

CALCHAS RESEARCH

White Paper

The Commoditisation of Truth

May 2026

The Commoditisation of Truth

Let's paint a picture:

In 2022, a research analyst is going through a semiconductor company's annual report. He's been doing this for the past 10 years, he's good at what he does. In about 4 hours he has deconstructed the report, surfaced the relevant financials, and written up an initial draft report.

Later that night, he gets sent a link from a fellow analyst to ChatGPT. It's the hottest new technology available, and it's revolutionary. Curious, he tinkers with ChatGPT, querying it on the same company he was researching. Both versions read side by side, it looked, for the most part, equally rigorous. He spotted some nuances in the filings that the machine did not catch, some broader macro news connections that the machine was not trained to draw, but the AI-generated report was almost identical to the human report.

Four years later, the AI-generated report is leagues ahead of the human-led research, no longer missing the nuances that only the analyst picked up. With the right prompt engineering and the right context window, AI tools surpass human analysis, many times over.

Some investors still believe the game is still about knowing more, faster data, more accurate modes, and sharper analysis. Unfortunately, the edge has shifted elsewhere. The quant revolution created an informational edge for a few, before information was made near universal. With AI tools, information acquisition has never been faster, broader, or deeper.

This diffusion of edge started with Bloomberg commoditising data. Previously, statistical arbitrage used algorithms to exploit mathematical pricing patterns at machine speed, commoditising pattern recognition. Research aggregators further commoditised analysis and interpretation for institutions, and then for the masses. Now, LLMs are finishing off the commoditisation stack. What took analysts hours in 2022 takes an LLM minutes today.

Gone are the days of information scarcity. We live in a world of interpretation parity.¹ When everyone has access to virtually the same data, analysis, and interpretation, what happens next?

Every technology humanity has ever produced eventually becomes universal. The printing press began as an advantage for whoever controlled it, and it's now the baseline infrastructure of literate society. Calculators took arithmetic from a professional skill to a free function on every phone. The personal computer, the internet, the algorithm, each begin as an asymmetric advantage. But as time goes on, the asymmetric advantage the technology once created is so dominant that it eventually disappears.²

Bringing things back to finance:

The first hedge funds to use quantitative strategies made extraordinary returns. Naturally, everyone built quantitative strategies. The niche became the baseline. The first firms to use Bloomberg terminals saw data nobody else did. Then every serious firm got its terminals, and the informational edge became the entry fee. The lifecycle of any competitive advantage runs like this: discovery, replication, and diffusion. The premium migrates up the stack to wherever genuine scarcity lives next.³

Artificial Intelligence presents a unique challenge to finance that other technology has not: it provides the retail investor access to what institutions have been building their edge on for years.

We're witnessing the collapse of informational asymmetry.



Data

Bloomberg launched its terminal in 1981, and SEC EDGAR opened free public access to corporate filings in 1993. Through the '90s, the internet slashed the price of data, economic releases, and regulatory disclosures to negligible amounts. In the early 2000s, alternative data providers proliferated, providing satellite imagery, credit card transactions, and web traffic scraping. By 2005, the data available to any well-capitalised fund was practically identical throughout the Street. Informational asymmetry was getting compressed, and an informational edge was harder and harder to find. Granted, dark networks like corporate access channels, private expert networks, and personal relationships built through professional tenures still exist, but these operate as closely guarded secrets, not a structural layer.


This only touches the institutional market. Data was largely inaccessible – or when they were accessible, uninterpretable – to the layman. That's changing now. Bloomberg's terminal still costs \$24,000 per year per user, practically unfeasible and economically unsound for individual investors. But functional equivalents are being built around the institutional-grade tool, giving retail investors access to information they once weren't privy to.

TradingView (2011) gave retail an interface that rivals those used at hedge funds. Koyfin (2016) delivers Bloomberg-style fundamental data for a few hundred dollars a year. Robinhood, Webull, and other successors bundle real-time quotes and options flow data into zero-commission apps that reach tens of millions. By the early 2020s, options flow, large off-exchange trade records, and institutional positioning were available on retail platforms for under \$50 a month. Bloomberg's moat did not disappear, but the data that drove trading decisions diffused downward.

Analysis

What happens when everyone has access to the same analysis?

In August 2007, the quantitative hedge fund industry suffered a crisis now known as the Quant Quake. A large fund, forced to raise cash to cover losses elsewhere, began liquidating its equity portfolio. What followed was a 12-standard-deviation event, a flash crash, and a lesson in what distributed analysis can mean for the convergence of positions. Over three days, the dominoes toppled; every other fund holding the same factor exposures had their risk limits triggered, forcing more selling, and further triggering other funds' risk limits.



It was pandemonium.

There are many hypotheses on the flash crash, many citing the immediate trigger as a liquidity-driven spiral, but the more structural issue lies at the heart of the diffusion of analytical edge. This resulted in the convergence of positioning, which was what permitted

the domino effect in the first place.^{3a} The underlying research driving many models had been diffusing for a decade. Anyone who could code and had access to a research paper could

build a systematic strategy. By 2010, Python, R, and open-source back testing infrastructure had made systematic analysis available to any graduate student with a laptop.

Then the retail barrier vanished with AI. What one once needed domain expertise to parse and understand, now needs an account, a prompt, and some patience. Projecting cash flow, screening by quantitative factors, decomposing balance sheets, building a comparison table – the advancement of AI did not just slash the finance practitioner's workflow, it granted the retail crowd analytical prowess they did not train. Granted, the query quality remains a variable; the AI output is highly dependent on how well we structure our inputs. But that gap is compressing too. LinkedIn handbooks, YouTube tutorials, and open-source templates now teach anyone how to spin up their own, personal agentic equity research workflow with their LLM subscription.

Interpretation

Research from investment banks was historically a premium paid service, but it was then bundled into trading relationships. Research aggregators, Bloomberg, AlphaSense, Factset, and CapIQ, to name a few, standardise thesis across reports, providing a scannable recommendation without having to manually read the research.

Like in the analysis stack, AI is dragging the interpretation edge in the same direction: to zero. Accumulated domain expertise, like knowing which line items mattered, which management guidance to discount, or which macro narrative to apply, is now accessible in your search engine's free AI mode. AI is not perfect at this, but the trajectory matters more than where we are now. We've seen maybe 5 years of AI iterations. What will AI output look like in 10 or 15 years? Even the institutional tier is facing a similar commoditisation cycle. Bloomberg is now rolling out ASKB, an AI interface that lets terminal subscribers query Bloomberg's news, filings, earnings transcripts, and research. Navigating the terminal traditionally required a steep learning curve, which is now flattened through turning Bloomberg's terminal into a thought partner, less like a tool. Bloomberg still costs \$24,000 a

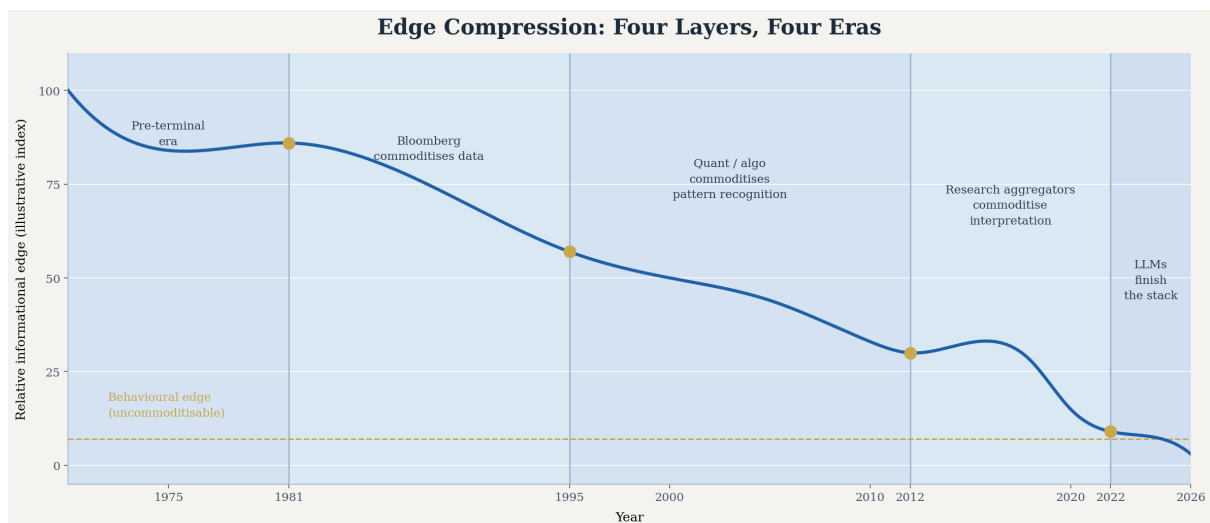
year, but it won't be long before other platforms build AI interfaces that are trained on their entire database.

Behaviour

The last node in the chain is completely different from the ones before. Every layer prior feeds into this terminal step, the final gap between information and action. Behaviour, however, is the only layer where the observer and the observed are in the same system, making it the only layer that can't be commoditised.

Think about what happens when the markets sell off sharply. Investors know that panic-selling is one of the most reliably destructive financial decisions one can make. Studies have shown this, financial advisors have cautioned against it, but none of it matters. When markets go red and keep going red, information gets shelved, and emotion takes control. Better information is not going to change this; better technology is not going to solve this. Behavioural biases, loss aversion, herding, and priming, these patterns run below the level of conscious deliberation. Knowing about them does not prevent them from operating or influencing your decisions.

The last non-replicable, uncommodifiable asset of finance is the ability to model what the crowd will do when it cannot model itself.



Decoding the Behavioral Layer

Behavioural alpha, as known in the literature, consists of both defensive and offensive edges. Defensive – protecting against your own biases – preserves alpha. Systems can be built to replicate behavioural alpha as a skill: risk protocols, pre-commitment rules, and process guardrails all work to protect the investor from their emotional selves. Offensive behavioural alpha – exploiting documented biases of others – is directly systematisable. To name a few, "buying the dip" is codified panic-selling exploitation and momentum strategies are codified herding. Once the pattern is known and traded, like all the layers before this one, the premium will compress by the same mechanism.

The third tier is what we call Talent Alpha. It is a capacity for second-order thought, modelling the crowd's behaviour, premises, and beliefs, and positioning for when the crowd is treating premises as fact. Every investor who correctly predicted financial crashes exhibits at least one of these three skills.

1. Reading the narrative before the numbers

Most analysts construct a narrative from the numbers, but Talent Alpha requires practitioners to ask what story the crowd is currently telling themselves, assessing whether that story is a product of marketing, if it's quietly broken, and only then checks whether the numbers substantiate the story.

2. Thinking a level above the consensus

It is never just "what will happen?", it is "given what everyone believes will happen, what actually happens, and what does the difference between those two things produce?" It requires asking questions the crowd has stopped asking, and predicting when and how the crowd will converge on that view. Most investors implicitly understand this, but never explicitly model this second-order chain. This means mapping what the crowd believes,

knowing what's keeping it intact, and finding the specific condition under which it would crack.

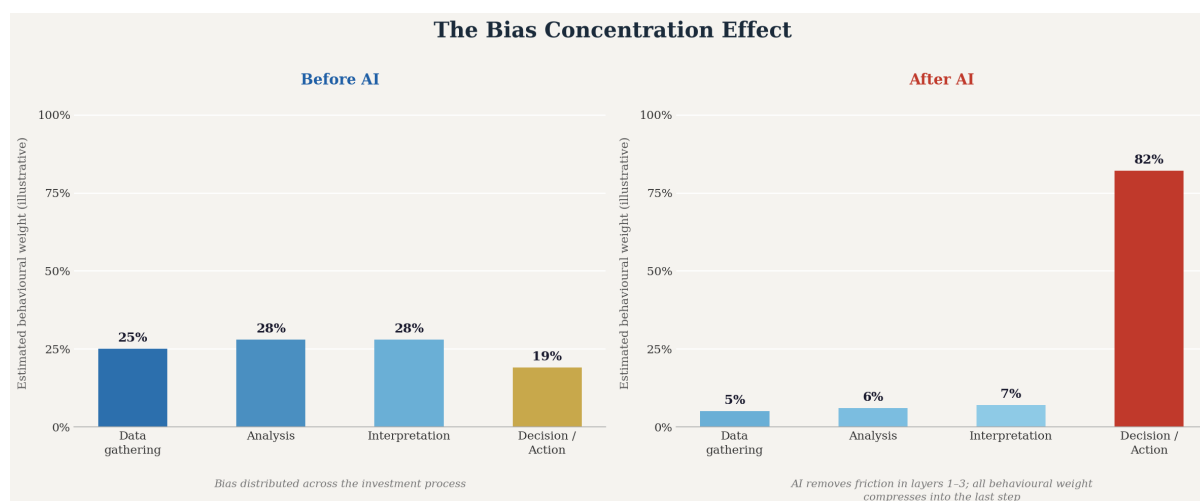
3. Knowing what you're actually responding to

The last habit is slightly different (we talked about it in our first white paper); Before acting on any signal, are you reacting because it's the most important signal, or because it's the loudest?^{3b} Whether a signal gets priced in depends first on whether it gets noticed, and what gets noticed is a function of salience. The practitioner who asks what is capturing their attention, and whether salience and relevance are pointing in the same direction, is working at a level that no systematic process approaches.

In some ways, artificial intelligence can ask the same questions, stress-test the same assumptions, and pick apart the same beliefs as humans can. Machines are very good at processing signals, but narrative-reading involves not processing what is being said, but identifying the assumptions so embedded in the consensus that they have ceased being a relevant data point. Invisible premises don't appear in data; they are the frame in which it was produced.⁴

Concentrating Bias

Before terminals, before aggregators, and before AI, behavioural mechanisms were diluted across the investment process. Confirmation bias shaped which data you looked for, anchoring bias manifests when you read analyst reports, and herding operates through calls and consensus, loss aversion surfaced at the moment of execution.

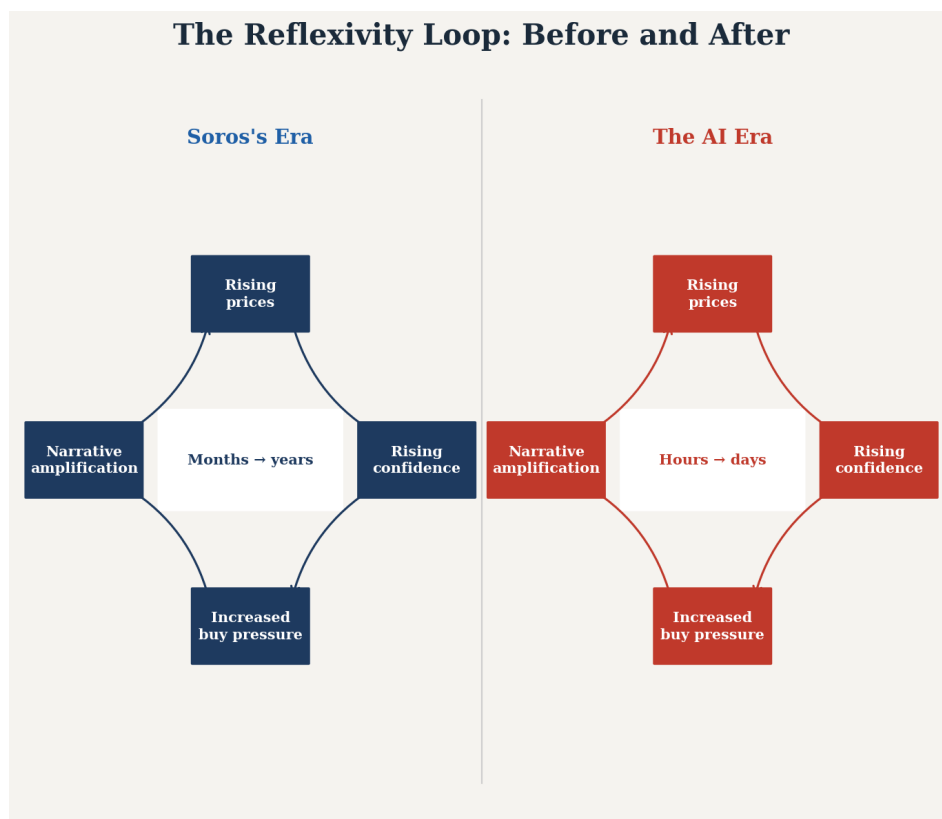


The commoditisation of the first three layers has second-order effects. It's easy to mistakenly believe that the advancement of technology removes the emotionality of investing. It does not. It compresses all the emotionality that once existed in the prior three layers into the last layer. The layer right before decision-making, one moment where the entire behavioural weight of the process now concentrates.

A 2025 study of 400 active traders using AI platforms found that AI usage simultaneously reduces loss aversion ($\beta = -0.41$) while amplifying overconfidence ($\beta = +0.56$) and anchoring ($\beta = +0.48$).⁵ If you think reducing loss aversion sounds like progress, it is not. Reduced fear combined with amplified overconfidence is not the emotional profile you would want before making financial decisions.⁶ An investor with the aforementioned psychological patterns anchors to the AI's price target, holds longer than the position warrants, and when loss aversion finally overrides the anchor, exits at the worst possible moment. Scale that across a whole market, and what looks like a market move is actually a synchronised

behavioural event.⁷ The mechanism is less of narrative speed, more of structural, psychological uniformity of human decision-making.

50 years ago, George Soros illuminated that market prices don't passively reflect reality; they influence it. His concept of reflexivity outlines a self-reinforcing loop of rising prices creating confidence, confidence creating buying, buying creating rising prices. This loop perpetuates until the gap between price and value becomes unsustainable, and the loop inverts. In Soros's era, these loops propagated through slow channels, broker calls, magazine covers, and the boom/bust cycle could run across years. Today, that same loop runs in time frames of days. The underlying mechanism remains the same, but the cadence of it has shifted by an order of magnitude. The gradual standardisation of input and output compresses the emotional sequence; it doesn't remove it.⁸



A popular belief is that AI processing information faster and more accurately will rationalise the market, but the concentration effect is the corrective. Further advancements in artificial intelligence will not make the crowd smarter; it makes us faster at being wrong in the same

ways we always have. Further, we become more uniform in the errors we make, because the inputs are now identical.⁹

Now, more than ever, investing has become a behavioural question, not an analytical one. When interpretation converges with the same summary, logic, and conclusion, the differentiating question when looking at any data point is not "What does this mean?", but "Given that everyone else is seeing this, what will they do next?" Technological advancements like Bloomberg's terminal have gradually shifted the landscape in this direction, but AI is exponentially reducing that gap. Now, everyone has access to the first-order question. That precisely makes the second-order question more valuable.

But if behavioural patterns become more consistent and legible, won't capital crowd into behavioural trades until the edge disappears? The prior three layers' commoditisation cycles are distinct from the fourth. Each one compressed an advantage external to the market participant. The edge was compressed through better, faster data, better models, and better interpretations. Behavioural patterns are immune to this commoditisation because it is endogenous. They are generated by an architecture that cannot be updated by awareness of itself. The January Effect decayed because knowing about it was enough to trade it away; loss aversion does not disappear when an investor learns they are loss-averse.

Why Behaviour Can't Be Modelled

Daniel Kahneman spent fifty years mapping the precise ways human judgment fails under uncertainty. In numerous interviews, he has been clear about the fact that, despite his knowledge of the flaws of human judgment, he is not granted immunity. He still experiences loss aversion. He still makes the errors he documented. Awareness and behaviour exist in different parts of the brain, and awareness does not govern behaviour.^{10, 11}

We assert that behaviour is an inherently different layer from the prior three that cannot be commoditised. A company's revenue figures do not change because an analyst models them. An analytical framework does not become invalid because it was published. Information, once accurate, stays accurate no matter how many people access it.

Behavioural patterns work differently, and it's why we believe Talent Alpha will be the remaining edge for years to come.

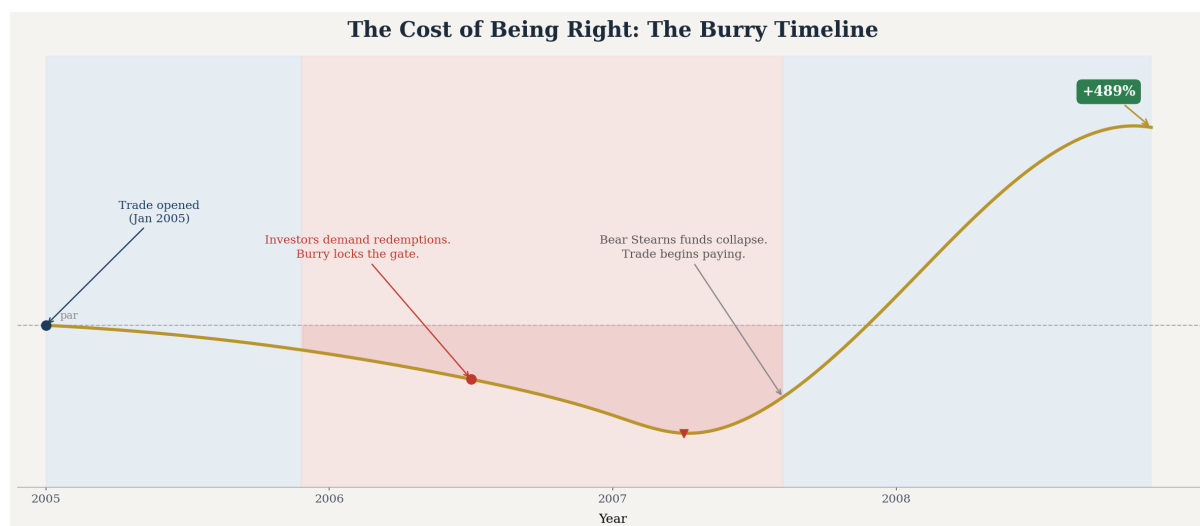
The moment a behavioural pattern or an artefact of human psychology becomes widely known, the pattern itself gets folded into the crowd's reasoning, shifting the behaviour to compensate for the anomaly. The January Effect is a mild version of this. Once enough investors learned that small-cap stocks historically outperformed in January, enough capital moved in in December to pre-empt the premium.

But these patterns don't just get arbitrated away; they mutate under observation. The observer effect occurs every moment in the stock market. Think about an investor who learns that loss aversion causes them to exit too early. They may incorporate that information and overreact, holding losses past any rational exit point. The incorporation of new information corrects the correction, and the model is always one step behind. The more precisely a behavioural model is specified, the quicker it fails at the moment of consensus.

Behavioural dynamics can be studied, partially anticipated, and exploited in the space between the pattern emerging and the crowd reaching a consensus about it. But it cannot be captured by a model. The accuracy of the model itself necessitates its inefficacy.

Burphy's Story Retold

In 2005, Michael Burry built a position against the housing market. As the market kept rising, his investors wanted out. Angry letters, lawsuit threats, because from the outside, being early is the same as being wrong. Had he not suspended redemptions, he would have been forced to close the position at exactly the moment it was about to pay.¹² 3 years later, the trade returned 489%.



We see an interesting relationship between managers, their decisions, and their investors. To succeed, it's not enough for you to be right, but you have to have the environment that lets you be right for long enough while others think you're wrong, so that your thesis will play out.

Keynes understood this in 1936. The winning strategy in markets is not to identify what's true; it's to identify what the crowd will believe next, and then hold that position long enough for the crowd to catch up to you. It's never been about finding out what's true.

The hard part is being willing to be early.¹³ Early means the position looks wrong before it looks right. Looking wrong means the people backing you get nervous.¹⁴ Nervous people pull out. And when they pull out, you close the position at exactly the moment it's about to pay. Which, ironically, the cascade comes full circle, and closing the position too early makes you wrong.

The understanding of one's own behaviour and biases, the modelling of others' reactions, and the management of others' expectations have always been a part of finance. The insertion of AI and other tools into the mix doesn't reduce the role psychology plays; it just shifts it. Beyond the quant era, beyond the AI era, investors will find that the behavioural era has been there all along. As crowd psychology shifts further and further downstream into the step right before action, behavioural finance and the modelling of second-order psychological reaction will start becoming more important, not less.

Endnotes

¹ The edge was never data access — it was processing. Bernard & Thomas (1989) showed that markets failed to incorporate the predictable serial correlation in quarterly earnings for decades, despite the information being public and the pattern statistically knowable. 25–40% of the drift concentrated at the next quarter's announcement.

² Minsky's Financial Instability Hypothesis (1992): stability is destabilizing. During prolonged prosperity, financial structures migrate from hedge (cash flow covers obligations) to speculative (covers interest but not principal) to Ponzi (requires asset appreciation). Each commoditization layer that becomes universal increases systemic confidence and systemic fragility at the same time.

³ McLean & Pontiff (2016) tested 97 published return anomalies and found that post-publication returns decay 58% on average. Decay was smallest in stocks with high idiosyncratic risk — where arbitrage is most costly. Commoditization is a measurable mechanism: once a finding is published, capital crowds in until the premium compresses.

^{3a} Khandani & Lo (2007) provide the forensic reconstruction. Using a simulated contrarian strategy on CRSP data, they document the Aug 7–9 losses as 12 daily standard deviations (!), statistically impossible under fundamental repricing, followed by an Aug 10 reversal of 11.4 standard deviations. The mechanism was probably that one large equity market-neutral portfolio was rapidly unwound (likely forced by credit-side margin calls), generating price impact across all funds holding the same factor exposures via shared BARRA/Northfield risk models. The "quant quake" was proof that the methods had replicated so completely that differentiation had ceased to exist. Alpha decay documented in the same paper: avg daily return of the contrarian strategy fell from 1.38% (1995) to 0.13% (2007), requiring ~9:1 leverage to maintain 1998-level returns. The leverage is what converted a liquidity event into catastrophic drawdowns.

^{3b} Hirshleifer & Teoh (2003): informationally equivalent disclosures have different price effects based on salience. Headline earnings receive attention; footnotes, segment breakdowns, and accrual decompositions are systematically underweighted. The practitioner reading the non-salient disclosures is exploiting the same attentional constraint that the crowd cannot escape.

⁴ Hong & Stein (1999): private information diffuses gradually through a population, even when the information is knowable. "Newswatchers" correct slowly; momentum traders overshoot past the correction. The housing prospectuses were public documents, but the consensus narrative meant nobody really took the time to process them. Narrative-first thinking identifies information that is public but not yet diffused.

⁵ Saidalavi, Rasheed K, and Afeefa Cholasserri. "Influence of AI Trading Tools on Cognitive Biases and Investor Behavior: An Exploratory Study." *IUP Journal of Accounting Research & Audit Practices*, Vol. 24, Issue 3, July–September 2025.

⁶ Kahneman & Tversky's prospect theory (1979) explains why this combination is dangerous. The value function is convex below the reference point: investors become risk-seeking in losses. Reduced loss aversion (from AI framing) combined with anchoring to an AI price target means the investor holds deeper into the convex loss domain, chasing recovery, until the pain overwhelms the anchor and forces a capitulation exit.

⁷ De Long, Shleifer, Summers & Waldmann (1990) formalized noise trader risk: the unpredictability of irrational sentiment is itself a priced risk that constrains arbitrageur position sizing. When AI synchronizes sentiment across participants, noise trader risk does not disappear — it becomes correlated, which makes it harder, not easier, for rational capital to stabilize the system.

⁸ Loewenstein et al. (2001): emotional reactions to risk diverge from cognitive assessments, and when they conflict, behavior follows emotion. Vivid narratives and proximity intensify

emotional response independent of probability. AI-generated narratives arrive simultaneously and vividly across all participants — compressing the reflexivity loop not just through speed but through synchronized emotional intensity.

⁹ Lerner & Keltner (2001): fear and anger produce opposite risk assessments despite identical negative valence. Fear (high uncertainty, low control) drives risk aversion; anger (clear attribution, agency) drives risk-seeking. AI platforms standardize not just the data but the narrative framing — which standardizes the type of negative emotion across participants, reducing the behavioural variance that would otherwise slow a herd.

¹⁰ LeDoux (2003): fear conditioning is encoded at the lateral amygdala through long-term potentiation (LTP). This is a synaptic-level mechanism that persists even after repeated non-reinforced exposure. The fear pathway operates in under 20 milliseconds, before cortical processing can intervene. This is the neuroscience of "awareness does not produce immunity."

¹¹ Phelps (2006): emotion regulation via cognitive reappraisal requires prefrontal cortex resources: the same resources that are compromised under acute stress. The crowd reads the behavioural models under exactly the conditions that degrade its capacity to act on them. Published models fail not just because the crowd incorporates them, but because the incorporation happens when regulatory capacity is at its lowest.

¹² Shleifer & Vishny (1997): performance-based arbitrage (PBA). Specialized arbitrageurs use outside capital. When mispricing deepens before correcting, the arbitrageur's short-term losses trigger withdrawals — forcing liquidation at the moment of maximum opportunity. The agency structure between fund manager and investor is the mechanism that converts being right into being forced out.

¹³ Brunnermeier & Nagel (2004), as synthesized in Statman (2019): hedge funds increased their technology holdings during the dot-com bubble. They were not fooled — they were

riding momentum rationally, because bubble duration is uncertain and short positions face capital risk during the holding period. Being early is not just psychologically hard. It is economically irrational in the absence of a capital structure designed to absorb temporary losses.

¹⁴ Scharfstein & Stein (1990): investment managers rationally discard private information and follow the crowd because the career consequences of being wrong alone are asymmetric. Shared bad outcomes are attributed to conditions; idiosyncratic bad outcomes are attributed to judgment. The manager who bets against the consensus and loses — even temporarily — bears the full reputational cost.