CoTrader logo

COTRADER
FREE GUIDE · 2026 EDITION

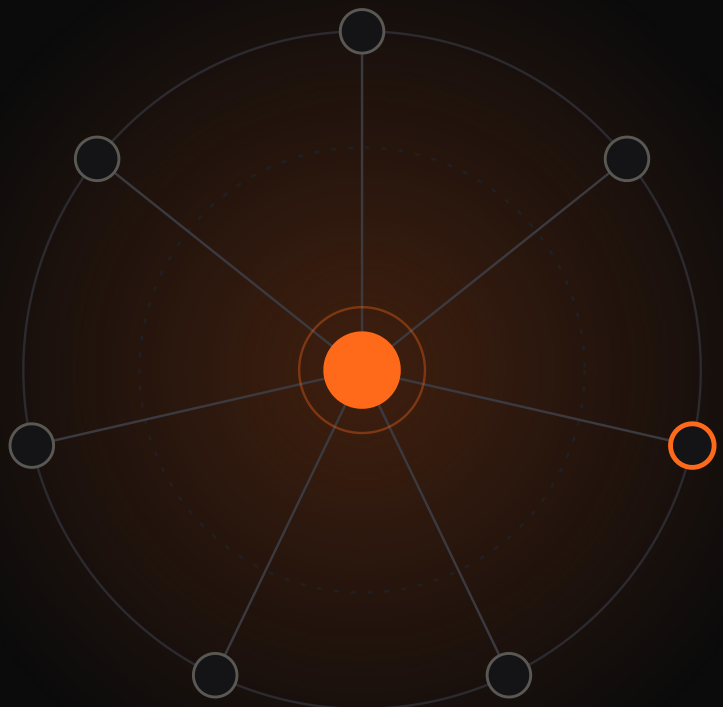

The 7-Agent Trading *System*

How a team of AI agents researches, checks and executes every trade, and why the most important agent is the one that can say no.

Written by **Prince**, founder of CoTrader

cotrader.pro

INSIDE

Contents

01 Why most traders fail 03

02 What an AI agent team is, and is not 04

03 The seven agents 05

04 How one trade flows through the team 08

05 The risk layer 09

06 What the research says 11

07 How to start 12

A NOTE BEFORE YOU START

Most AI trading content sells prediction. This guide is about something less exciting and far more useful: **process**.

I build automated trading systems for funded traders and private capital. The systems that hold up have one thing in common. The AI does the reading, the structuring and the arguing. Plain code holds the line on risk. Nothing reaches the market until every check has passed.

Read this in one sitting. By the end you will know exactly how a team of agents works, where it helps, where it does not, and whether to build your own or have one built.

Prince

01 · THE REAL PROBLEM

Accounts rarely die from a bad strategy. They die from *one bad day.*

Ask anyone who has lost a funded account what happened. It is almost never "my edge stopped working". It is usually one session where the plan went out of the window.

One trade too many

A loss, then a revenge trade, then a bigger one. The strategy was fine. The fourth trade was not part of it.

A single rule break

Prop accounts close on one breach of the daily loss or drawdown limit, no matter how good the rest of the month was.

Emotion at the worst moment

Moving the stop "just this once". Taking a setup that only half matches. Skipping the trade that would have worked.

No written record

Without a journal, the same mistake repeats. You cannot fix a pattern you never wrote down.

ILLUSTRATIVE · NOT REAL ACCOUNT DATA

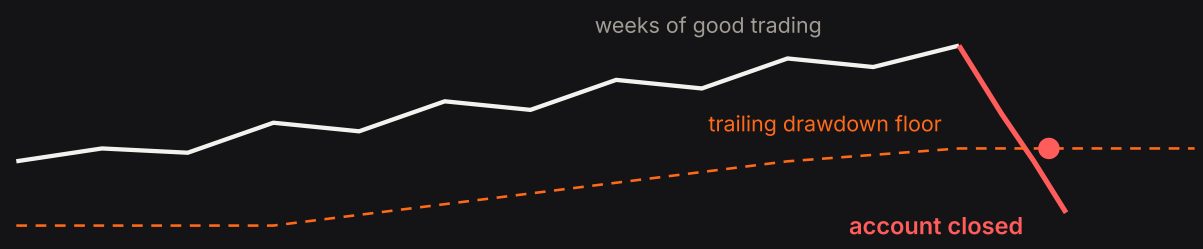

> A funded account is a rules game first, and a prediction game second.

That is why this guide starts with discipline, not signals. An AI agent team is useful mostly because it makes the boring parts impossible to skip.

02 · PLAIN ENGLISH

What an AI agent team *actually* is

An AI agent is a language model given **one job**, the tools to do it, and a fixed format to report back in. A team is several agents with narrow jobs, passing structured notes to each other, with a coordinator deciding what happens next.

Think of a small, well-run trading firm. Analysts read the market. A strategist checks the setup against the playbook. A risk officer can veto anything. An execution trader places the order. Someone keeps the records. A head of trading makes sure every step happened in order. Research frameworks such as *TradingAgents* are built exactly this way, with separate analysts, a bull-versus-bear debate, a trader, a risk team and a final approver.

The difference with agents: every step runs the same way, every time, and every decision is written down.

WHAT IT IS

- A process that runs your rules the same way on every trade
- A second and third opinion before any order is sent
- A hard risk layer that can block a trade, written in code
- A complete record of what was decided, and why
- Alerts to your phone so you stay in control

WHAT IT IS NOT

- × A money printer or a guaranteed pass
- × A signal group or copy-trading service
- × A crystal ball. AI cannot predict the future
- × A replacement for having a real strategy
- × Something you switch on and stop watching

REGULATOR VIEW The US CFTC put it plainly in its 2024 customer advisory: *"AI won't turn trading bots into money machines."* Anyone promising guaranteed returns from an AI bot is a red flag. We agree.

03 · THE TEAM

Seven agents. *One* decision.

Six specialists do the work. One coordinator runs the sequence. The risk agent sits in the middle of the line, and nothing gets past it on a "no".

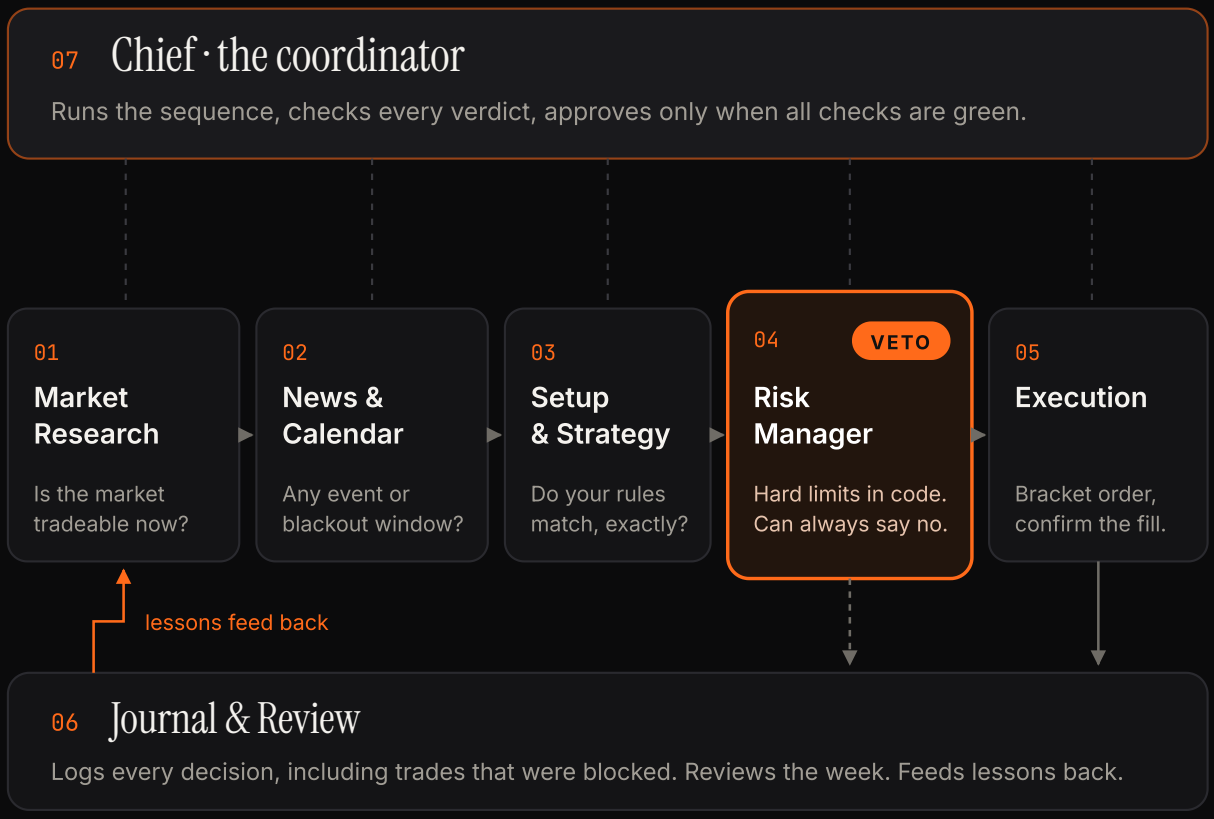

ANY "NO" STOPS THE TRADE · THE REASON IS LOGGED

Agents 01 to 03 are AI. They read, summarise and argue. They produce a proposal, never an order.

Agent 04 is deterministic code. Same inputs, same answer, every time. The AI cannot talk it round.

Agent 05 only acts on an approved, risk-checked ticket, with the stop attached.

Agents 06 and 07 keep the team honest: one remembers, one runs the order of play.

> Agents propose. Code decides what is allowed. *You* supervise.

The *specialists*

Research shows narrow, concrete jobs work better than vague personas. So each agent has one question to answer and a fixed format to answer in.

01 Market Research

AI · CAN BLOCK

Answers one question: **is this market in a condition my strategy is built for?**

LOOKS AT	Trend and structure on your timeframes, key levels, volatility versus normal, the session, related markets.
HANDS OVER	A short structured note: regime, bias, levels, confidence, and anything unusual.
BLOCKS WHEN	Conditions are outside what the strategy was tested on, such as extreme volatility or a dead, range-bound session.

02 News & Calendar

AI + DATA · CAN BLOCK

Answers: **is anything scheduled or breaking that changes the risk right now?**

LOOKS AT	The economic calendar, scheduled high-impact releases, central-bank events, headlines on the instrument.
HANDS OVER	Clear or blocked, the next event and its time, and a one-line headline summary.
BLOCKS WHEN	Inside your blackout window around a major release, or your prop firm's news rule applies.

03 Setup & Strategy

AI · CAN BLOCK

Answers: **does this match my written rules, exactly, and what is the case against it?**

LOOKS AT	Your rulebook, item by item, plus the notes from agents 01 and 02.
HANDS OVER	A trade proposal: direction, entry, stop, target, which rules matched, and the strongest argument *against* taking it.
BLOCKS WHEN	Any rule is only "nearly" met, or the counter-case is stronger than the case for. This is the bull-versus-bear debate from the research, built into one step.

The *gatekeepers*

04 Risk Manager

CODE · FINAL VETO

The agent that can say no, and the only one that is not an AI. **Plain code, hard limits, no exceptions.**

CHECKS	Position size from the stop distance, max risk per trade, daily loss used, distance to drawdown floor, trades taken today, max contracts, session hours, consistency.
HANDS OVER	Approved with the exact size, or rejected with the rule that failed.
WHY CODE	Language models can get arithmetic wrong and can be persuaded. A limit has to give the same answer every single time.

05 Execution

CODE

Turns an approved ticket into an order, and makes sure it landed correctly.

DOES	Sends a bracket order with stop and target on the broker side, confirms the fill, checks the stop exists, handles rejects and partial fills.
NEVER	Sends an order without a stop, changes size, or acts on anything not approved by agent 04.

06 Journal & Review

AI + DATABASE

The team's memory. Every decision is written down, including the trades it did *not* take.

RECORDS	Time, inputs, each agent's verdict and reason, the order, the result, screenshots.
REVIEWS	Weekly summary of rule breaks, blocked trades, and patterns. Suggests changes; a human approves them.

07 Chief

CLAUDE · COORDINATOR

Claude as CEO: it runs the order of play, reads every verdict, and **only moves forward when every check is green**.

DOES	Starts each cycle, routes work, stops the line on any "no", sends you alerts, and writes the daily brief.
CANNOT	Override the Risk Manager, raise a limit, or skip a step. Its authority ends where the code begins.

04 · ONE TRADE, START TO FINISH

Seven checks before *any* order

Here is one setup moving through the team. Every step must pass. If one fails, there is no trade, and the reason goes in the journal.

ILLUSTRATIVE WALK-THROUGH · NQ FUTURES · NOT A REAL TRADE

1 **System check** — PASS

Chief confirms the session window is open, the broker is connected, the data feed is fresh, and the kill switch is off.

2 **Market check** — PASS

Research: trending session, volatility normal, price pulling back into a level your rules care about.

3 **News check** — PASS

Calendar: no high-impact release inside the blackout window. Next event is hours away.

4 **Rules check** — PASS · 5/5

Setup: every rule in the playbook matched. Entry, stop and target proposed from your rules, not from opinion.

5 **Counter-case check** — PASS

Setup argues the other side: "Price is near yesterday's high." Judged weaker than the case for, with the stop placed beyond it.

6 **Risk check** — APPROVED · size set

Code: risk per trade within limit, daily loss well inside your buffer, drawdown floor far away, trade count under the cap. Size calculated.

7 **Order sanity check** — PASS

Chief and Execution: stop attached, size within max contracts, price still within tolerance of the planned entry.

✓ **Order sent**

Bracket order placed with stop and target held by the broker. Fill confirmed. Journal entry written. Alert sent to your phone.

If step 6 had said no (say, daily loss already near your limit) nothing else matters. No order. The Chief logs "blocked by risk: daily buffer", tells you, and waits for tomorrow.

05 · THE RISK LAYER

Rules are code. *Not* AI judgement.

An AI is good at reading a chart and arguing a case. It is the wrong tool for deciding whether you are allowed to lose another dollar today.

Numbers

Studies on financial documents show language models make more mistakes as calculations get more complex. Position size is a calculation.

Consistency

Ask a model the same question twice and the wording, sometimes the answer, can change. A limit cannot.

Persuasion

Models respond to the text they read. A dramatic headline should never be able to argue a limit away.

Prop-firm rules, as code

Rule types found across firms such as FTMO, Topstep and Lucid. Exact numbers differ by firm, account and date. Always check your firm's current terms.

RULE TYPE	TYPICAL PUBLISHED EXAMPLE	WHAT THE CODE DOES
Daily loss limit	FTMO 2-Step: 5% of initial balance, floating losses included	Locks new trades at your own, tighter buffer, well before the firm's line
Max / trailing drawdown	Topstep $50K: $2,000 limit, trailing on end-of-day balance	Tracks the floor live; cuts size as you approach it
Consistency	Best day under 50% of total profit (evaluation)	Caps the day's profit target so one day cannot dominate
Size limits	Max contracts per account size	Rejects any order above your max, every time
Session times	Flat by a set time before the close	Blocks entries late in session; flattens on schedule
News / conduct	Some accounts restrict trading around major releases	Blackout window enforced by agent 02 and the risk check

Sources: firm-published rules and help-centre articles, checked October 2026. Firms change rules often; firm names belong to their owners and CoTrader is not affiliated with them.

A kill switch, and *demo first*

What a risk check looks like

Simplified and illustrative. Real systems add more checks, but the principle is the same: plain rules, no model involved.

```
# risk_manager.py  (illustrative)
def check(ticket, acct, limits):
  if acct.kill_switch:
    return reject("kill switch on")
  if acct.daily_pnl <= -limits.daily_buffer:
    return reject("daily buffer used")
  if acct.room_to_floor < limits.min_room:
    return reject("too close to DD floor")
  if acct.trades_today >= limits.max_trades:
    return reject("trade cap reached")
  if ticket.stop is None:
    return reject("no stop")
  size = limits.risk_per_trade // ticket.risk_per_contract
  size = min(size, limits.max_contracts)
  if size < 1:
    return reject("stop too wide")
  return approve(size)
```

The kill switch

Professional guidance (FIA) describes a manual kill button that disables trading *and* cancels resting orders. Treat it as the last line, not the first.

- **Stops** new orders immediately
- **Cancels** anything still working
- **Flattens** open positions if you choose
- **Alerts** you on your phone, with the reason

Fires automatically when

- The daily buffer is reached
- Connection or data feed is lost
- Orders are rejected repeatedly
- Anything behaves unexpectedly
- You press it. One command, any time

> In 2012, one firm's faulty code sent millions of erroneous orders in about 45 minutes. Its controls did not stop it. That is what a missing risk layer costs.

Knight Capital, August 2012. Company-reported pre-tax loss of about $440 million; SEC order, October 2013.

Demo first. Always.

Run the full team on a simulation or practice account until its behaviour matches what you expect, every day, for weeks. Then go live at the smallest size your account allows. Topstep, for example, advises testing API strategies on a practice account first.

06 · EVIDENCE

What the research *says*

The honest summary: agent teams are a sound way to organise decisions. They are not proof of an edge. The best papers say so themselves.

2024

TradingAgents · Xiao et al., UCLA / MIT / Tauric Research

Models a trading firm: four analysts, a bull and bear researcher debate, a trader, a three-voice risk team and a final approver. Agents pass structured reports, not long chats.

Limit: tested over about three months on three stocks. The authors note their own Sharpe ratio was above what they would expect.

2025

FINSABER · Li, Kim, Cucuringu, Ma

Re-tested LLM strategies over two decades and 100+ stocks. Advantages reported in shorter studies largely faded. Strategies were too cautious in bull markets and too aggressive in bear markets.

Takeaway: the authors call for regime-aware risk controls over more complexity.

2025

StockBench and LiveTradeBench

Benchmarks on data the models could not have memorised. Most LLM agents struggled to beat simple buy-and-hold, and the best general chatbot was not the best trader.

Takeaway: a smarter model is not a strategy.

2026

Toward Expert Investment Teams · Miyazaki et al., Oxford

Splitting analysis into fine-grained, concrete tasks gave better risk-adjusted results than agents given vague job titles, in a leakage-controlled backtest.

Takeaway: one narrow, checkable job per agent.

2012

Knight Capital, SEC order, and FIA guidance

Automation without independent controls can fail fast. Industry guidance calls for pre-trade limits, connection monitoring and a kill switch that cancels resting orders.

Takeaway: risk lives outside the strategy, in code you can test.

Full references are available on request: founder@cotrader.pro. Paper results are backtests or benchmarks reported by their authors, not expected results.

07 · NEXT STEP

Build your own, or *have it built*

Build your own

A checklist, in the order that matters.

- [ ] Write your strategy as rules a stranger could follow
- [ ] Write down your prop firm's limits, then set your own, tighter
- [ ] Code the risk layer **first**. Test it with fake losses
- [ ] Add a kill switch: one command that cancels and flattens
- [ ] Give each agent one job and a fixed output format
- [ ] Log every decision, including trades you did not take
- [ ] Run on sim or practice until behaviour is boring
- [ ] Go live at minimum size. Review every week

Typical stack

Claude API for the agents · Python for the risk layer · a broker connection such as MT5, NinjaTrader, Tradovate or ProjectX · a Telegram bot for alerts · a simple database for the journal.

Have it built

CoTrader builds private systems around your rules and your firm's limits: the agent team, a coded risk layer, kill switch, journal and phone alerts. Tested on sim before it touches a live account.

1. **Apply**
   Tell us your strategy, platform and account.
2. **Review**
   We reply in writing on whether it can be built well.
3. **Scope**
   Clear spec, rules and limits agreed before any build.
4. **Build & sim**
   Built, tested, run on simulation with you.
5. **Handover**
   You own and run it. Kill switch in your hands.

Private builds, by application only

`cotrader.pro/apply`

Not ready yet? Start with the free library

Guides, templates and breakdowns on building with AI agents.

`princeofagents.app`

DISCLAIMER

This guide is for education only. It is not financial, investment, tax or legal advice, and it is not a recommendation to buy or sell any instrument. Trading futures, forex and CFDs involves substantial risk of loss and is not suitable for everyone. You can lose more than you expect. Any figures from research papers are backtests or benchmarks reported by their authors; hypothetical and past results do not indicate future results. Nothing here guarantees profits or that you will pass or keep any evaluation or funded account. Prop-firm rules described are examples from published sources and may have changed. FTMO, Topstep, Lucid, MetaTrader, NinjaTrader, Tradovate, ProjectX, Telegram and Claude are trademarks of their respective owners; CoTrader is not affiliated with or endorsed by them. CoTrader builds software to your specification and does not provide personalised investment advice or manage client funds.

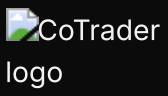

The edge is the *process.*

Read. Check. Argue. Limit. Execute. Record. Every trade, the same way. That is what a team of agents gives you. Not prediction. Discipline you cannot skip.

Prince

Founder, CoTrader

Have a private system built

`cotrader.pro/apply`

Free library

`princeofagents.app`