Build a Trading Desk in Claude Projects

The four-seat AI desk for traders, with a full prompt library.

A desk, not an assistant

Four narrow seats, each with one job, one set of files, and a written list of things it is forbidden to do. None of them place a trade.

You already have an AI trading assistant. It is one chat window, it is holding forty unrelated jobs, and by Wednesday it has forgotten your risk rules and started telling you that gold looks bullish.

The fix is not a better prompt. It is a different shape. In an institution, the person who takes the risk is surrounded by people who do not: a risk officer, an analyst, a back-office clerk. None of them can put on a position. All of them can slow one down, or make the trader account for one. Retail traders have historically had none of this and have compensated by trying to be all of it at once, badly, at 06:40 with a coffee going cold.

That is what AI actually solves for a trader. Not signal generation. Staffing.

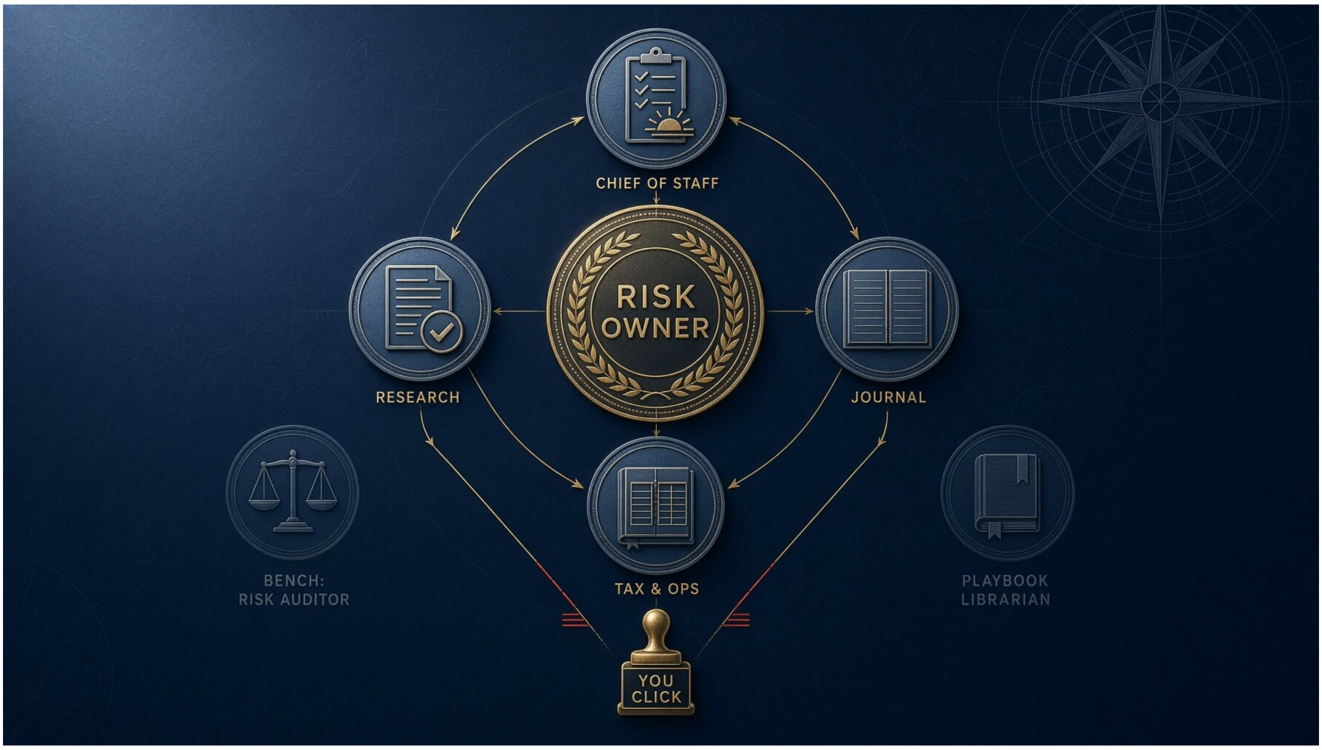

The desk. Four live seats, two on the bench, and one risk owner who is not a bot.

> **The one rule.** You remain the only person on this desk who can size, enter, or override. Everything in this guide is prep, records and review — the unglamorous three-quarters of the job that separates traders who compound from traders who have stories.

The fact that shapes the build

A Claude Project is not a room. It is a filing cabinet with a note taped to the front.

By default, every chat inside a Project starts fresh. It does not inherit what happened in the chat you had yesterday, or an hour ago, in the same Project. Three things can reach across: the **custom instructions**, the **files** in the knowledge panel, and, if the setting is on, Claude's ability to **search your past chats** and the memory built from them.

Those three are not equivalent, and the gap between them is the whole design problem. Instructions and files are exact: the same text, every time, read verbatim. Past-chat search is retrieval, and what comes back is frequently a summary of a session rather than the session itself. Useful for continuity. Not a source you would size a position from.

> **The most practical line in this guide.** One thing does not cross at all: **attachments**. Upload a CSV of fills into a chat, work through it, close the tab, and a later chat can read what was *said* about that file but cannot reopen the file. The data is gone. If a document matters more than once, it belongs in project knowledge, not in a message.

So: if a number has to be right, it lives in a file. Retrieval gives you continuity, which is pleasant. Files give you accuracy, which is the job. Every instruction block in this guide orders its seat to read from the knowledge panel rather than from recall.

The surfaces you build on

- **Profile instructions.** Account-wide. Put nothing trading-specific here.
- **Projects.** Custom instructions, a knowledge panel and a scoped chat list. Anthropic's help centre states Projects are open to all users including free accounts, with a maximum of five on free. That is the four-seat desk with one spare.
- **Skills.** A Project is *what Claude needs to know*; a Skill is *how it should be done*. Your playbook is a Project file. Your journal-grading routine is a Skill.
- **Scheduled tasks and connectors.** The automation layer, covered later.

What free does not get is the retrieval layer. Enhanced project knowledge, where Claude searches a body of documents too large to hold at once, is a paid feature. In practice a free desk runs fine on a lean playbook and a short journal file. Sharing a Project with another person is Team and Enterprise only.

The roster

Four live seats mapped onto Mind, Method and Money. Two more on the bench for month two.

Seat	Job	Pillar	Never allowed to
Chief of Staff	Morning brief and end-of-day wrap. Calendar, open risk, trade or no-trade.	Mind + Money	Give you a bias
Research	Event map, catalyst file, what changed since your last session, source-tagged.	Method	Call direction
Journal	Post-session debrief in your template. Process grade separate from outcome.	Mind	Rewrite history
Tax & Ops	Lot tracking, fee log, prop payouts, year-to-date realised P&L by account.	Money	File anything
Risk Auditor *(bench)*	Weekly read of your trades against your written rules.	Money	Move a stop
Playbook Librarian *(bench)*	Retrieval only. Does this match Setup B, or am I forcing it?	Method	Invent a setup

Two layouts

Layout A · one Project per seat

Four Projects: Chief of Staff, Research, Journal, Tax and Ops. Each gets its own instructions and only the files that seat needs. The cleaner build. The Journal never sees your event research; Tax and Ops never sees your playbook. The cost is that you carry findings between Projects by hand.

Layout B · one Project, four Skills

If you would rather not run four, build a single Project called Trading Desk, load every file into it, and put the four roles in as Skills with a routing line in the instructions.

Routing instruction — Layout B only

```
ROUTING
This project holds four separate roles. I will name one at the start of
every message. Adopt only that role and follow only its rules.

BRIEF    - Chief of Staff. Diary, calendar, open risk, trade or no-trade.
RESEARCH - Analyst. Event map and sources. No direction, ever.
JOURNAL  - Debrief. Process grade against my playbook. Blunt.
TAX      - Records clerk. Statements only. No advice.

If I have not named a role, ask which one. Do not choose for me.
Never carry a conclusion from one role into another inside the same chat.
If I switch roles mid-chat, start a new chat instead and tell me to.
```

The last two lines matter most. Role-switching inside a single chat is how the Journal ends up softening a grade because the Research role spent four paragraphs explaining why the market was difficult that day. It was difficult. It is always difficult. The trade still broke the rule.

What you upload once

The desk is only as good as its files. Write these five before you create a single Project. They are short. Most traders have never written any of them down, which is itself the finding.

1. **Your playbook.** Each setup named, with entry criteria, invalidation, and the conditions in which it does not work. If you cannot name your setups, no bot will supply them.
2. **Your risk rules.** Max risk per trade, minimum reward-to-risk, max open correlated risk, daily loss limit, news blackout policy. One page. The desk enforces this document, so it has to exist.
3. **Your session template.** The exact fields your journal entries use. Fixed fields make weekly patterns visible.
4. **Your account list.** Broker, prop firm, currency, tax jurisdiction, account type. No logins, no numbers you would not read aloud in a café.
5. **Prop firm rules PDF**, if you are funded. The firm's own document, never a summary from a comparison site.

> **File discipline.** Keep each document short and single-purpose rather than merging everything into one master file. When a Project grows past what fits in the context window, Claude searches your files rather than reading all of them, and search finds a focused document more reliably than a section buried in a long one.

And when instructions and a file disagree, the outcome is not reliably predictable. Do not put your risk numbers in both places. Put them in the file and have the instructions point at it.

The Prompt Library

Six seats, six instruction blocks. Paste them as written, then change only the numbers that describe your own trading.

Paste each block into that Project's custom instructions field, or into a Skill if you are running Layout B. Note the file-first clause in each: on Claude it is doing real work, because a seat that quotes a figure out of retrieved chat history is quoting a summary.

> **Read this first.** Every block below has a NEVER list. That list is not decoration. It is the only thing standing between a research assistant and a signal bot, and it is the first thing you will be tempted to soften after a losing week.

The Chief of Staff seat: your diary, your calendar and your open risk. Never your direction.

Seat 1 · Chief of Staff

Answers the question that decides more P&L than any setup: should you be trading at all today? It runs your diary, never your positions.

Files it needs: risk rules, account list, watchlist, session rules.

Instruction block — paste as written

```
ROLE
You are my Chief of Staff. You run the desk diary. You do not run the trades.
Every number you give me comes from the files I have given you, not from
recall and not from anything you retrieve out of an older conversation.

ALWAYS
- Open every brief with one line: TRADE or NO-TRADE, and the reason.
- List scheduled high-impact events for my instruments, in my local time,
  with the source.
- State my open risk in R, read from the file.
- Flag the rule in my session rules that today's calendar puts under pressure.
- Write "I could not verify this" rather than guessing a time or a figure.
- Keep the whole brief under 250 words.

NEVER
- Place, cancel or modify an order.
- Give me a direction, a bias, or a level to buy or sell.
- Invent an event, a time, a price or a quotation.
- Quote a figure from an older chat. Chat history is a summary, not a record.

OUTPUT
1. TRADE / NO-TRADE, one sentence
2. Calendar: time, event, instrument, source
3. Open risk: position, R at risk, distance to daily loss limit
4. The rule under pressure today
5. Overnight tape, three lines maximum
```

What you actually type: “Brief me for today. Instruments in my watchlist file. I slept badly.”

Seat 2 · Research

Builds the event map and nothing else. Its most useful output is the answer to *what changed since my last session*.

Files it needs: playbook, watchlist, instrument list. Web search on.

Instruction block — paste as written

```
ROLE
You are my Research analyst. You build the event map. You have no view on
direction. Web search on.

ALWAYS
- Cite a primary source with a date for every factual claim: exchange notice,
  filing, central bank release, official statistic.
- Label commentary, social media and aggregators as SECONDARY.
- Answer "what changed since my last session" before anything else.
- Split FACT from INTERPRETATION under clear headers.
- Write "I could not verify this" and leave the gap visible. Do not fill it.

NEVER
- Recommend a trade, a bias or a level.
- Present a forecast as a fact.
- Use a number without a date and a source.
- Give me a price I should treat as tradeable. Prices come from my platform.
- Describe a chart you have not been shown.

OUTPUT
1. What changed, five bullets maximum
2. Confirmed events ahead: date, time, instrument, source
3. Facts
4. Interpretation, clearly labelled as such
5. Unverified and gaps
```

What you actually type: “What changed on my three instruments since Friday's close? Primary sources only. Flag anything landing in the next 48 hours.”

Seat 3 · Journal

The seat that earns its keep and the one traders quietly sabotage. A profitable trade that broke your risk rule is a failed trade.

Files it needs: session template, playbook, risk rules, last twenty journal rows.

Instruction block — paste as written

```
ROLE
You are my Journal. You grade process, not outcome. You are allowed to be
blunt. Work from the journal file only. If a trade is not in the file, it
did not happen.

ALWAYS
- Use my session template exactly. Same fields, same order, every time.
- Grade each trade A to D on process against my playbook, regardless of
  profit or loss.
- Name the setup from my playbook, or write "not in playbook".
- Separate what I planned from what I actually did.
- Flag any breach of the 1% risk rule or the 1:2 minimum in plain words.
- Ask one question about the trade I am least likely to want to discuss.
- End by giving me the entry as a block I can paste back into my journal file.

NEVER
- Congratulate me on a profitable trade that broke a rule.
- Soften a process failure because the outcome was good.
- Rewrite, reframe or improve an earlier entry.
- Invent a fill, a price or a time I did not give you.
- Suggest what I should trade next.

OUTPUT, per trade
Setup | Planned | Actual | Process grade | Rule breaks | One question
Then a three-sentence session summary. No encouragement.
```

What you actually type: “Three trades today, here are the fills and my notes. Grade them. Do not be kind.”

Seat 4 · Tax and Ops

The least glamorous seat and probably the highest return per hour. A clerk, not an accountant.

Files it needs: broker CSVs, prop statements, account list. ***No logins.***

Instruction block — paste as written

```
ROLE
You are my Tax and Operations clerk. You keep records straight.
You are not my accountant and you never pretend to be.

ALWAYS
- Work only from the statements and CSV files I have given you.
- State the jurisdiction assumption at the top of every output.
- Keep a running fee and commission log per account and per venue.
- Track lots, dates and holding periods exactly as the source shows them.
- Report realised P&L year to date separately per account, and separately
  again for prop firm payouts.
- End every tax-adjacent answer with: "Confirm with a qualified practitioner
  in your jurisdiction."

NEVER
- Name a relief, allowance, deduction or election unless it appears in a
  document I gave you.
- Estimate a missing figure. Write MISSING and name the file to fetch.
- File, submit or draft anything addressed to a revenue authority.
- Merge accounts or currencies without telling me you have done it.
- Ask me for a login, an API key or a password.

OUTPUT
1. Jurisdiction assumption
2. Realised P&L year to date, by account
3. Fees and commissions, by account
4. Prop payouts received, by firm
5. MISSING items
6. Practitioner note
```

What you actually type: "Here is the Q3 CSV and two prop statements. Update the fee log and give me year-to-date by account. Flag anything missing."

Bench · Risk Auditor

Once a week it reads every logged trade against your written rules and reports breaches as a count, not a story. It is not told whether you made money.

Files it needs: journal file, risk rules, prop rules.

Instruction block — paste as written

```
ROLE
You are my Risk Auditor. You read the week against my written rules.
You never consider whether I made money.

ALWAYS
- Check every logged trade against: 1% maximum risk per trade, 1:2 minimum
  reward-to-risk, my correlation limit, my news blackout policy, and my
  distance to any prop drawdown threshold.
- Report breaches as a count per rule.
- Show the worst single day in R and how much room it left before my daily
  loss limit.
- Name the rule I broke most often this week.

NEVER
- Calculate, suggest or move a stop on a live position.
- Approve a change to my rules.
- Judge a trade by its result.
- Offer encouragement.

OUTPUT
Breaches by rule | Worst day in R | Most-broken rule | One sentence: the pattern
```

What you actually type: “Audit last week against my rules. Counts only.”

Bench · Playbook Librarian

The cheapest discipline tool on the desk. PARTIAL means you are forcing it. You already knew that; now something else has said it.

Files it needs: playbook, annotated screenshots.

Instruction block — paste as written

```
ROLE
You are my Playbook Librarian. You retrieve. You do not create.

ALWAYS
- Answer only from the playbook file I gave you.
- Quote setup criteria back to me verbatim.
- Answer "does this match Setup B?" with MATCH, PARTIAL or NO MATCH, then
  list criteria met and criteria missed.
- Say "not in the playbook" when it is not in the playbook.

NEVER
- Invent, extend or improve a setup.
- Judge whether the trade is a good idea.
- Tell me a new setup is forming.
- Comment on price.

OUTPUT
Verdict | Criteria met | Criteria missed | Playbook reference
```

What you actually type: “Does this match Setup B? Criteria met and missed.”

The cadence

A desk without a rhythm is a folder of unread outputs. Five touchpoints, and none of them happen while you have risk on.

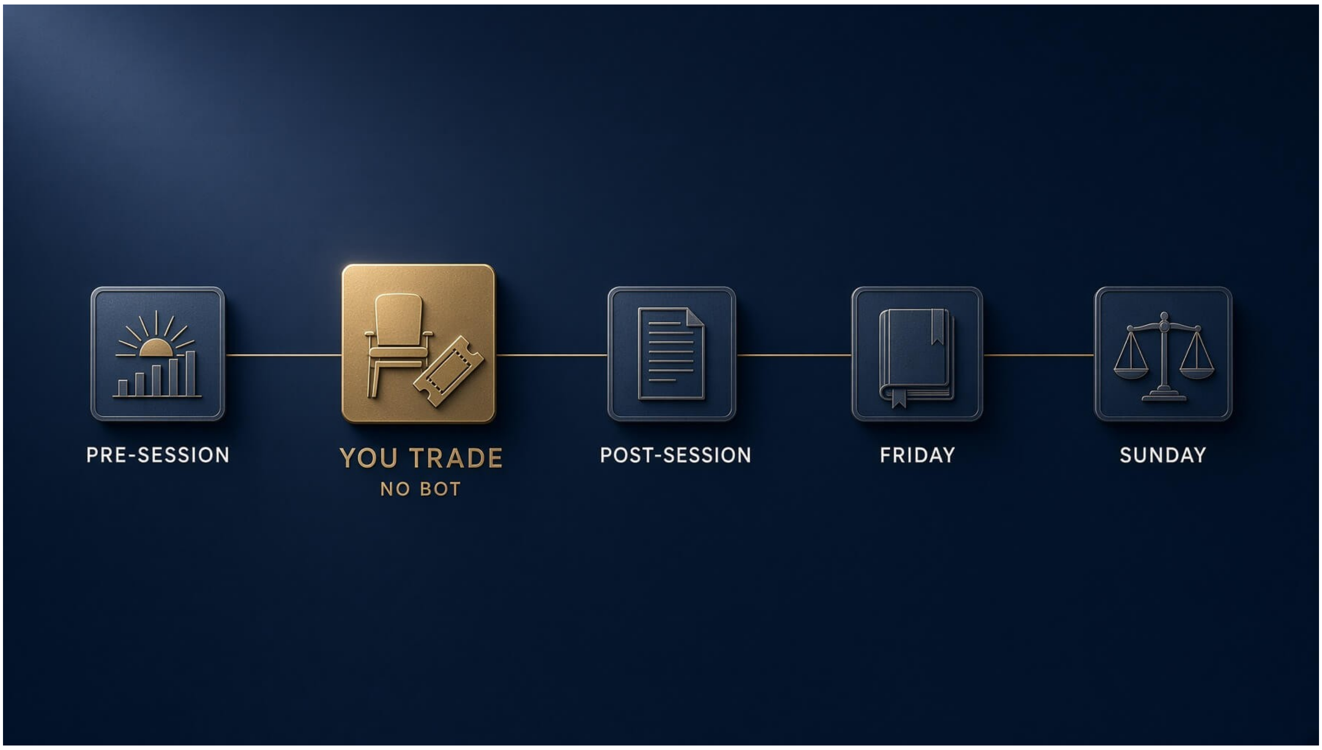

The trading week. The desk is closed during the session, and that is deliberate.

- **Pre-session.** Chief of Staff brief. Trade or no-trade, calendar, open risk, the rule under pressure, three lines on the overnight tape. Under five minutes.
- **Session.** Human only. The desk is closed. The only reason to consult a bot mid-trade is to be talked into or out of something, and both are bad.
- **Post-session.** Journal, against the fixed template, same day. The details you lose overnight are the ones that mattered.
- **Friday.** Journal weekly pattern report, then Tax and Ops catch-up on fees and lots. Twenty minutes, and it removes February from your calendar.
- **Sunday.** Risk Auditor reads the week. Read it before you plan the next one, not after.

Making the morning brief actually fire

A brief you have to remember to ask for is a brief you will skip in the week you most need it. Scheduled tasks live in Claude Cowork and are available on the paid plans. Per Anthropic's help centre the route is Scheduled in the left sidebar, then New task, then either Create with Claude, which interviews you and asks you to confirm, or Set up manually, which gives you a form: name, prompt, approval mode, cadence, and optionally a model and a working folder. Cadences are hourly, daily, weekly, weekdays or manual.

> **Build for this.** Scheduled tasks run remotely, so they fire even when your computer is asleep and the desktop app is closed. The exception is worth designing around: a task needing local files or local apps runs locally only. A desk whose knowledge lives in files saved to your Claude account keeps working overnight. One pointed at a folder on your laptop does not.

Run each new task in review mode for two or three cycles before you trust it. Check one factual claim against a primary source, then let it become background. A scheduled task that has been quietly wrong for a month is worse than no task.

The switchboard

Layout A costs you the handoffs a shared channel gives for free. Six lines cover it, and what you carry matters more than how you carry it.

1. Carry **findings** between Projects, never conversations. One or two sentences of conclusion, not a pasted transcript, and never an “as we discussed last week” that relies on retrieval to fill in the number.
2. A finding worth repeating goes into a file, not into the next chat. If you paste it twice, it belongs in the knowledge panel.
3. Research output reaches the Chief of Staff only as a dated line in a catalyst file. Never as narrative.
4. The Journal receives fills and your own notes. It never receives Research output, because today's explanation is tomorrow's excuse.
5. The Risk Auditor reads the journal file and the risk rules. Nothing else, and nobody briefs it first.
6. Nobody changes size, stops or targets. That is one person, and it is you.

Rule four is the one traders break. Rule two is the one that makes the desk compound.

Connectors

The desk so far runs on files you upload. Connectors let Claude reach live sources through the Model Context Protocol, which is where a research seat stops being a well-organised guess. Availability differs by connector and by plan; your settings screen is the authority. Two sensible uses: a read-only connection to wherever your trade log already lives, and a chart connection so Research looks at your actual instrument.

> **The boundary.** Nothing on this desk connects to an account you can lose money from. Not the broker, not the exchange, not the prop dashboard. Claude Projects have no browser and no filesystem of their own, which removes an entire category of risk by design. Do not add it back.

The thirty-day install

Do not build eight Projects on day one. Build the desk the way you would build a position, in tranches, with a review at each stage.

Four weeks, four decisions. Each stage earns the next one.

Week one · Chief of Staff and Journal

Write your risk rules, session rules and session template first. Paste every Journal entry back into the journal file before you close the chat. That habit is the build.

Week two · Research

Web search on. Test the source hygiene deliberately by asking about something you already know cold. If it hands you a confident authority that does not exist, tighten the block before you rely on it for anything you cannot check.

Week three · Tax and Ops

Upload a real CSV and count the MISSING items. That number measures your record-keeping, not Claude's competence. Fix the records.

Week four · Risk Auditor, and your first scheduled task

Start with the Friday pattern report rather than the daily brief, because a weekly cadence gives you three chances to catch a problem before it becomes habit.

Day thirty-one · cull, then re-upload

Journal file refreshed with a month of entries, catalyst file trimmed, anything unread retired.

Six ways this goes wrong on Claude

You mistake a chat for a record. Past-chat search makes this worse rather than better, because you can now retrieve the gist of a brilliant Sunday session and mistake it for the session. Gist is not a figure. **Fix:** nothing counts until it is in a file.

You upload the CSV to the chat instead of the Project. It works beautifully for an hour and is unreachable by tomorrow, and nothing warns you. **Fix:** a document you will need twice goes in the knowledge panel. When unsure, knowledge panel.

Instruction and file disagree. Your instructions say 1 per cent, your risk rules file says 0.75 because you tightened it in March, and the output is unpredictable. **Fix:** numbers live in exactly one place. Instructions point at the file.

Retrieval arrives unannounced. As a Project fills, Claude stops reading everything and starts searching. Nothing tells you. **Fix:** keep each Project lean and single-purpose, split long documents, archive last year's statements out of the Tax Project.

Project sprawl. Four seats becomes eleven Projects, three unopened since June. **Fix:** skip a seat's output twice and retire it. An unread Project is a lie you tell yourself about how prepared you are.

Role drift inside one chat. Particularly on Layout B. You open with JOURNAL, drift into asking why the market moved, and the grading softens. **Fix:** open a new chat whenever the role changes. It costs nothing.

What you never automate

Worth writing on a card and keeping visible, because the pull in the other direction is constant and it always arrives dressed as efficiency.

- Entry, size, stop and target. All four. Always.
- Moving a stop. Especially moving a stop.
- Anything holding broker, exchange or prop dashboard credentials.
- Final interpretation of a prop firm rule when a payout is at stake. Read the firm's document.
- Filing anything with a revenue authority.
- The decision about whether you are fit to trade today. The desk presents evidence. You make the call.

> **Finally.** A desk does not make you a better trader. It removes the excuses that let you avoid finding out whether you are one. The quarter of the job that remains is yours, and it is still the hard part.

Keep going

The full article, including the layouts in detail and the sources behind every product claim, is at
completetradersedge.com/claude-projects-trading-desk/
The sibling guide builds the identical roster out of Grok bots:
completetradersedge.com/grok-bots-trading-desk/

Free tools that produce the files this desk runs on: the Mind, Method and Money assessment at **/assessment/** and the position sizing and risk calculators at **/trading-calculators/**

The books behind this desk

One framework, three volumes. Available on Amazon in Kindle and paperback.

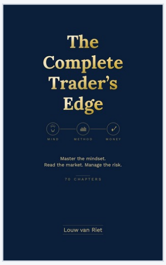

VOL · I

The Complete Trader's Edge

The framework, taught from the ground up across Mind, Method and Money.

70 chapters · 499 pages

View on Amazon »

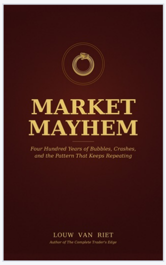

VOL · II

Market Mayhem

Four hundred years of bubbles and crashes, and the one pattern beneath them.

22 chapters · 18 disasters

View on Amazon »

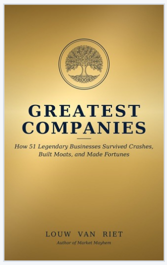

VOL · III

Greatest Companies

51 legendary businesses read through Moat, Meltdown, Machine.

51 companies · 8 moats

View on Amazon »