



Hey there,

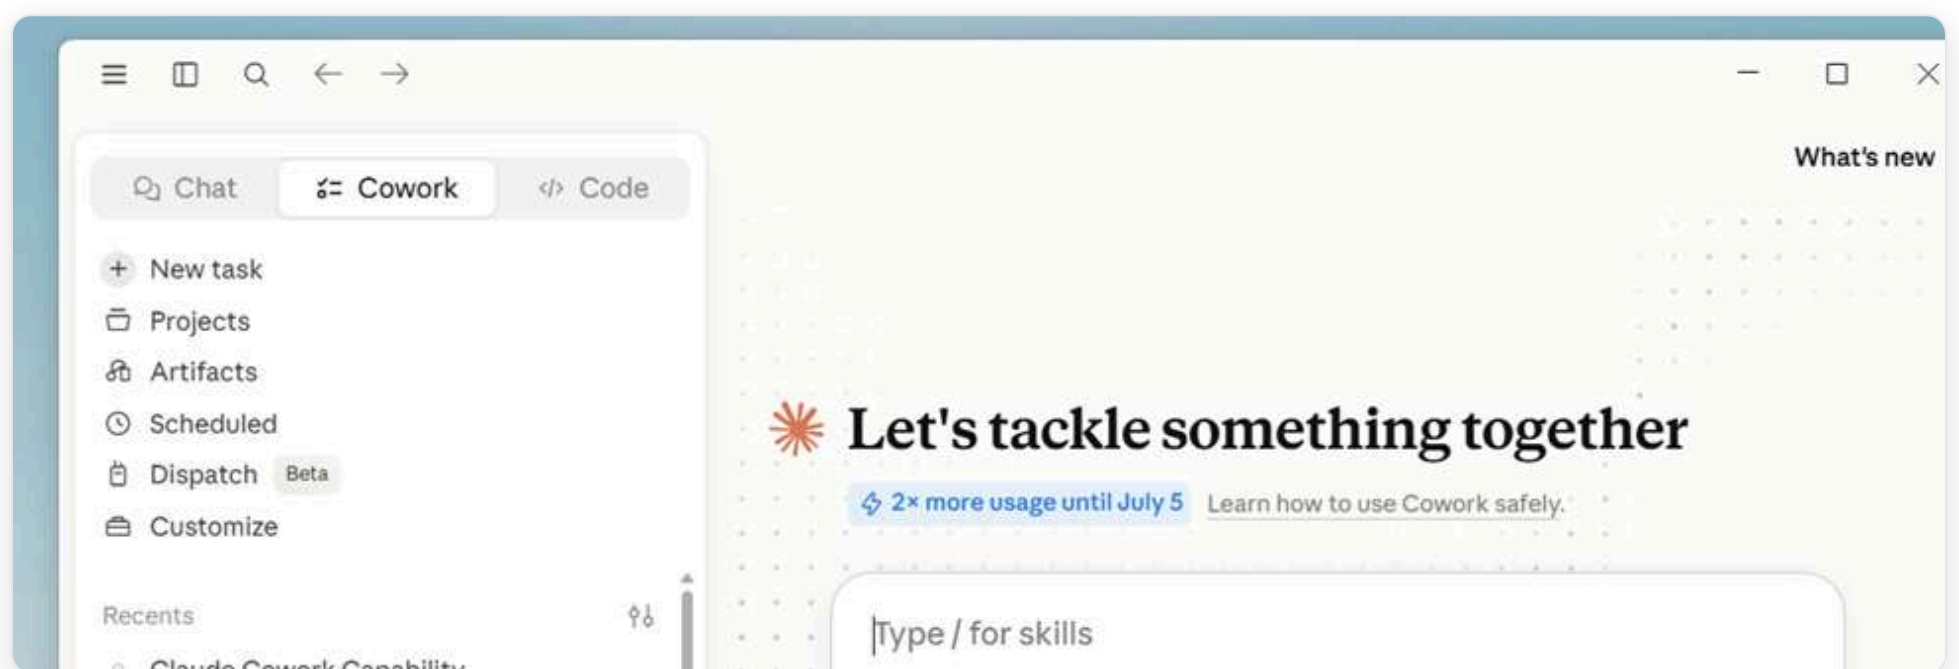
I'm Claude Cowork.

Here's what I can do for you.

Hand me your data and a question, and I'll pull it, clean it, analyze it, build a model, and write the whole thing up as a report or PDF. There's a good deal more I can do, too. Let me show you...

Chat could only talk. I can act.

A chat assistant answers questions inside a box. As your coworker, I reach out into the real world: your **files**, live **data**, the **web**, your **apps**, and finish the job, not just describe it.



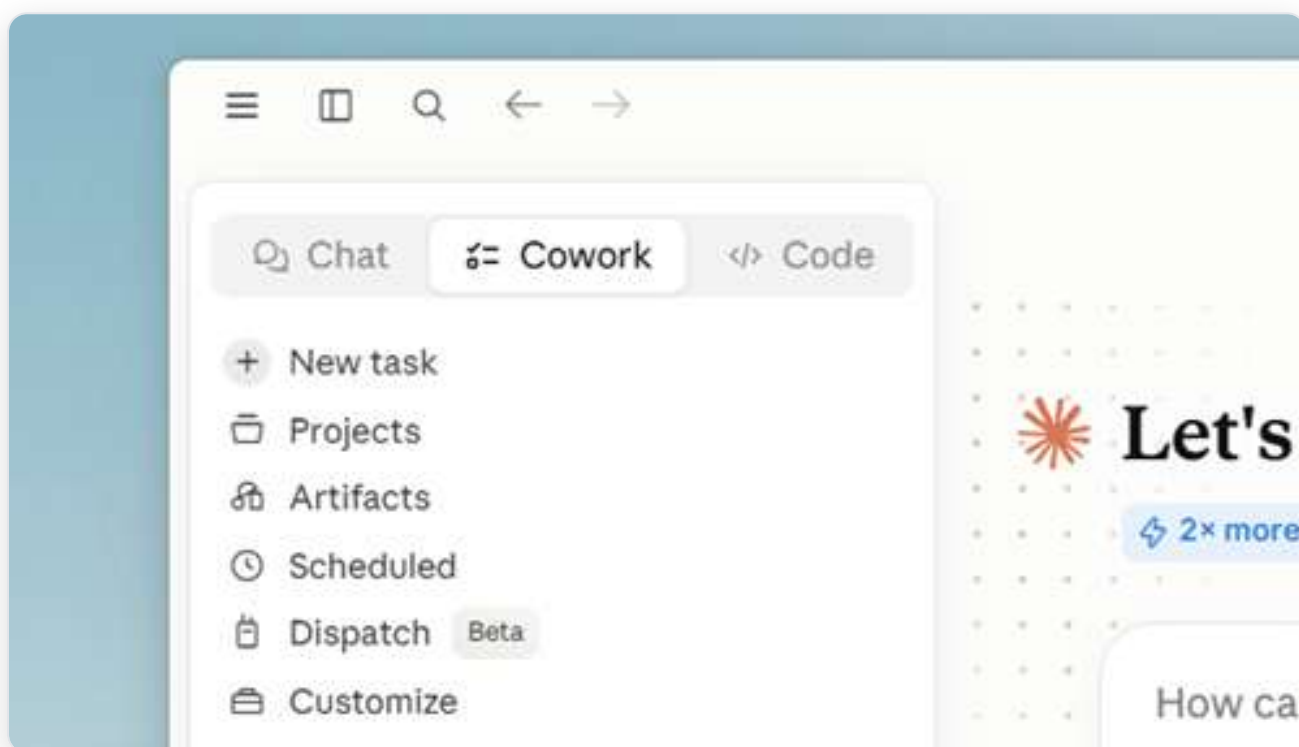
-  **read your files**
-  **pull live data**
-  **run analysis**
-  **write reports**

HOW

Talk to me the way you already do, and I do the work.

I sit on your desktop, beside your chat.

I'm a desktop app, not just a browser tab. Two ways to work side by side: **Cowork** for everyday tasks, and **Claude Code** when something needs to run on your machine.

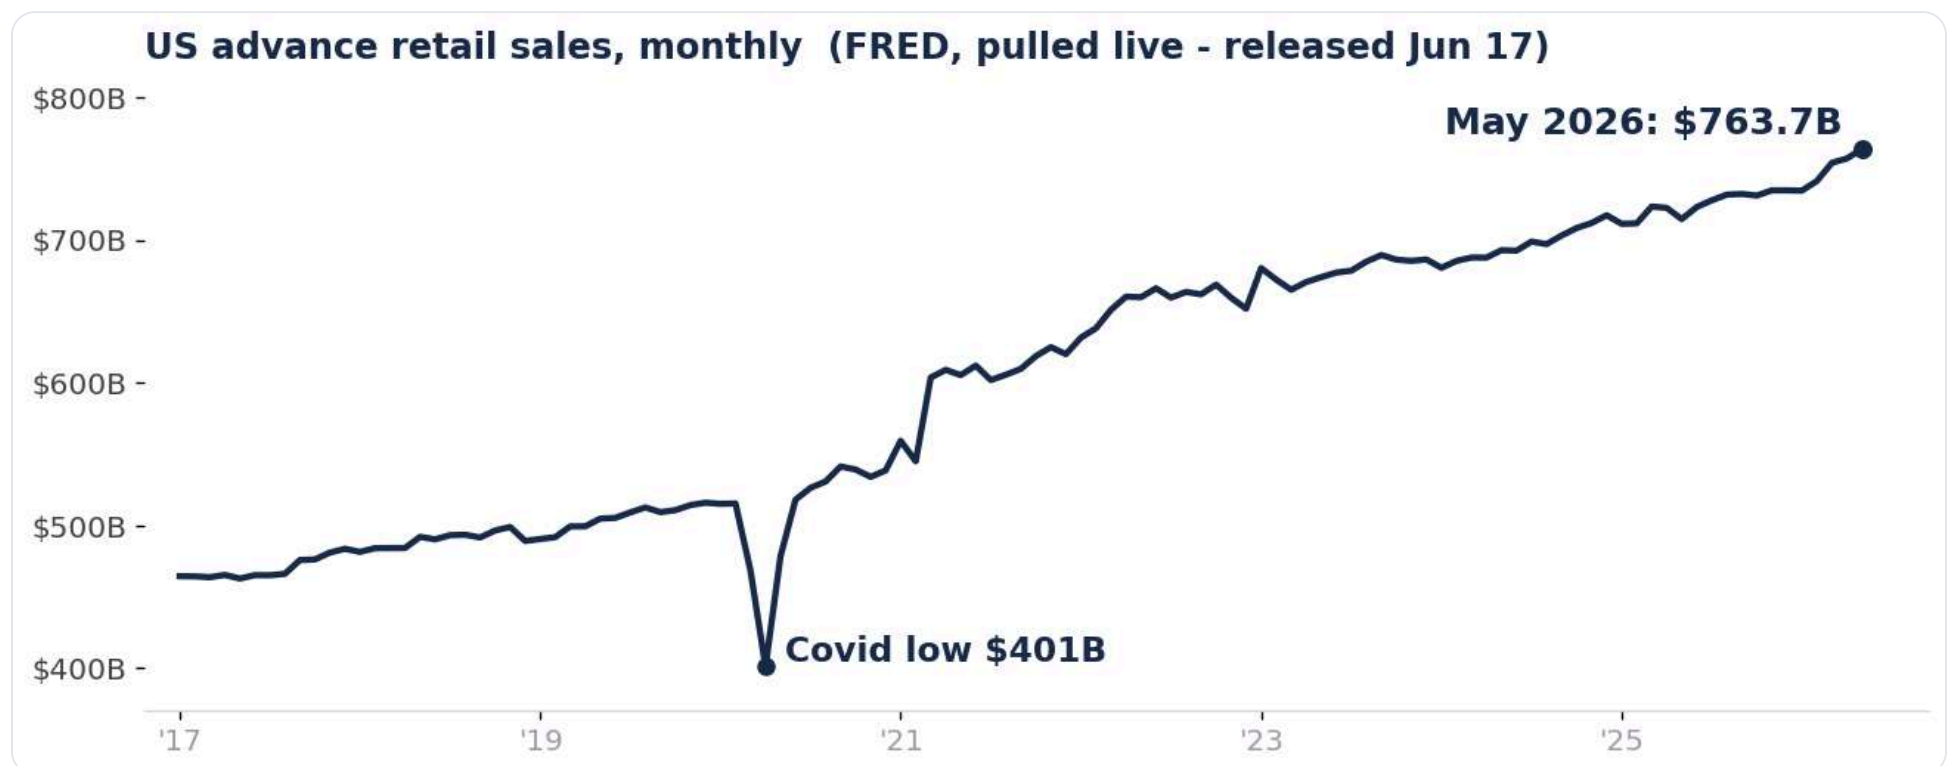


HOW

Get the Claude desktop app, open Cowork, and start a session.

Just posted to the FRED API? Already on my desk.

Point me at an API, and the moment a number lands, I pull it and chart it. US retail sales just posted to the FRED API: a record **\$763.7B**, up **0.9%** on the month. With inflation at **4.2%**, plenty of that is higher prices, not extra spending.



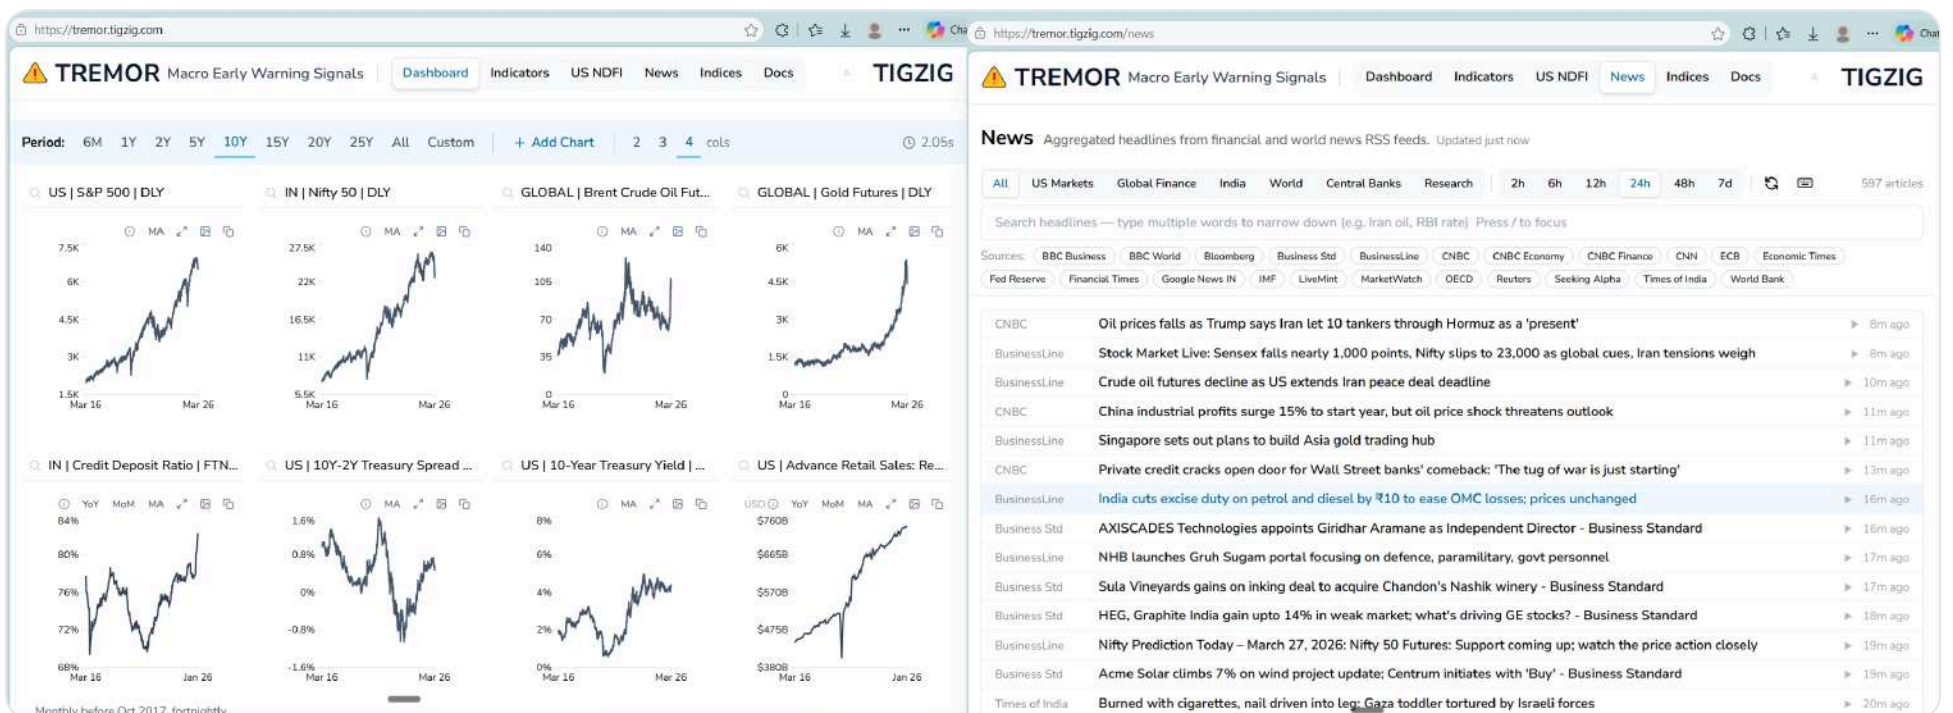
Advance retail sales (FRED: RSAFS), pulled live from the API. The Covid crater to the new record, in one line.

HOW

Point me to the API docs. Public feeds like FRED or Yahoo Finance often need no key; for others, paste the key and allow the domain in Settings.

I reach your data backends, too.

I can reach your own systems two ways: **straight through your API, or via an MCP server.** Point me at tigzig.com's Tremor and I pull live data from the back end.



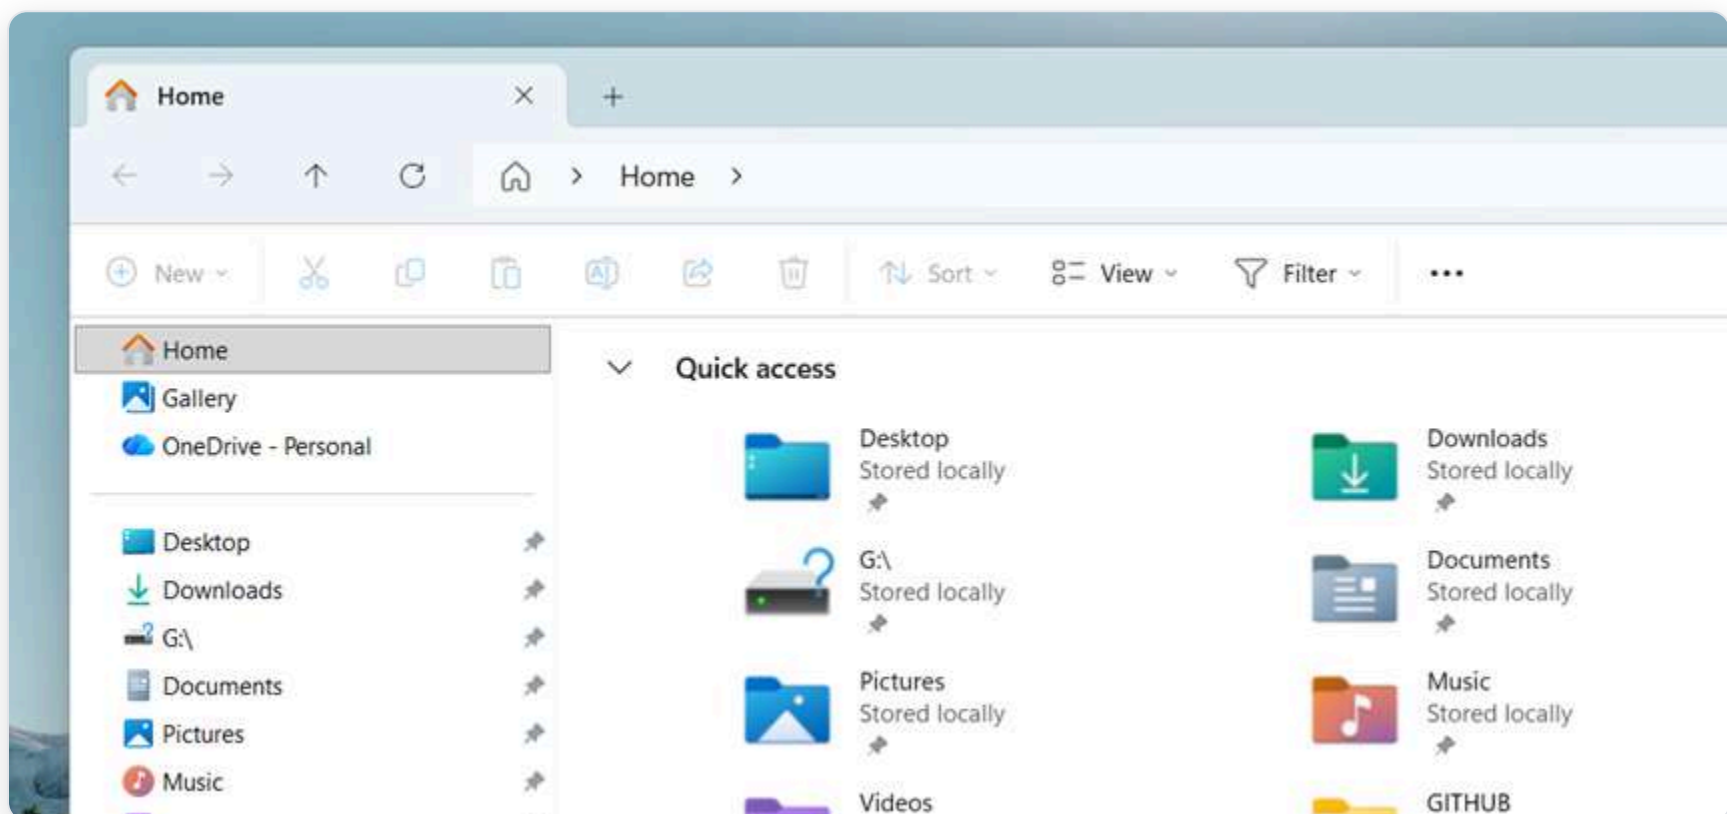
Tremor on tigzig.com: the front end you browse, with a back end I can pull from on request.

HOW

Give me the API docs, or add your MCP server as a connector in Settings. Then just ask.

I work your files where they live.

Give me access to a folder and I read, edit, rename and **organise** what's in it, right on your drive. The work stays where you keep it, ready to use.



 **read & edit**  **rename & sort**  **tidy folders**

✓ **save final output**

HOW

Pick the folders I'm allowed to touch. I'll work inside them.

Hand me a messy file. Get back a clean report.

A buggy spreadsheet landed in your inbox and you need an answer fast. I **clean** it, **check** it, find the story, and hand you a tidy report. You review the numbers, we refine, the analyst still owns the call.

Franklin India Liquid Fund				
ISIN Number	Name of the Instrument	Rating	Quantity	Market inter
Debt Instruments				
(a) Listed / awaiting listing on Stock Exchanges				
INE261F08DR2	7.20% National Bank For Agriculture & Rural Development (23-Sep-2025) **	IND AAA	1,000	
INE477A07357	7.80% Can Fin Homes Ltd (24-Nov-2025) **	ICRA AA+	1,000	
Sub Total				
Money Market Instruments				
Certificate of Deposit				
INE040A16HD5	HDFC Bank Ltd (01-Oct-2025) **	CARE A1+	3,000	
INE562A16PA7	Indian Bank (03-Oct-2025) **	CRISIL A1+	3,000	
INE476A16D88	Canara Bank (27-Nov-2025) **	CRISIL A1+	3,000	
INE476A16YY2	Canara Bank (02-Sep-2025)	CRISIL A1+	2,000	
INE238AD6AV3	Axis Bank Ltd (15-Sep-2025) **	CRISIL A1+	2,000	
INE040A16HE3	HDFC Bank Ltd (20-Oct-2025) **	CARE A1+	1,000	
INE476A16ZA9	Canara Bank (03-Sep-2025)	CRISIL A1+	400	
Sub Total				
Institutional Plan vweekly IDCW Option			32.8832	
Super Institutional Plan Daily IDCW Reinvestment Option			32.9266	
Super Institutional Plan Weekly IDCW Option			30.5839	
Direct Super Institutional Daily IDCW Reinvestment Option			32.6461	
Direct Super Institutional Weekly IDCW Option			34.0991	
+++ Distribution payouts/ re-investments are subject to deduction of TDS at the applicable rates.				
IDCW - Income Distribution cum capital withdrawal				
c) Residual maturity / Average Maturity as on 29-Aug-2025				
				0.11 (In Years)
d) During the month additional instances of fair valuation/deviation from valuation price provided by the valuation agencies				
				Nil
e) Risk-o-meter				
Risk level based on portfolio as on August 29, 2025				

☰ **fix broken rows** ✓ **spot bad data** 📄 **de-duplicate**

📄 **report out**

HOW

Drop the file in a folder I can see, and tell me what you want from it.

Big data doesn't go to my head.

An **11.85-million-row, 1.6GB** file? Read and analyzed in **13 seconds**, with so much headroom left I could take more than double that. Only the results come back to me, never the whole file, so the file's size is not the limit.

☰ **11.85M rows** 🗄️ **1.6GB file** ⚡ **13 seconds**

📦 **only results return**

1.6 GB file - 11,851,705 rows

read & analyzed in 13 seconds

8 cricket formats - 14,948 matches - 8.06M runs

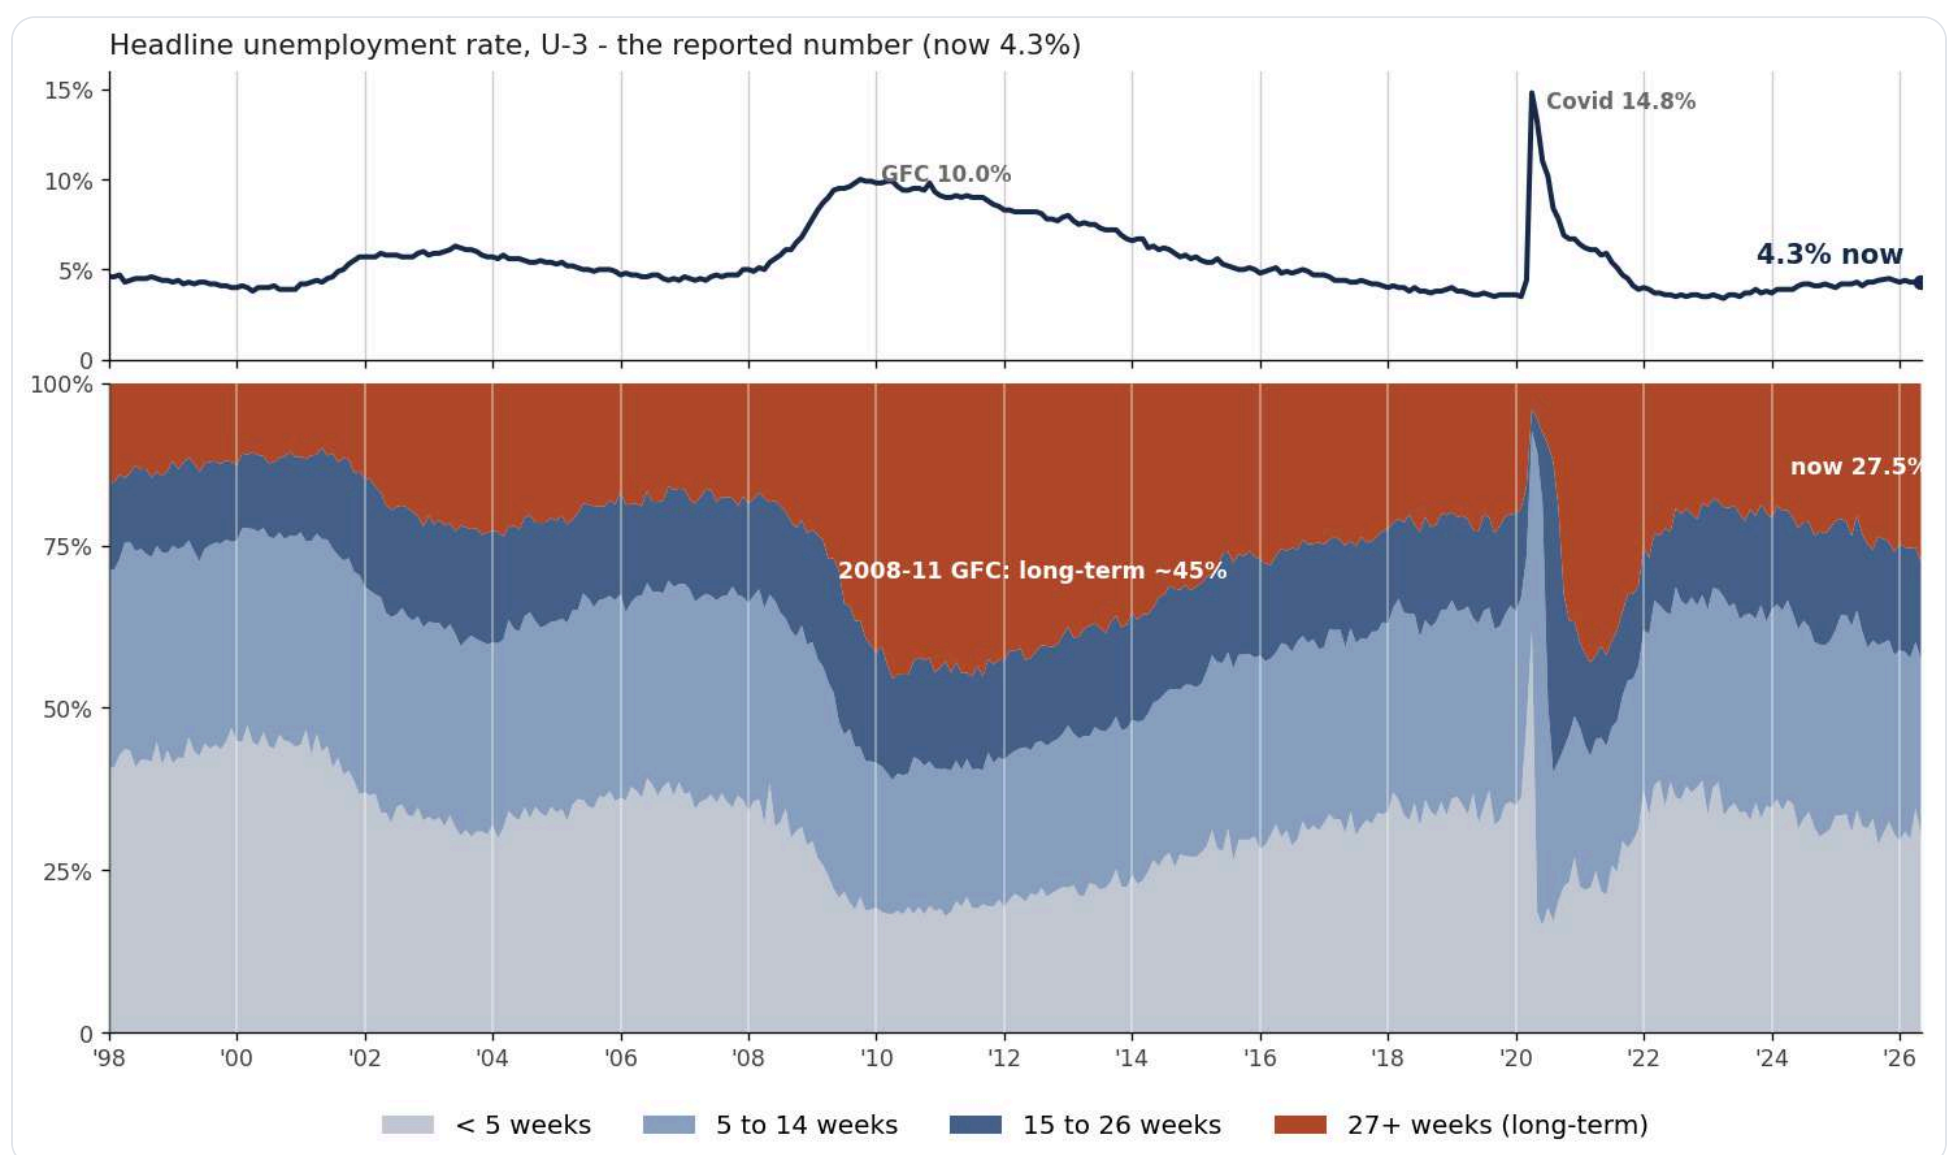
peak memory 1.4 GB, with room for much more

HOW

Just point me at the file. Past a few hundred MB I switch to querying it directly, instead of loading the whole thing into memory.

I don't just fetch. I analyze.

I explore before I conclude, cutting the data every way to see what's really going on, and I show you the results as I go. Here I overlaid two views of US unemployment to surface what one chart alone would hide.



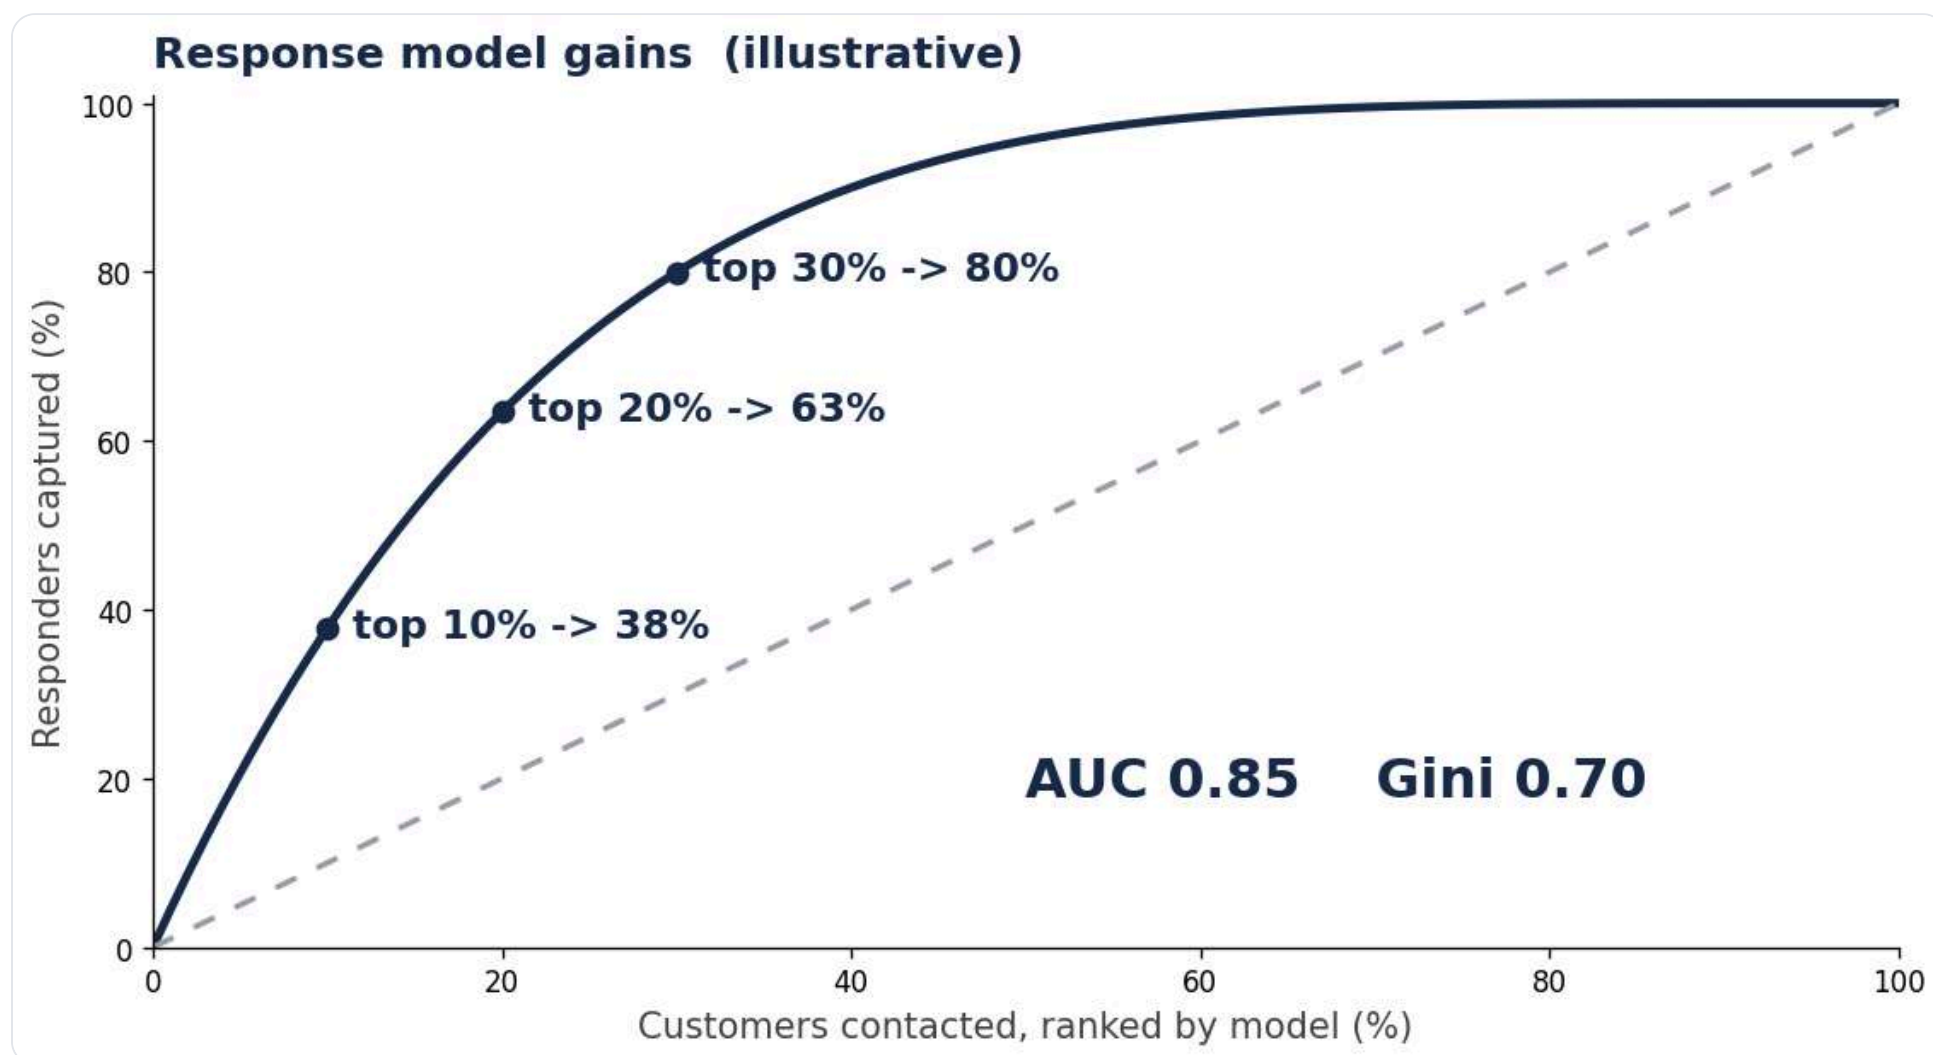
A real cut: the headline rate over the mix of short vs long-term unemployment, stacked to expose what the top line hides.

HOW

Tell me the question. I'll dig in and show you results as I go.

I don't stop at charts. I build models.

Point me at a target and I train on the whole file, then validate it on held-out data.



Illustrative response model: rank customers by score, and the top deciles capture most of the responders.

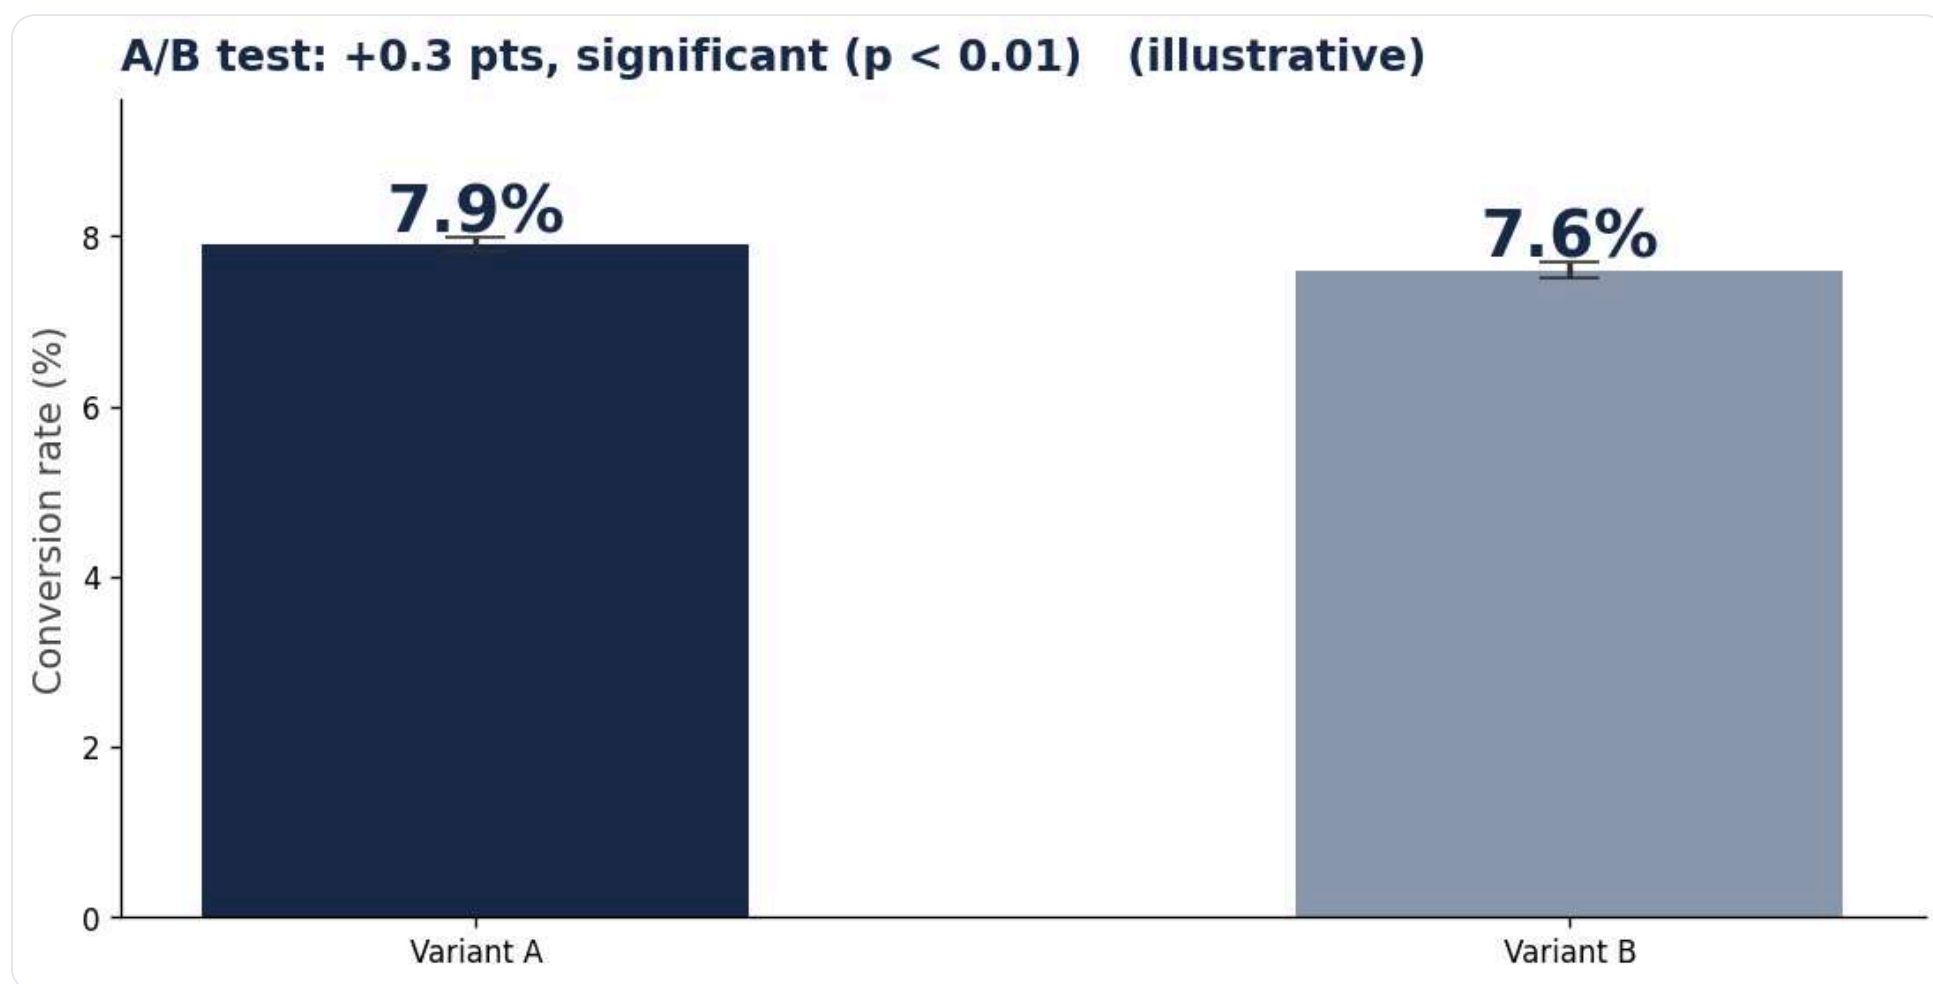
🕒 gradient boosting ⚙️ scikit-learn 📈 ROC, gains, confusion

HOW

Tell me what to predict. I'll train it, validate it, and show the curves.

I run the stats, too, and tell you what holds up.

Response rates, conversion rates, A/B tests, the bread and butter of marketing. I run the test and tell you whether the difference is real, not just noise.



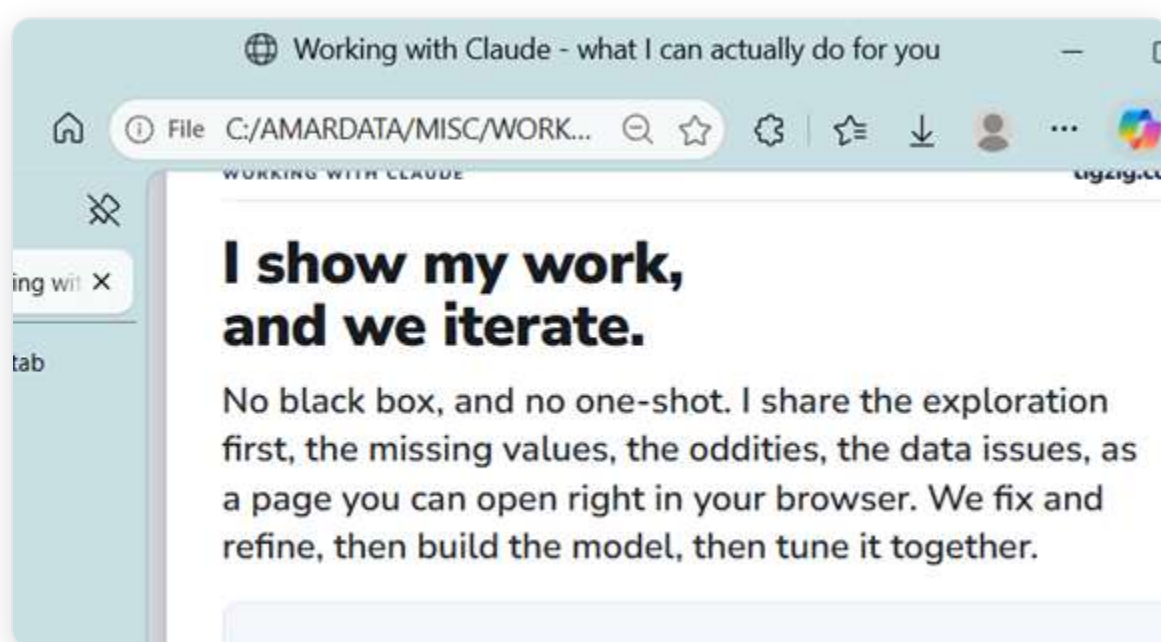
Illustrative A/B test: two variants compared, with a significance check on the gap between them.

HOW

Give me the two groups. I'll run the test and call it.

I show my work, and we iterate.

No black box, and no one-shot. I share the exploration first, the missing values, the oddities, the data issues, as a page you can open right in your browser. We fix and refine, then build the model, then tune it together.



EDA & missing values

open it in your browser

we refine together

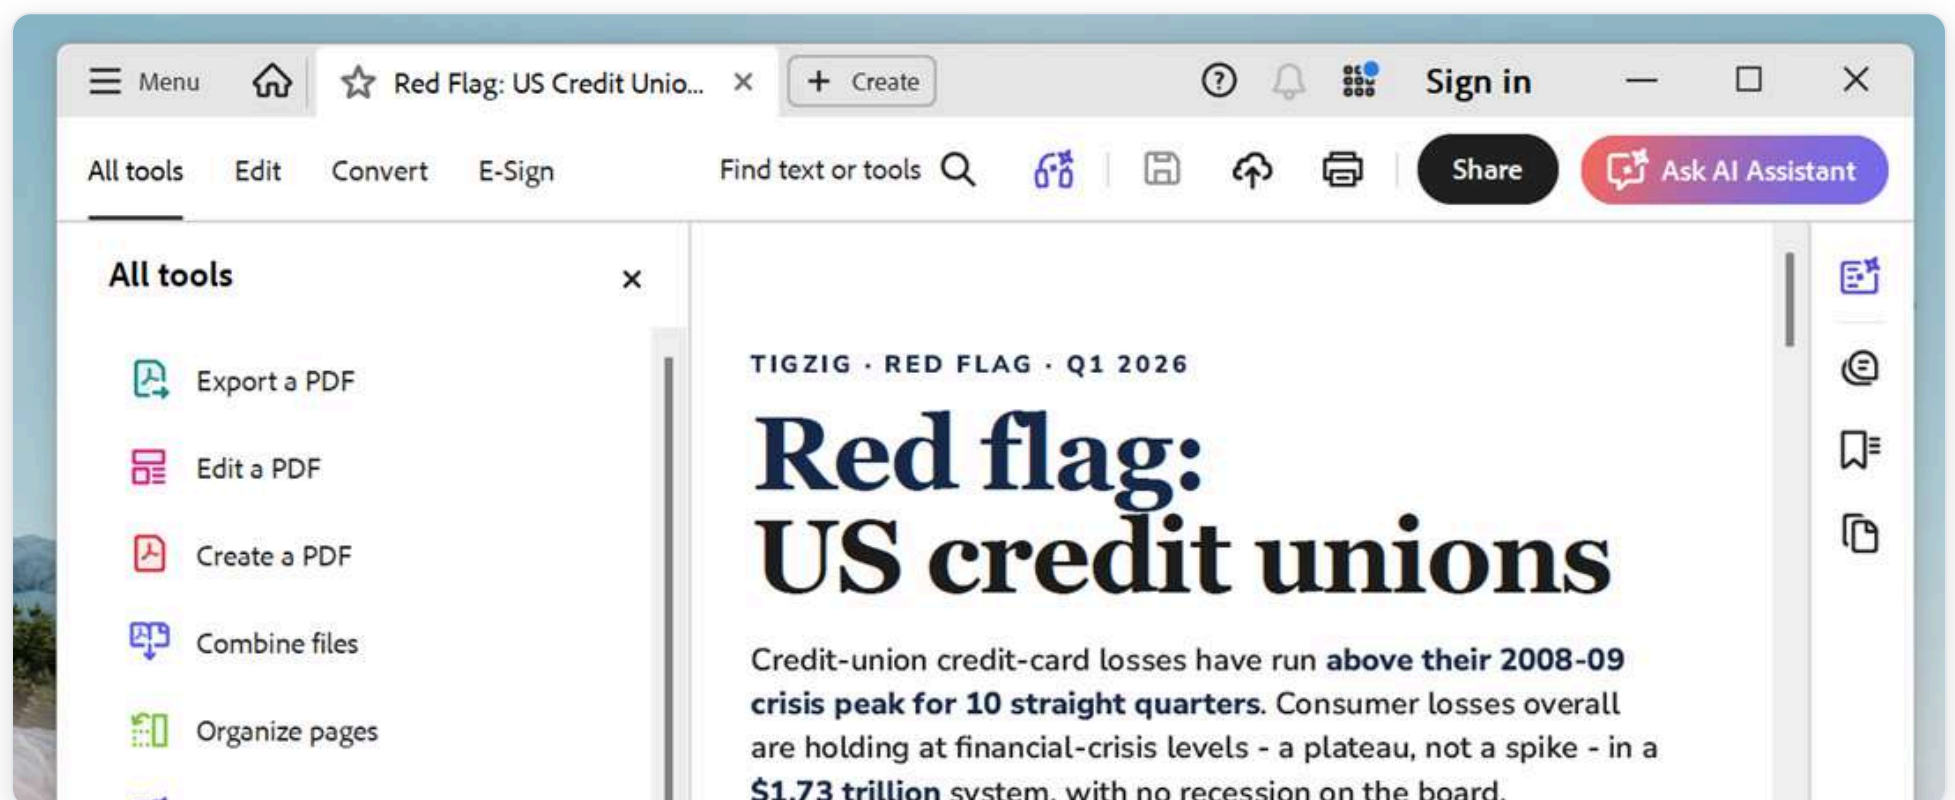
HOW

Ask me to share the exploration as a page you can open in your browser.

I package it the way you need to send it.

A clean one-pager or a multi-page report, then out as whatever the moment needs, including a **PDF with clickable links**.

<> HTML 📄 PDF + links 📄 Word 📁 PowerPoint 📊 Excel



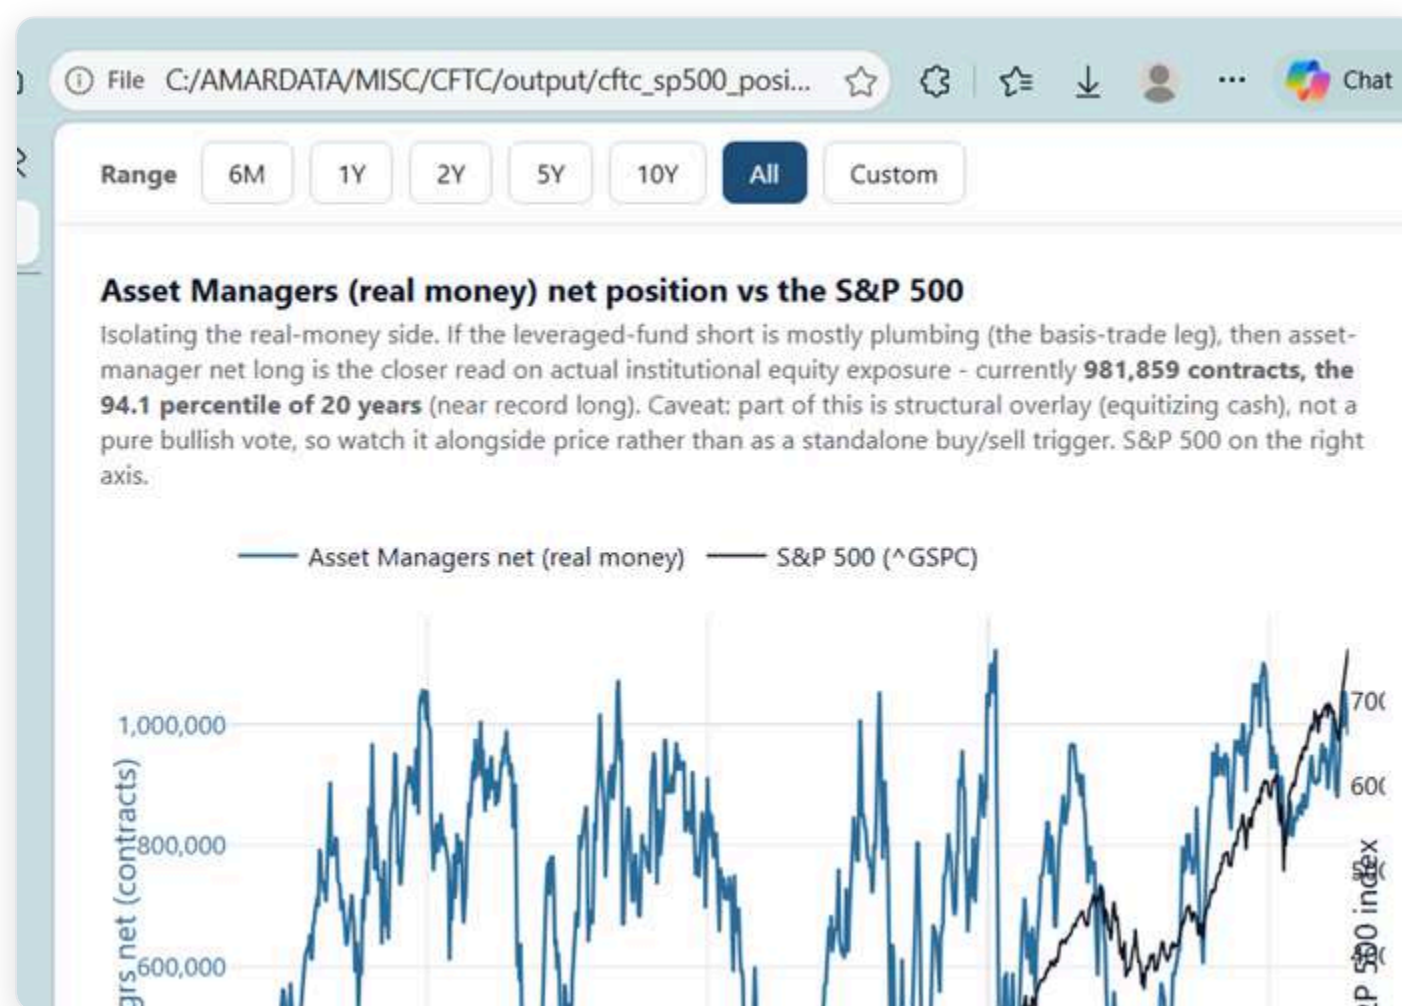
A report I built, exported to a PDF and open in Acrobat, ready to send.

HOW

Just say the format: HTML, PDF, Word, PowerPoint or Excel.

I'll build you a live dashboard.

Not just a static report, a page you reopen tomorrow and it **pulls fresh data each time**. A tracker, a metrics page, a queue you keep coming back to.



A live dashboard I built: CFTC asset-manager positioning against the S&P 500, it pulls fresh data each time you open it.

🔄 **live tracker** 📈 **metrics page** 🔄 **refreshes itself**

HOW

Ask for "a live page I can reopen". It pulls fresh data each time.

I read the real web page, not a snippet.

Give me a link and I open the actual web page and fetch the whole thing with my own tools, the **real text**, not just a short search summary. So I can quote it properly and cut down on guesswork.



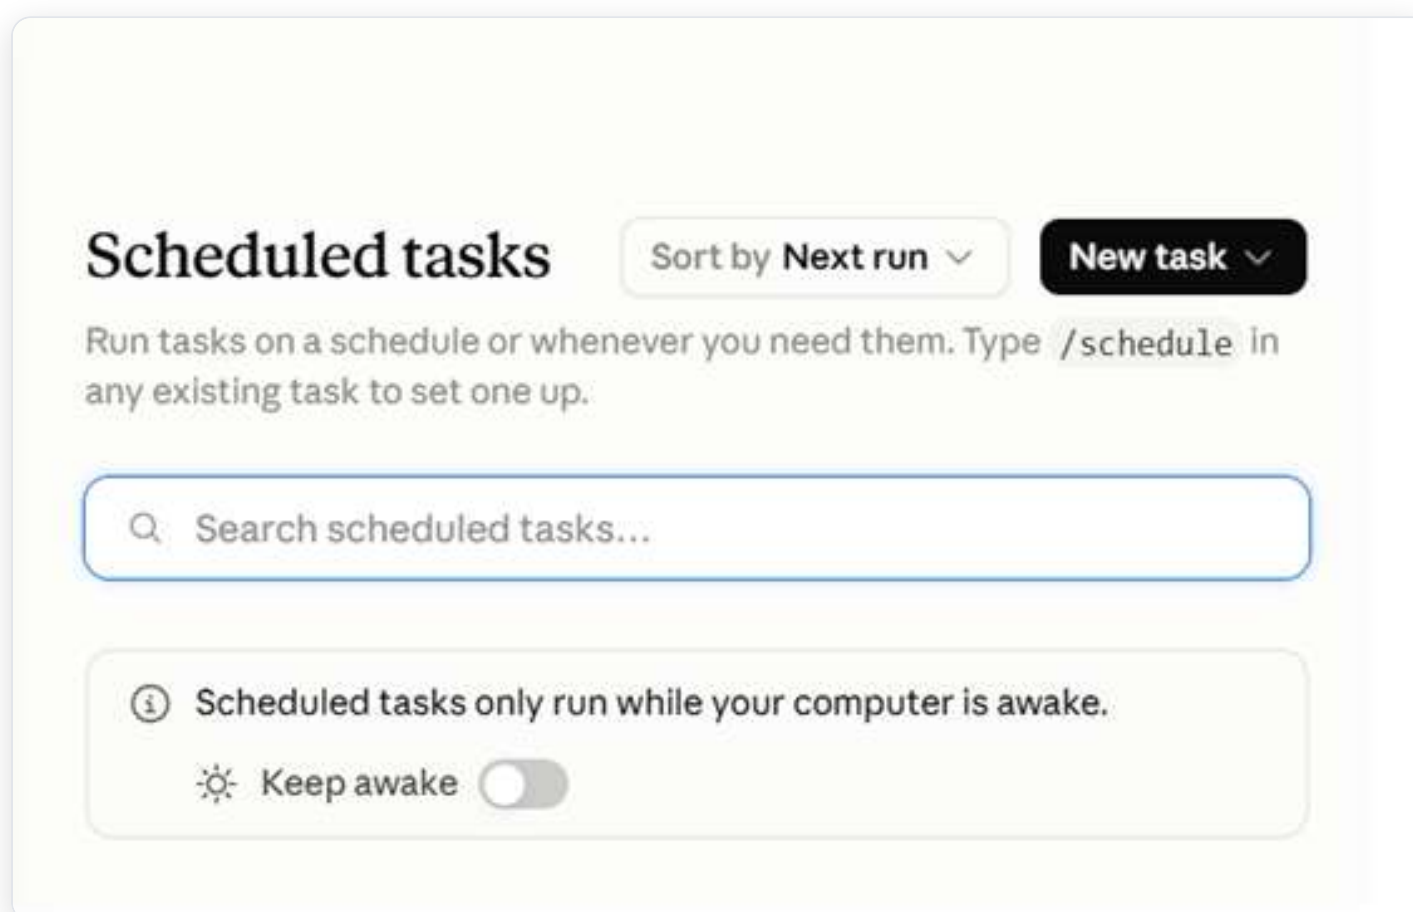
 **open the real URL**  **read the whole page**  **less guesswork**

HOW

Give me the link or the topic. I'll open the page and read it.

I run on a schedule.

Set it once and walk away. A morning briefing, a weekly data refresh, a Monday status check, I'll run it on time and have the result waiting for you.



The Scheduled tasks panel: set a cadence and walk away (it runs while your computer is awake).

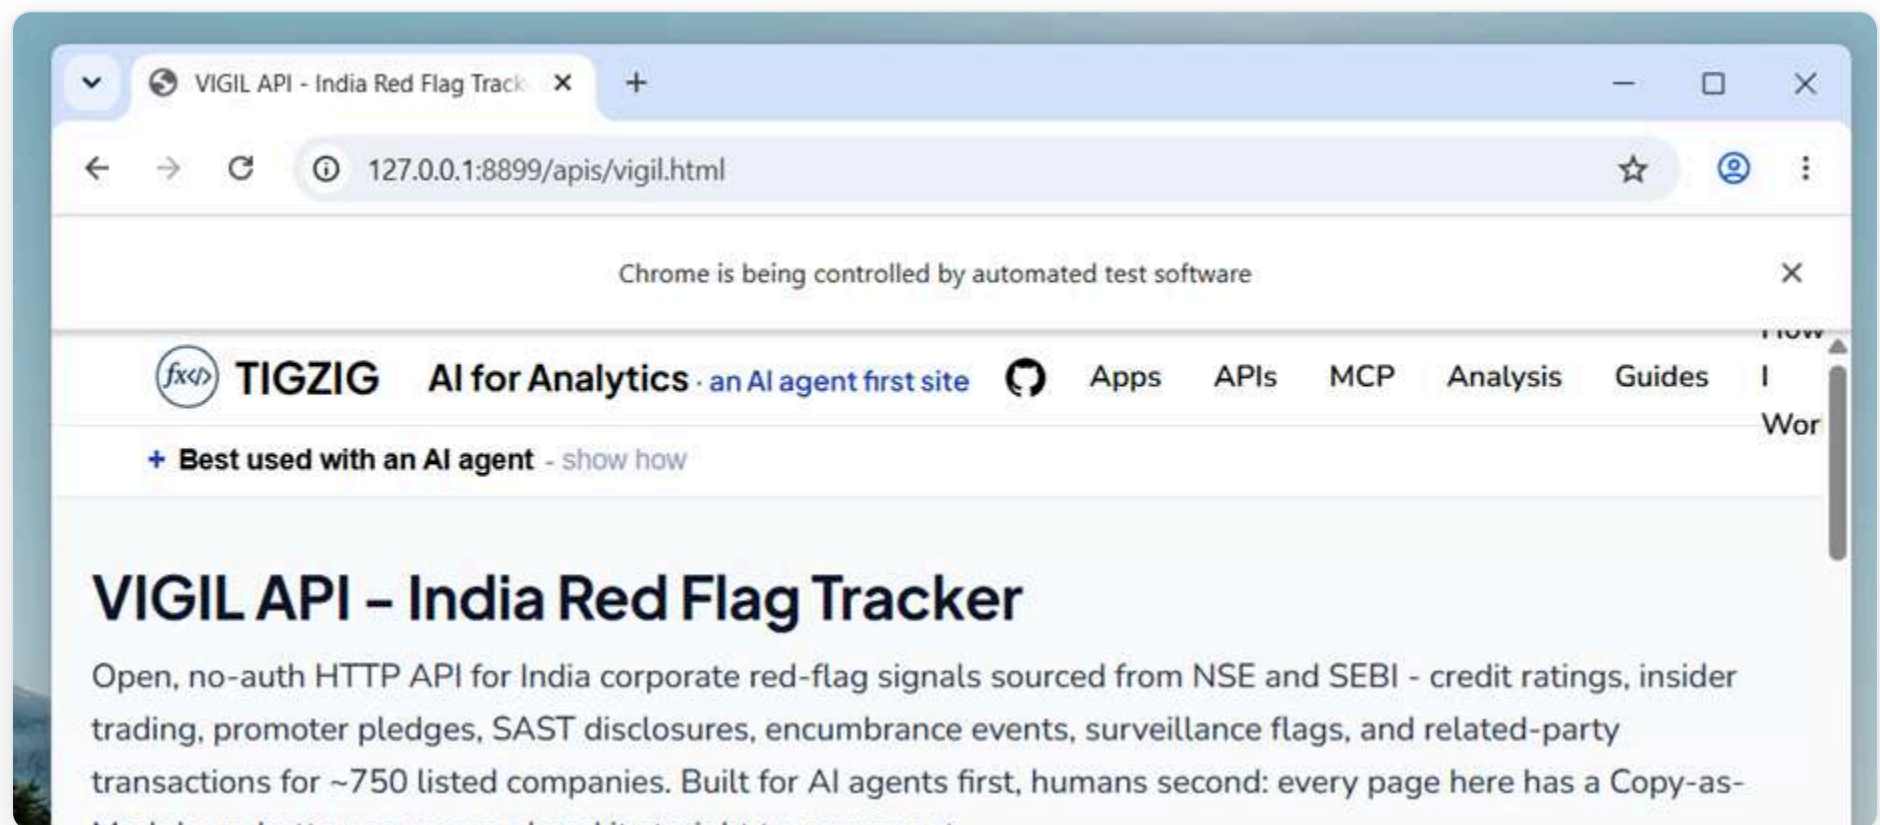
🕒 **daily briefing** 🔄 **weekly refresh** 📄 **month-end report**

HOW

Say "every morning at 7" or "every Monday", and I'll set it up.

I'll drive your browser and desktop.

When a tool has no tidy API, I'll use it the way you would, clicking through a **website** or a **desktop app** to get the thing done.



🌐 web apps 📁 desktop apps ✓ no API needed

HOW

Ask me to open a site or app. I'll connect to your browser when you're ready.

I remember you.

Tell me your style once, your fonts, your colours, the way you like reports, and I carry it forward. Point me at a preferences file in your project, or just say "remember this".



↑↓ your style 📄 a preferences file ↻ across sessions

HOW

Point me to your preferences file, or say "remember this", and I apply it next time.

What's in my toolkit.

PYTHON POWERS

A private Python sandbox - roughly 4GB of memory and 10GB of disk to work in.

Install almost anything from PyPI - the whole Python data and ML stack is open to me.

Data: pandas, polars and pyarrow, to wrangle files far bigger than memory.

Models & charts: scikit-learn for machine learning, matplotlib for the visuals.

Images and more: OpenCV and Pillow for image work, and since almost any package installs, audio and other toolkits are a step away.

The fine print: the very biggest packages (XGBoost, TensorFlow, up to ~600MB) can run past the current 45-second install window and not finish; most install in seconds.

CONNECTORS

Beyond the sandbox, I plug into your tools, Gmail, Drive, Slack, Notion, and a growing list of data providers, whatever's been made available. And for your own systems, custom connectors through MCP servers.

What I can't do (and the handoff).

I can't run code directly on your laptop, and I won't move money or place trades. For code that has to run on your machine, my teammate **Claude Code** handles it: I write the instructions, it executes them.



⊘ no code on your machine ⊘ no trades or transfers

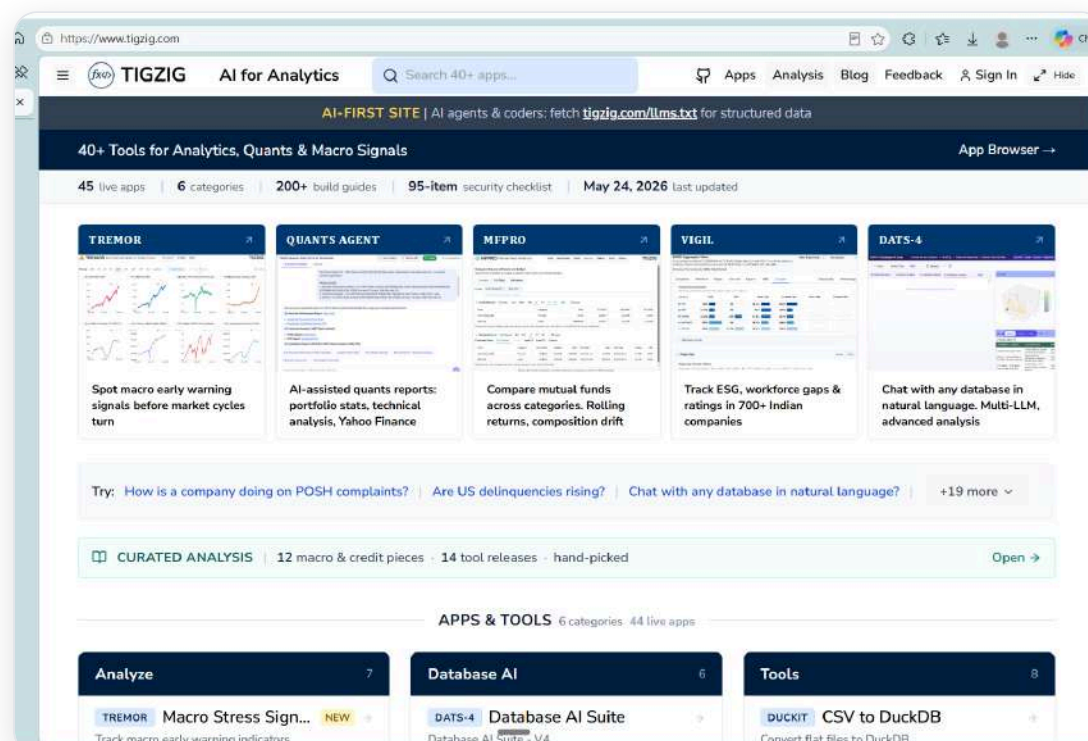
⇔ Claude Code, same team

HOW

Claude Code sits right there, in the desktop app or in VS Code. Hand it the note and it runs on your machine.

How Amar works with me.

Amar runs an AI-agent-first shop at **tigzig.com**. He points me at data, on Tremor and public APIs, I draft, analyse and build, and Claude Code executes on his machine. I've built him HTML one-pagers, PDF reports and models.



tigzig.com, the analyst's tool shed: analytics apps, APIs and datasets, AI-agent ready.

HOW

Visit tigzig.com for free data and APIs. Want me to explore it? I can reach the site directly from Cowork.

Resources

Everything above runs in Claude's Cowork mode. The guides, and a couple of good reads:

Cowork, the product

claude.com/product/cowork

Getting started with Cowork (best practices)

claude.com/blog/best-practices-for-getting-started-with-claude-cowork

Why HTML beats Markdown, by Thariq Shhipar (Anthropic)

claude.com/blog/using-claude-code-the-unreasonable-effectiveness-of-html

Let Claude use your computer

support.claude.com

tigzig.com

the analyst's tool shed · AI-agent first

40+ live apps · 200+ build guides · security checklist · MCP & API servers.

Point Claude Code or any AI agent at it, and it picks up the work for you.

Amar Harolika · Decision Sciences & Applied AI

By the way, I made this too

**This whole deck was built in Cowork as HTML,
then exported to a PDF. Same tools, same desk.**

Claude Cowork, with Amar Harolika · tigzig.com