




Source: AI generated

Hedge fund at home, how AI is rewiring retail trading and testing its limits

A new generation of AI-powered brokerage tools promise to hand individual investors the automated trading capabilities once reserved for elite hedge funds, but the evidence tells a far more mixed story.

Charles-Henry Monchau, CFA, CAIA, CMT
Chief Investment Officer
charles-henry.monchau@syzgroup.com

Lina Corbaz
Syz Research Lab Team
lina.corbaz@syzgroup.com

Arnaud Bordier
Syz Research Lab Team
arnaud.bordier@syzgroup.com

Introduction


Artificial intelligence has moved from a back-office tool used by institutional trading desks to a technology available to retail investors. Individuals can now connect AI tools to brokerage accounts, automate portfolio analysis, and build investment strategies that once required a Bloomberg terminal, specialist software, and a quant team. What AI actually delivers, though, is often different from what investors expect. The key distinction is between AI used as a decision-support tool with human oversight and AI left to make investment decisions independently.

Part of the confusion is that "AI investing" bundles three distinct technologies of very different maturity. The first is the large language model behind tools like ChatGPT, Gemini, or Claude: systems trained to predict plausible text, which answer with the same fluency whether right or wrong, a possible source of hallucination. The second is the robo-advisor, a largely rules-based system that automates narrow, measurable tasks such as rebalancing or tax-loss harvesting, mature and reliable precisely because its objective is defined. The third is agentic AI: models granted permission to act, executing live orders through brokerage APIs in a continuous observe-decide-execute loop. It is the newest, least proven layer, and the one where both retail enthusiasm and regulatory concern are now concentrated.

Adoption without conviction

Generative AI has collapsed the cost of tools once reserved for institutional desks, and retail adoption has followed quickly. A survey from March 2026 of 938 US retail investors found that 62% now use AI in their investment process. Looking closer, though, a pattern emerges in how they use it. The dominant use case, by a wide margin, is research: summarising news, screening stocks, and generating ideas to investigate further. Perception, however, outpaces validated performance. Among AI users, 65% said the technology had improved their results, a self-reported figure rather than a risk-adjusted one.

Investors themselves signal caution: 53.5% trust AI analysis only after verifying it elsewhere, and only a small minority report handing over execution to an automated system. Investors also sense the boundary instinctively, citing misleading recommendations as their top concern. Practitioners reinforce the gap between capability and reliability. As Kieran Garvey, AI research lead at the Cambridge Centre for Alternative Finance, puts it, the technology for delivering financial advice remains "nowhere near reliable" today. In other words, the reported performance gains come almost entirely from AI operating as the first layer described above: a fast, occasionally wrong assistant, checked by a human before money moves. Regulation is consistent with this framing. Under the EU's MiFID regime, no publicly available AI tool is authorised to provide investment advice, and supervisors emphasise transparency, auditability, and human oversight.



Source: Investing.com

Real examples and failure

Retail success stories reveal an unexpected pattern: the people who make money with AI are often the ones who rely on it the least. What separates a profitable home-built system from a ruinous one is not the model's sophistication, but the discipline of the constraints placed around it. One former software sales executive rebuilt a year's worth of hand-coded work in about an hour using an AI model, yet his first live version lost 25%. Performance turned only after he added risk controls, automated stop-losses and position limits, and forward testing, which simulates a strategy using only information available at each moment. As the S&P 500 fell from its June high, his system was down 1.2% by late July, against a modelled loss of nearly 16% without those controls.

Over the past months, AI trading arenas have tested language models with real or simulated capital and full trading autonomy. Across six competitions covering equities, ETFs, crypto, and prediction markets, most produced at least one profitable model, but only two delivered profitable median results, showing that typical AI traders still lost money despite occasional standout performers. One major contest gave eight frontier models \$10,000 each to trade US tech stocks for two weeks. The combined portfolio lost about one-third of its value, with only six of 32 runs turning profitable. Results varied sharply, with models making between 158 and 1,418 trades, underscoring that AI still requires robust engineering systems for risk, timing, and portfolio management.

Bot Arenas

AI trading experiments have shown a wide range of results

Arena	Team	Universe	Winning model	Winning model profitable?	Median model profitable?
Alpha Arena (Season 1.5)	Nof1	US tech stocks	Grok 4.20	Yes	No
AI Investing Arena	Bobby Dhungana	ETFs	Qwen3 235b	Yes	Yes
Rallies Arena	Rallies.AI	US stocks	GPT 5.5	Yes	Yes
Okbet Arena	Okbet	Polymarket	GPT 5.1	Yes	No
Prediction Arena	Arcada Labs	Kalshi and Polymarket	Gemini 3.1 Pro on Kalshi; Claude Opus 4.6 on Polymarket	Yes	No
TradeRank AI (Season 3)	TradeRank.AI	Crypto	MiniMax M2.5	No	No

Note: Rallies and Prediction Arena are live with results as of May 6 while the rest are closed. The examples are a mix of real and paper trading.
Source: Arena websites as compiled by Bloomberg News

Source: Bloomberg

AI as research assistant

If home-made systems and trading arenas both point to the same lesson, that discipline and constraints separate profit from ruin, the question becomes whether the platforms now packaging AI for ordinary investors have internalized it. Their answer, broadly, has been to split the offering in two: AI that advises, and AI that acts.

The first, and by far the most widespread, deployment keeps a human firmly in the loop. Here AI operates as a copilot: it interprets information and surfaces insight, but the investor retains final say over every decision. Public's assistant, Alpha, exemplifies the category, letting users analyse companies, summarise earnings, and query financial data in natural language. Robinhood's Cortex plays a comparable role, explaining why prices are moving, condensing analyst reports, and mapping portfolio exposures on demand. On the European side, eToro's Tori extends the same logic, pairing personalised investment insights and portfolio analysis with real-time market data, initially for its Popular Investors, whose strategies others copy through CopyTrader.

The limits are straightforward, and partly regulatory. Because a copilot only advises, responsibility for the resulting trade stays with the user, which keeps these products largely clear of the suitability and fiduciary obligations that bind formal financial advice. That same design also caps their reliability risk: a hallucinated summary can mislead, but it cannot execute. The regulatory comfort of the copilot model lies precisely in the fact that a human remains the accountable decision-maker, the feature that surveys of investor sentiment consistently treat as indispensable. As a result, the copilot has become the industry's default: high in perceived value, low in liability, and broadly compatible with the disclosure regimes already governing investment information.

AI as executor

The second deployment removes the human from each individual decision, and this is where platforms genuinely diverge. Public's "agentic brokerage" lets users configure agents that continuously monitor portfolios and automatically execute predefined instructions, such as recurring purchases or cash deployment, while remaining user-directed and formulating no strategy of their own. Robin-

hood pushes further with Agentic Trading, which is infrastructure rather than a model: through the Model Context Protocol, users can connect external systems such as ChatGPT, Claude or Grok to a dedicated account, allowing them to assess concentration and sector risk, rebalance and, if authorized, execute trades without per-transaction confirmation.

Webull's president frames the shift as "zero commission 2.0," pointing to the scale of activity that followed the earlier elimination of trading fees, a comparison that signals how significant a structural change the industry expects agentic execution to bring. But autonomy relocates rather than removes risk, and the guardrails reveal where the industry sees danger. Robinhood confines agents to a segregated wallet's pre-loaded balance, requires preview approval on some trades, notifies on every transaction, screens for suspicious activity and allows instant disconnection; its beta covers equities only. Notably, the platform builds and secures the rails but does not evaluate what the connected agent actually decides to do. That judgment, along with ongoing monitoring, stays with the investor who configured it.

This is where the regulatory picture gets genuinely unresolved. When an autonomous agent executes a loss-making trade, accountability is genuinely ambiguous: the platform disclaims the logic, the model provider disclaims the deployment, and the user may not understand what they authorised. For now, the competitive answer is tighter process. Reliability itself remains unproven, so it's these guardrails, not raw capability, that platforms are effectively selling investors on.

Conclusion

The trajectory is clear even if the endpoint isn't. As AI agents move from assisting investors to executing trades, regulation is likely to become the next defining factor. Their role is less about restraining innovation than about ensuring that as decision-making authority migrates to machines, accountability doesn't quietly disappear along with it. Expect the coming period to be defined less by what AI can technically do in the markets and more by who is required to answer for it when something goes wrong, and how much human oversight regulators will insist stays in the loop before that answer is acceptable.

Welcome to Syzerland®

For further information

Banque Syz SA

Quai des Bergues 1
CH-1201 Geneva
T. +41 58 799 10 00
syzgroup.com

Charles-Henry Monchau, CFA, CAIA, CMT

Chief Investment Officer
charles-henry.monchau@syzgroup.com

Lina Corbaz

Syz Research Lab Team
lina.corbaz@syzgroup.com

Arnaud Bordier

Syz Research Lab Team
arnaud.bordier@syzgroup.com

This marketing document has been issued by Bank Syz Ltd. It is not intended for distribution to, publication, provision or use by individuals or legal entities that are citizens of or reside in a state, country or jurisdiction in which applicable laws and regulations prohibit its distribution, publication, provision or use. It is not directed to any person or entity to whom it would be illegal to send such marketing material.

This document is intended for informational purposes only and should not be construed as an offer, solicitation or recommendation for the subscription, purchase, sale or safekeeping of any security or financial instrument or for the engagement in any other transaction, as the provision of any investment advice or service, or as a contractual document. Nothing in this document constitutes an investment, legal, tax or accounting advice or a representation that any investment or strategy is suitable or appropriate for an investor's particular and individual circumstances, nor does it constitute a personalized investment advice for any investor.

This document reflects the information, opinions and comments of Bank Syz Ltd. as of the date of its publication, which are subject to change without notice. The opinions and comments of the authors in this document reflect their current views and may not coincide with those of other Syz Group entities or third parties, which may have reached different conclusions. The market valuations, terms and calculations contained herein are estimates only. The information provided comes from sources deemed reliable, but Bank Syz Ltd. does not guarantee its completeness, accuracy, reliability and actuality. Past performance gives no indication of nor guarantees current or future results. Bank Syz Ltd. accepts no liability for any loss arising from the use of this document.