


The AI Trading Verdict

Companion Guide

What Artificial Intelligence Can—and Cannot—Do for Traders



Artificial intelligence can do remarkable financial work. It can read filings, summarize earnings calls, compare companies, organize research, retrieve current documents, help write code, and turn more information than any one person could hold into something one person can actually use.

That does not make it a market oracle.

The central conclusion of *The AI Trading Verdict* is simple:

AI has made markets easier to read — not easy to predict.

This guide distills the book’s practical lessons: what AI does well, where the evidence thins, how to evaluate a trading claim, and where AI belongs inside a responsible trading process.

1. “AI Trading” Is Not One Thing

The label AI trading describes very different systems. A firm may build a predictive model specifically to rank securities. Another may adapt an existing language model to a narrow financial task. A retail trader may simply prompt an off-the-shelf chatbot about a stock. Researchers may use a frozen language model to turn news into numbers and train a separate predictor on them. An agentic system may connect a language model to data, code, memory, risk rules, and even a brokerage account.

These systems should not inherit one another's evidence. If a purpose-built machine-learning model predicts returns in commodity futures, that does not prove a general chatbot can tell you what stock to buy tomorrow. If a language model reads an earnings report accurately, that does not prove it can forecast the stock price.

So the first question is never "Does AI work?" It is: what AI, doing what job?

2. AI Is Already a Powerful Financial Reader

One of the clearest findings in the book is also the most useful: language models can be remarkably useful with financial language. They can identify what changed between two filings, locate management guidance, separate reported facts from expectations, ask the same questions of many companies, and organize large volumes of text into structured evidence. Retrieval makes this stronger, because the model can work from current, dated documents instead of relying only on what it learned in training.

That capability materially expands what an individual investor can accomplish.

But reading is not prediction. A model may correctly recognize that management lowered guidance without knowing whether the stock will fall. It may see that margins improved without knowing whether investors expected even more. Understanding the company and predicting the stock are different jobs — and that distinction may be the most important one in the book.

3. Prediction Exists — but the Evidence Narrows

AI and machine learning do show genuine predictive value in some settings. Language-based systems have identified short-lived information in news. Purpose-built models have produced strong results in defined markets. Hybrid systems have used language models to read text while separate statistical models estimate returns.

Those findings deserve credit. But their success depends on very specific conditions: the market studied, the information supplied, the target, the horizon, the model, the portfolio rules, the speed of execution, and the testing procedure. As the claim broadens, the evidence weakens.

There is a large difference between "this model extracted economically meaningful information from this news event" and "this AI can predict the market." The first may be well supported. The second requires far more evidence than the current research supports.

4. Accuracy Is Not Profit

A high win rate is the easiest trading statistic to misunderstand.

Imagine a strategy that wins 73 trades out of 100. Each win makes \$1, so the wins produce \$73. The other 27 trades lose \$4 each — \$108 in total. The strategy wins 73 percent of the time and still loses \$35 before costs.

Performance depends on more than how often the prediction was right. It also depends on the size of the wins and losses, position sizing, leverage, drawdown, spreads, slippage, commissions, and liquidity. That is why “90% accurate” and “95% win rate” are never enough. The missing question is always: accurate at what — and what happened to the money?

5. A Backtest Is the Beginning of the Audit

A beautiful historical result can be real and still fail as a trading strategy. The model may have used information that was not available at the time. The model itself may have encountered later information during training. A calculation may be wrong. The opportunity may vanish before a trader can act. Costs may consume the edge. The result may depend on one favorable market period.

And one problem deserves special attention: how many things were tried before the winner appeared? If one strategy is tested and succeeds, that means something. If thousands of prompts, models, settings, horizons, and portfolio rules are tested and only the winner is shown, the same chart means something very different. Searching hard enough can find impressive patterns in pure noise.

AI makes experimentation faster and cheaper. That is valuable — and it makes the search itself part of the evidence.

6. Autonomous Does Not Mean Dependable

Modern AI agents can retrieve documents, analyze information, use tools, remember earlier actions, propose trades, manage a portfolio’s state, and sometimes submit orders. That is a real engineering achievement. It is not automatically an investment achievement.

The research reviewed in the book shows both sides. Some agents performed encouragingly in short portfolio tests. Broader, longer evaluations were less convincing. Agents made arithmetic errors and formatting errors, overtraded in some designs, and behaved very differently across market regimes.

The lesson is not that autonomous trading is impossible. It is that the authority given to a system should not grow faster than the evidence supporting it. The more authority the system has, the stronger its controls must be. An AI allowed to suggest an idea is one level of risk. An AI allowed to resize positions, reinterpret limits, and place trades is another.

7. Retail AI Creates Access — and Illusion

For retail investors, AI's greatest benefit may not be prediction at all. It is access. AI cuts the time, skill, and cost of serious analytical work. One investor can now examine more filings, compare more companies, explore unfamiliar industries, write analytical code, and ask sharper questions.

But the same technology makes weak analysis look finished. A polished interface, a personalized explanation, a confidence score, recent winners, testimonials, and an "AI-powered" label can make uncertainty feel settled before it is. Interacting with a fluent system can even make the user feel more capable — while the underlying evidence has not improved at all.

Retail AI creates three effects at once. Access is real. Influence is real. Illusion is real. The goal is to capture the first without being fooled by the third.

8. Put Each Tool in the Right Seat

The book's strongest practical conclusion is not that traders should avoid AI. It is that no single tool should be asked to do everything.

Language models read. Use them to retrieve, summarize, compare, explain, organize, question, and help develop research.

Code calculates. Use deterministic tools when arithmetic, reproducibility, and precision matter.

Predictive models estimate. Use tested models for defined probabilities, rankings, and forecasts.

Portfolio systems constrain. Use explicit rules for sizing, exposure, diversification, drawdown, and risk.

Humans verify. Keep responsibility for deciding what evidence is sufficient and when the process should stop.

The model is not the system. The complete process is what gets judged.

9. The Four Questions to Carry With You

The book asks four questions of every AI trading claim.

Filter 1 — The Job. What exact job is the AI performing, and what was actually measured?

Filter 2 — The Information. What information did it receive, when did it receive it, and could it already have known the answer?

Filter 3 — The Reality. Was the result still available after time, costs, competition, risk, and market change?

Filter 4 — The Search. How many prompts, models, settings, features, horizons, and portfolios were tried before the reported winner was selected?

A printable Four Filters card is available as a separate download.

10. The AI Trading Verdict

AI has earned an important place in financial research. It helps investors read more, understand more, compare more, build more, and test more than was practical for an individual only a few years ago. It can also accelerate weak research, bad backtests, unsupported conclusions, and poorly controlled trading systems.

That is why capability alone is not enough.

AI can make a good process more capable and a bad process more dangerous.

Use AI aggressively where it improves the work. Keep calculations reproducible. Require prediction to earn its place through testing. Make risk rules explicit. Verify important claims outside the conversation. And keep responsibility where it belongs.

Let AI expand your reach, not replace your judgment.

Define the job. Audit the information. Follow the edge across time and cost. Count the hidden trials. Judge the entire system. Then make the trade yourself.

Companion Downloads

The Four Filters. A one-page reference for evaluating any AI trading claim.

AI Trading Claim Checklist. A printable checklist to use before paying for an AI trading product — or trusting one with money.



ProbabilityTraderPro.com/ai-trading-verdict

The material in this guide is for educational and informational purposes only. It is not investment advice, trading advice, or a recommendation concerning any security or financial product. Historical results, published research findings, statistical relationships, and AI outputs do not guarantee future performance.