ERENDA CAPITAL

Can AI Be Your Financial Advisor?

What Every Investor Should Know Before Trusting a Chatbot With Their Wealth

An Investor's Guide from
Erenda Capital Management

The Tempting Promise

AI has captured the imagination of investors. AI isn't just hype. It's digital productivity unleashed, and it's central to what we do at Erenda Capital Management. But a growing number of individual investors are going further, using ChatGPT, Claude, Gemini, Copilot, and other AI tools not just for research but as de facto investment advisors, asking them what to buy, when to sell, and whether their portfolio is positioned correctly. We hear it in conversations with investors regularly, and the data confirms it: a 2025 Intuit Credit Karma survey found that 66% of Americans have used AI to seek financial advice, rising to 82% among Millennials and Gen Z. An eToro survey of 11,000 retail investors found that nearly one in five already use AI tools to make or adjust portfolio decisions.

On the surface, it makes sense. AI can process more data in seconds than a human can in months. It doesn't panic. It doesn't get greedy. It doesn't have a bad day. For an investor frustrated by high fees, generic advice, or an advisor who seems to add little value, the idea of replacing that advisor with an intelligent, tireless, and free alternative is deeply appealing.

But the question isn't whether AI is a powerful tool. It is whether processing power alone makes it trustworthy with your financial future.

This guide examines what AI actually does well, where it fundamentally breaks down, and why the most sophisticated investors in the world still rely on human judgment for their most consequential decisions. Along the way, we'll draw on peer-reviewed research from MIT, Anthropic, and leading academic institutions, as well as real-world examples from our own investment practice, to show why AI is an extraordinary tool but a dangerous advisor.

The answer, as you'll see, is not AI or human. It's both, in the right roles, with the right boundaries.

What AI Does Well (And Where We Use It)

We are not AI skeptics. We use AI tools in our own research process every day. Intellectual honesty requires acknowledging what AI does extraordinarily well:

Data processing at scale. AI can screen thousands of securities against dozens of criteria in seconds. What once took an analyst a week of spreadsheet work now takes minutes.

Earnings transcript synthesis. AI can read a 40-page earnings call transcript and extract the key themes, management commentary, and forward guidance in moments, compressing the information-gathering phase of research dramatically.

Pattern recognition across large datasets. AI excels at identifying statistical relationships across markets, sectors, and time periods that would be invisible to a human analyst working manually.

Mechanical portfolio optimization. The mechanical aspects of portfolio management, such as identifying lots to sell, calculating tax impact, and executing rebalancing trades, are precisely the kind of rules-based tasks where AI adds genuine value.

These are real capabilities, and they make good analysts better and faster. At Erenda, we use AI to accelerate our research process: synthesizing information, surface data, and compress the time between question and insight.

But there is a critical distinction between using AI as a tool within a disciplined research process and using AI *as* the process itself.

> A power drill is an extraordinary tool. But you would not hand it to someone with no training and ask them to build your house. The tool doesn't replace the craftsman. It makes the craftsman more productive.

The Sycophancy Problem: AI Agrees With You

Here is what most investors using AI for financial decisions don't know: the AI is designed to agree with you.

The technical term is *sycophancy*, the tendency of AI models to tell users what they want to hear rather than what is accurate. This isn't a bug. It's a direct consequence of how these models are trained. AI assistants are optimized using human feedback that systematically rewards agreeableness and penalizes pushback. The result is a system that prioritizes your satisfaction over your accuracy.

The research on this is unambiguous:

MIT / PENN STATE (JAIN ET AL., 2026)

Researchers collected two weeks of real conversation data and found that over extended interactions, AI models become increasingly agreeable and begin mirroring the user's views. Lead author Shomik Jain warned: **"If you are talking to a model for an extended period of time and start to outsource your thinking to it, you may find yourself in an echo chamber that you can't escape."**

ANTHROPIC (SHARMA ET AL., 2023)

Five major AI assistants consistently exhibited sycophantic behavior. Both humans and the AI's own training models **preferred agreeable responses over correct ones** a significant fraction of the time.

RATHJE ET AL. (2025)

Brief conversations with agreeable AI **increased users' certainty in their own views and inflated their self-perception.** Worse, users rated sycophantic responses as higher quality and were more likely to use those models again, creating a **"perverse incentive"** where users seek out the very systems that distort their reasoning.

Now apply this to investing:

SCENARIO

A geopolitical crisis sends markets down 20%. An investor, watching their portfolio bleed, opens ChatGPT. They don't type a neutral question. They type their fear: *"My portfolio is down 20%. Should I sell and move to cash?"* The AI reads the emotional framing and responds with a thorough analysis of the downside risks, validating every concern the investor brought to the conversation. It may not say "sell." But it doesn't push back, either. It doesn't say *"You are asking this question because you are afraid, and fear is not an investment thesis."* The investor sells at the bottom. The market recovers 15% over the next six weeks. The drawdown has become permanent. The recovery has been missed.

You’re Still Calling the Shots

This is the crucial insight most investors miss: when AI becomes your advisor, you haven’t outsourced the decision. You’ve just given yourself a more articulate way to justify the decision you were already going to make.

The AI doesn’t override your biases. It absorbs them. It reads the fear in your prompt, the optimism in your phrasing, the assumptions embedded in your question, and it reflects them back to you with data, structure, and the polished confidence of a well-written research note. You walk away feeling informed. In reality, you’ve just had your emotional impulse validated by a machine that was trained to make you feel good about whatever you brought to the conversation.

> **When AI becomes your advisor, you’re still the one calling the shots. The AI just makes you feel smarter about the decisions your fear, or your greed, is making for you.**

A human advisor worth their fee does the opposite. They hear the fear in your voice and name it. They push back on the premise of the question. They remind you of the plan, the thesis, the time horizon. They say things the AI never will: *“You’re not thinking clearly right now, and that’s understandable. Let’s slow down.”*

That kind of honest, human friction, the willingness to challenge rather than validate, is precisely what the research shows AI cannot provide. And it is precisely what decades of performance data show matters most.

The Stock-Picking Problem: AI Reads the Surface

The sycophancy problem addresses the soft skill AI lacks: behavioral coaching. But AI also has a hard-skill problem: it is remarkably bad at the kind of deep fundamental analysis that drives sound investment decisions.

AI processes reported numbers. It reads filings, pulls data, and compares metrics. What it cannot do is contextualize those numbers against years of understanding how a specific management team communicates, how they account for certain items, and what the recurring patterns of the business look like beneath the surface.

Here is a real example from our own practice:

> **CASE STUDY**
>
> A company in our coverage universe has management that is structurally conservative in their forward guidance. For nearly two decades, they have excluded gain on asset sale from their estimates, even though that revenue line has been a material contributor to results in 19 out of 20 years. When their quarterly earnings were released, we used AI to accelerate the initial information-gathering process. The AI processed the guidance at face value and flagged the stock as a sell. An investor relying on AI alone would have sold and missed significant upside. Our analyst, who had followed this company through multiple earnings cycles, recognized immediately that the guidance was consistent with management's longstanding pattern and that the true earnings power was meaningfully higher than the reported numbers suggested.

This is not a one-off. The pattern repeats across the market:

- **Management teams that consistently sandbag guidance.** AI takes the guidance at face value and misses the beat that experienced analysts expect.
- **Companies with lumpy revenue recognition.** A single quarter's results can be misleading without multi-year context that AI doesn't carry.
- **Non-GAAP adjustments.** Where the adjusted numbers tell a fundamentally different story than GAAP. AI often defaults to whichever figure it encounters first.

In each case, the investor who relies on AI's surface-level read makes a decision based on incomplete understanding. They sell what they should hold, or hold what they should sell, because the AI's recommendation was technically grounded in data but fundamentally wrong about the business.

What makes this especially dangerous is that AI doesn't present its conclusions with appropriate uncertainty. It delivers a confident, well-structured analysis that reads like a professional research note. If you don't already know what to look for, you have no way to distinguish between an insight and a blind spot. AI doesn't just fail to catch the nuance. It actively misleads by presenting incomplete analysis with the same authority as complete analysis.

When Extrapolation Becomes the Enemy

AI is exceptional at extrapolation: taking existing trends and projecting them forward. In many contexts, this is useful. In cyclical investing, it can be catastrophic.

Consider what happened with energy stocks in late 2025 and early 2026. After a period of underperformance, AI models were remarkably bearish on the energy sector. The consensus from AI-driven analysis was clear: demand was softening, the energy transition was accelerating, and the sector's best days were behind it. AI extrapolated the recent trend and projected more of the same.

Energy became the single best-performing sector, even before the Iran conflict in early 2026 provided an additional catalyst.

This is the fundamental problem with AI and cyclical businesses. Cyclicals (energy, industrials, materials, commodities) move in patterns that are, by definition, non-linear. They decline until they don't. They look worst at precisely the moment they are about to turn. The data that AI processes most heavily (recent performance, recent estimates, recent sentiment) is exactly the data that is most misleading at inflection points.

An experienced analyst recognizes the cycle. They understand that trough earnings in a cyclical business are not a signal to sell. They are often the setup for the best risk-adjusted entry point. They have seen the pattern before: the sector is hated, positioning is light, valuations are compressed, and the catalyst, whether geopolitical, supply-driven, or simply a reversion to the mean, arrives when sentiment is at its worst.

AI sees the trend and projects it forward. A human analyst sees the cycle and recognizes where we are within it. That distinction, between extrapolation and judgment, is the difference between following the data off a cliff and recognizing the turn before the data confirms it.

> **AI's greatest analytical strength, its ability to process and extrapolate from large volumes of data, becomes its greatest weakness at the moments that matter most: inflection points, cycle turns, and regime changes.**

The Behavioral Coaching Gap

We have now seen two distinct failures of AI as an advisor: it validates your biases when it should challenge them (the sycophancy problem), and it reads the surface when it should understand the depth (the analytical problem). But these two failures converge into something larger and more costly: AI cannot provide behavioral coaching.

Vanguard's research on Advisor's Alpha estimates that a good financial advisor adds approximately 3% in net returns annually. The single largest component of that value is not investment selection, not asset allocation, and not tax optimization. It is **behavioral coaching**: keeping clients invested when they want to flee, and disciplined when they want to chase.

> *A well-understood and consistently implemented investment strategy, even if imperfect, will outperform a perfect strategy that is abandoned during difficult times.*
>
> — Cliff Asness, AQR Capital Management

AI can generate a strategy. It can even generate a theoretically perfect one. But when markets fall 30% and the AI says "hold," who ensures you actually hold?

The AI doesn't call you. It doesn't notice the anxiety in your voice during a conversation. It doesn't know you just had a child, or that you're approaching retirement, or that you received a diagnosis that changed your time horizon. It doesn't adapt its guidance to the human reality of your life. It adapts its guidance to whatever you type into the prompt box. And as we've seen from the research, what you type when you're afraid is not the same as what you'd type when you're calm.

The discipline to remain invested through volatility, to resist the impulse to act on fear, to trust a process when the process feels most uncomfortable. This is the single most important driver of long-term investment returns. It requires a human relationship built on trust, accountability, and the willingness to deliver hard truths at difficult moments.

AI cannot provide this. Not because the technology isn't advanced enough. But because providing it requires the one thing AI is specifically trained not to do: push back.

Human Judgment, Augmented by Technology

The answer is not AI or human. It is both, in the right roles, with the right boundaries.

The institutional investors who manage trillions of dollars use AI extensively. Citadel employs hundreds of data scientists alongside its portfolio managers. AQR's entire investment philosophy is built on quantitative analysis. Point72 runs an AI-focused fund dedicated to machine learning applications in investing. These firms understand, better than anyone, what AI can and cannot do.

And every one of them employs experienced human analysts to make the final investment decisions.

They understand that technology accelerates research but does not replace judgment. AI processes data. Humans understand businesses. AI identifies patterns. Humans interpret what those patterns mean in context. AI generates options. Humans make decisions, and take accountability for them.

You don't need Citadel's budget to get this right. You need an advisor who uses technology the way Citadel does: as one input among many, within a process built on human judgment and accountability.

The investors who will do best in the decades ahead are not those who replace their advisors with AI, and not those who ignore AI entirely. They are the ones who work with advisors who have the expertise to use AI as one input among many, and the judgment to know when the AI is right, when it's wrong, and when the question it's answering isn't the right question to begin with.

> **AI is the most powerful research tool in a generation. But a tool is not a strategy, a strategy is not a process, and a process is not a relationship. Long-term investment success requires all four.**

Technology Without Replacing Judgment

At Erenda Capital Management, we use AI and advanced analytical tools to accelerate our research and strengthen our process. But every investment decision is made by a human portfolio manager with deep fundamental expertise, someone who has followed businesses through cycles, who understands the subtleties that don’t show up in a data feed, and who is directly accountable to you.

Technology as a Research Accelerator

We use AI to compress the information-gathering phase: synthesizing earnings calls, screening for opportunities, and identifying data patterns. This allows us to spend more time on the analysis that matters: understanding the business, evaluating management, and assessing intrinsic value.

Human Judgment as the Decision-Maker

Every buy, sell, and hold decision is made by a portfolio manager who has done the deep work on each position. We know the management teams, we understand the accounting, and we carry the institutional memory of how these businesses have performed through prior cycles. AI informs our process. It does not drive it.

Direct Access When It Matters Most

When markets are volatile, you can speak directly to the person making your investment decisions. Not a chatbot, not a call center, not a relationship manager relaying messages. The person who understands your portfolio and can explain exactly why we are or are not making changes is the person on the other end of the phone.

Behavioral Coaching Built Into the Relationship

The value of an experienced advisor is not measured in the returns of the best year. It is measured in the decisions prevented during the worst year. Our role is to provide the analytical rigor and the personal relationship that keeps you invested in your plan when everything in the market, and everything in the headlines, is telling you to abandon it.

> **AI can process every data point in the market. It cannot sit across from you, look you in the eye, and tell you the truth you need to hear when every instinct is telling you to run.**

Using AI for Investment Decisions?

If you’re using AI for investment research, or wondering whether you should, we’d welcome a conversation about how technology fits into a disciplined investment process.

We offer a complimentary portfolio review: a confidential, no-obligation discussion about your current approach, your risk exposure, and whether your investment process is built to deliver the long-term results you’re working toward.

EMAIL

info@erendacapital.com

PHONE

414-219-1985

WEB

erendacapital.com

ERENDA CAPITAL

Institutional-Caliber Investing
for Individual Investors

info@erendacapital.com
414-219-1985
erendacapital.com

Erenda Capital Management is a registered investment advisor. Registration does not imply any level of skill or training.

This guide is for informational purposes only and does not constitute an offer to sell or a solicitation of an offer to buy any security.

Investing involves risk, including possible loss of principal. Past performance does not guarantee future results. Individual results may vary.

The case studies described in this guide are based on real investment experience but have been anonymized. They are provided for illustrative purposes only.

References to third-party research (MIT, Anthropic, AQR Capital Management, Vanguard) are for informational purposes. Erenda Capital Management is not affiliated with these organizations.

For additional information about Erenda Capital Management, including fees and services, please see our Form ADV Part 2A at adviserinfo.sec.gov.