched to users whose workflows span multiple sessions, require reproducible logic, or depend on automation that runs without a user present. The three-tier framework on the next page maps those distinctions precisely.

02 WHY THIS MATTERS NOW

The decision risk underneath the noise.

The volume of AI announcements in the spreadsheet space has made it genuinely difficult to know where to invest your time and money. In the last few months, major vendors have added AI capabilities to their spreadsheets, each announcement arriving with similar claims: faster analysis, smarter formulas, easier workflows. The noise is not a minor inconvenience — it is a real decision risk.

The cost of getting it wrong runs in two directions. Teams that dismiss AI entirely miss genuine productivity gains in formula generation, error tracing, and data cleaning. Teams that adopt features without understanding their limits end up with processes that have problems: the AI may automate a task but require the same human verification it was supposed to eliminate.

> *Spreadsheets are evolving from passive tools that wait for instructions into platforms that can act on a schedule, connect to live data sources, and preserve the reasoning behind every result.*

03 WHAT "CLAUDE IN EXCEL" IS

Two architectures, two ceilings.

AI is built into spreadsheets using two fundamentally different design approaches. Which approach a tool uses explains why certain limitations are structural rather than bugs to be fixed in the next update.

Add-in vs. native AI — what that changes

The first approach is to add AI as an external layer on top of an existing spreadsheet application. CiE is the clearest current example. It installs through Excel's add-ins menu and accesses your existing Claude account. It operates as a sidebar alongside the spreadsheet, reading the workbook and taking actions within it, but running as a separate process rather than as a component of Excel itself. It supports .xlsx and .xlsm files and preserves formulas, cell relationships, and existing formatting.

The second approach is to build AI in as a native component from the start. In this model, the AI layer, the spreadsheet engine, and the data connections share the same environment. They were not integrated after the fact because they were designed together. The practical consequence is that the native AI has access to the full capability surface of the application — which may be much more than the operations the add-in vendor has enabled. The distinction matters most when a user reaches the edges of what the add-in can do.

04 FRAMEWORK

The workflow-maturity framework — who it's for.

Not every spreadsheet user needs the same things from AI. The features that matter and the limitations that create problems depend almost entirely on the complexity of the work being done. Three usage tiers define the relevant distinctions. They are anchored to *workflow complexity* rather than technical sophistication.

TIER A

Convenience users

Ad-hoc analysis: quick cleanup, basic pivots and charts, simple summaries. Success metric: faster answers with minimal setup.

Each task is self-contained, so the absence of session persistence barely registers. Output is verifiable before you accept it.

STRONG FIT ●

TIER B

Builder users

Repeatable workflows: advanced formulas, structured models, custom dashboards, refreshable reporting. Success metric: a reliable, reusable analysis that holds up over time.

Workable, but the reasoning behind a formula is hard to recover when it was built through a conversation CiE does not store between sessions.

WORKABLE, WITH GAPS ●

TIER C

Power users

Automation, integration, and governance: VBA, Power Query, complex models at scale, auditability, and integration with other systems.

CiE does not support VBA, macros, and data tables within the AI layer. For workflows depending on these, it is not the answer for that part of the work.

ARCHITECTURAL MISMATCH ○

Tier B — where the add-in pattern starts to show its limits

A user who opens the workbook the next day faces internal logic that cannot be interrogated by continuing the original chat. CiE offers a logging function that prints actions to a sheet. Whether that log captures the *reasoning* behind those actions or only the cell-level changes is a question to verify before treating it as a full audit trail.

If this starts to feel like rework — rebuilding context each session, or handing models across teammates — it may be worth using an AI-native spreadsheet where conversations and logic persist alongside the workbook.

Tier C — explainability is not optional

A Tier C user building a compliance model cannot accept the absence of reasoning that supports the recorded actions. Explainability and reproducibility are not uniform requirements across users. They *increase with tier*, and the tool needs to match.

05 WHAT WORKS

What Claude in Excel does well.

Across all three user tiers, certain capabilities hold up consistently in real use. These are reliable, available now, and worth adopting — with minimal caveats.

Formula generation and explanation

The most widely used and most understood capability. AI generates formulas from natural language descriptions, explains what existing formulas do, and traces errors back to their source.

A user who inherits a workbook with 30 tabs, 200 formulas, no documentation, and a deadline can ask the AI to explain what each formula does in plain English, trace dependencies across sheets, and map how data flows through the workbook. That is a genuine capability that saves real time, and it is available to users at every paid tier.

Data cleaning and normalization

Holds up equally well. When a spreadsheet arrives with dates in five different formats, names split inconsistently across columns, and duplicate rows throughout, the cleanup is hours of manual work.

AI handles it reliably because the task is pattern recognition applied to repetitive structure. The verification responsibility does not go away, but the time required drops significantly.

Document and PDF extraction

PDF extraction is one of the least glamorous and most immediately useful recent developments. Upload a receipt or financial report and the AI understands the document structure and imports the data, eliminating the copy-paste-clean chore. Extracted data should be spot-checked, particularly for financial documents where source formatting may be inconsistent.

Installation and account requirements

CiE installs through the Microsoft Marketplace via Excel's add-ins menu and requires an active Claude account at Pro, Max, Team, or Enterprise. For organizational deployments, IT admin deployment through the Microsoft 365 Admin Center is needed; if an organization has disabled "Let users access the Office Store," IT must use the manifest XML method. The number of troubleshooting workarounds documented online indicates installation may be problematic — worth checking before a team rollout.

06 WHAT BREAKS (AND WHY)

Where the architecture stops cooperating.

Session persistence and auditability

> *Anthropic's documentation states that "chat history is not saved between sessions" and that "each time you open the add-in, you start a fresh conversation with Claude."*

Every conversation about the workbook, every instruction refined through an hour of work, and every decision about how to handle an edge case is gone when the file closes.

A separate auditability gap affects enterprise users. Documentation states that observability and auditability are not currently available for CiE, that CiE does not inherit custom data retention settings, and that it is not included in Enterprise audit logs or the Compliance API. For regulated industries, this is not a workaround problem — it is a hard limit.

SILENT AUTO-COMPACTION

Long sessions discard detail without notice.

In long sessions, Claude automatically summarizes earlier exchanges to make room for new ones, discarding detail in the process. The compaction is silent: no notification, no record of what was lost. The first sign that compaction has affected a session may be that Claude contradicts a decision made an hour earlier without knowing it is doing so.

Scheduled runs and MCP-connected data

For users whose goal is to build processes that run without them — for example, a monthly report that generates on a schedule — CiE is not the answer. Instead, the user writes Excel scripts that run on files saved in SharePoint or OneDrive.

One operational consideration is relevant specifically for Tier B and Tier C users: Anthropic documents that MCP connectors configured in your Claude account activate automatically inside Excel. For an account with connections to financial data providers — S&P Global, LSEG, Pitchbook, Moody's — data from those sources can move into Excel workbooks through the AI layer. This may be an important constraint for users managing workbooks in environments with data governance requirements.

Automation surface area

CiE's layered add-in approach means Anthropic must specifically enable each Excel-native operation. As of March 2026, Claude can apply a defined set directly: sorting and filtering, editing pivot tables and charts, applying conditional formatting, setting data validation, and preparing workbooks for printing. Each capability had to be built and granted individually — not the same as native access to the full Excel object model.

Reproducibility and traceability

Financial modeling from scratch is one of the most promoted CiE use cases and one of the most genuinely risky. AI can generate a working model structure faster than most analysts can build one manually. The more fundamental issue is what happens when the model needs to change. A model built in a conversation that no longer exists, using assumptions refined through exchanges Claude cannot remember, is difficult to modify correctly. The model persists; the conversation that created it does not.

SECTION SYNTHESIS

The pattern across all three failures is the same.

Reasoning lives in the chat session. The session ends with the file. If reasoning needs to outlive the session — for handoff, re-work, or audit — the architecture fights you. Workarounds substitute for capability the tool does not currently expose.

07 DECISION GUIDANCE

Choose Quadratic if any of these are true.

Seven decision criteria. They apply across all three user tiers. If any single one applies, the add-in pattern will cost you more time over twelve months than it saves you in the first week.

01	You need the reasoning to persist across sessions.
02	More than one person will maintain the file.
03	You want repeatable reporting that refreshes on a schedule.
04	You need to combine data sources or use code for analysis.
05	You're spending time on exports / CSV workflows and rebuilding the same analysis.
06	Your analysis depends on large datasets or complex joins that become brittle or slow in traditional spreadsheets.
07	You care about auditability and traceability.

CiE is fine for simple tasks; Quadratic is simpler in the long term. If you expect to reuse the workflow, share it, or build it into reporting, start with Quadratic so you don't rebuild later.

When native AI matters

Quadratic is built on the native-AI model. The AI layer, code cells, and data connections were designed in the same environment from the beginning. There is no add-in to install and no separate account to manage. Complete chat sessions are stored in chat history, which means the reasoning behind any analytical conversation is available to return to and continue.

The practical consequence shows up most clearly for Tier B and Tier C users. A Tier B user whose reporting model needs to be maintained and understood by more than one person has access to the full conversation history that built it. A Tier C user who needs automation can use Quadratic's scheduled tasks to build processes that run without anyone present to trigger them.

Formula fluency carries over

Quadratic preserves what a user already knows about Excel formulas — relative A1 notation, absolute locking, named ranges. Ask for the Excel formula for a cell and the AI will provide it if one exists, or explain how to use tools in Excel if it does not. An example is in the Appendix.

08 APPENDIX

Two examples of formula equivalence.

One concern specific to finance users who are power users of Excel is worth naming directly. A common objection to AI-native spreadsheet tools is that finance professionals do not trust Python because they do not know it. Quadratic addresses this in two ways. First, the platform's AI can write formulas and code (Python and SQL). Users who prefer to work in familiar Excel-style syntax can do so; power users who want to see or write code can do that.

Second, when a result has been calculated using code, you can ask the AI to show you the equivalent Excel formula (if it exists) and explain exactly what the code is doing in plain English. The reasoning behind any result is available, not hidden. Two examples follow.

EXAMPLE 01 — EQUIVALENT EXCEL FORMULA EXISTS

"The value in H14 is the percentage change (`pct_change` *) of the closing price. It's calculated as:"*

```
=(F14-F13)/F13
```

"This takes the current close price (`F14 = 176.3` *) minus the previous close price (* `F13 = 179.36` *), divided by the previous close price, yielding approximately* `-0.01706` *(or about* `-1.71%` *)."*

EXAMPLE 02 — NO EQUIVALENT EXCEL FORMULA

"Cell H29 contains the Positive sentiment score (currently `0.89` *) computed by NLTK's* `VADER SentimentIntensityAnalyzer`*. There is no direct Excel formula equivalent. … This kind of NLP processing is beyond what Excel formulas can do natively. To replicate it in Excel, you would need:"*

- Excel + VBA macro calling a Python script or API.
- Excel + Power Query connecting to an external sentiment analysis API (e.g., Azure Cognitive Services, Google Cloud NLP).
- Excel Labs' Python in Excel feature (which could run the same NLTK code).

"In short, this is a task that requires a programming language or external service — it can't be reduced to a spreadsheet formula."

These examples demonstrate the difference between a tool that uses AI to manipulate numbers and a tool that unlocks the potential of AI to go beyond traditional number crunching.

END ONE LAST THOUGHT

Your raw data, your existing spreadsheets, and your spreadsheet formula skills *transfer instantly* from Excel or Sheets to Quadratic.

The add-in model has a hard ceiling defined by what the host application exposes and what the add-in vendor has enabled. The work outgrows that ceiling when, for example, the absence of session reasoning is a workflow problem or its auditability is a requirement. Then the answer is not a better add-in. It is a different architecture.

TAKE IT FURTHER · 01

Learn more about Quadratic.

Read documentation, see the AI-native spreadsheet model in detail, and explore example workbooks.

quadratichq.com

TAKE IT FURTHER · 02

Request a demo.

See your own workflow inside Quadratic. We'll bring a workbook of your shape, or build with one of yours.

quadratichq.com / book-demo

AUTHOR
DR. M. J. LINDEMAN
EDITORIAL CONTRIBUTOR

BRIEF
QUADRATIC RESEARCH
№ 01 · SPRING 2026

PUBLISHED
APRIL 7, 2026
FINDINGS CURRENT TO MARCH 2026