CALCHAS CAPITAL

White Paper

The Behavioral Finance Shaped Hole

2026

Abstract

The "behavioral" dimension of finance has been explored extensively in academia, but few practitioners operationalise the consistent findings of behavioral economics and investor psychology from the past 50 years. It is a genuinely under-utilised tool, and this paper bridges that gap of academia and profession.

Attention has two relevant properties here: it is unidirectional, and its target is determined by salience. Macro and micro signals are treated as independent and unrelated — when they should be held in the same frame, for the same company, at the same time. We give them different weights at different times depending on what is capturing our attention, producing systematic mispricings that traditional finance and behavioural theory have not operationalised into a practitioner framework.

Mispricing lives in the intersection between signals. Before acting, a practitioner should ask: what is capturing my attention, why is it capturing my attention, and is that because it is actually the most relevant signal — or simply the most salient one? That discipline of metacognition is what behavioural finance has never operationalised into daily practice. It is the shaped hole.

1. Markets, War, and What Gets Lost

The financial markets are often seen through indices. We judge their health based on the underlying businesses, the economics, and the geopolitics. But the markets are rarely seen as a collection of human decisions. There is a human aspect behind price action that gets lost in the newscycle. When we reduce finance to a number or a percentage point, we are missing a more human dimension, a dimension that explains many market anomalies, and will be the underpinnings of many more.

An expression of this — when markets receive news of war, they front-load fear. We see this in real time. They price it in, get used to it, and move on to the next thing that hits the news cycle. This is not a flaw in market architecture. It is a consequence of who is doing the driving. Humans price the tail risk when they are fearful, we systematically overshoot.

Another expression of this is that the people who run, govern, and invest in the financial markets have finite attention. When a macro event fills all available attention, the micro signals that swim beneath it do not get priced correctly. This crowding effect is not random. It is structural — a systematic crowding effect we will call the geopolitical override. It has always happened and it always will — but it has never been operationalised into a practitioner framework.

2. The Mechanism: Attentional Tunnelling

Attentional tunnelling is the cognitive tendency to fixate on the most salient stimulus in the environment, effectively crowding out competing signals regardless of their relevance. It is easy to imagine this phenomenon having far reaching consequences in driving and aviation, but it occurs daily for finance professionals too.

When dominant macro news fills the newscycle, our attention is captured by it — the war that is ongoing, the oil prices that are skyrocketing, geopolitical tensions rising (Shiller 2017). But the moment we read a ticker and its earnings, this frame collapses into a micro one. These two frames are treated as independent and unrelated when they have clear interaction effects. We are not accounting for them because our attention cannot hold both simultaneously.

Kahneman's What You See Is All There Is (WYSIATI) is a psychological bias that has been well documented. For an investor, it means that the way we think about a situation is built entirely from whichever signals that are currently salient. We fail to incorporate other signals not because they are absent, but because the mind does not register their absence. We construct a complete-feeling picture from incomplete inputs without flagging the gap because of this bias.

Consider a company that releases a strong earnings beat on the day a major geopolitical shock dominates the news cycle. The stock price underreacts to the underlying economics, resulting in post-earnings-announcement drift (PEAD) — where the stock drifts in a certain direction for weeks or even months instead of fully pricing in the information at the time of the earnings call. The macro event did not change the company's fundamentals. It changed where attention was directed.

Attentional tunnelling does not operate alone. Belief perseverance works simultaneously against the investor (Kahneman 2011). Tunnelling determines what arrives. Perseverance determines what gets prioritised and what gets suppressed. Once a frame is adopted — the macro narrative, the risk-off posture, the war premium — contradictory evidence that arrives later is not evaluated on its merits. It is discounted. The micro signal does not register as conflicting; it simply does not register at all.

Attention precedes decision. When the market is not paying attention to a signal, it does not get priced in. It is at the intersection of macro and micro signals that mispricing lives.

3. *The Evidence*

When we read macro news, we tend to think at a global scale — geopolitics, international relations, macro regimes. When we read a ticker, we tend to think at a company scale — economics, product offering, partnerships. These are two fundamentally different frames that we cannot hold in our minds at the same time. When the frame switches, the prior context collapses (Kahneman 2011; Hirshleifer and Teoh 2003). We see this frame-switching effect in PEAD. It is the most direct measurable trace of this disconnect, documented consistently since 1985 and replicated across decades and markets (De Bondt and Thaler 1985).

Post-earnings announcement drift is the finding that after a company reports earnings that beat or miss analyst expectations, the stock price continues drifting in the same direction for weeks or even months afterward, rather than adjusting fully on the announcement date. Under the Efficient Market Hypothesis, this should not happen — all information in an earnings announcement should be priced in immediately. This drift is one of the most persistent and replicated anomalies in financial markets, documented since 1968 and surviving decades of scrutiny. Trading strategies exploiting PEAD have generated substantial risk-adjusted returns even after controlling for transaction costs (Hirshleifer, Lim and Teoh 2009), raising the question: if the alpha is there and has been documented for over fifty years, why hasn't it been arbitraged away?

Investors underreact to earnings announcements because other, more salient signals are capturing their attention. The micro signal gets drowned out by whatever is more salient at that moment. The price correction happens gradually as attention returns to the stock and the market slowly incorporates what it should have priced in on day one.

Hirshleifer, Lim and Teoh (2009) showed that PEAD is stronger when competing stimuli are present on the same day. This is precisely the mechanism at work in the geopolitical override — concrete and salient signals capture attention and drag it away from competing information.

We term this the geopolitical override — the systematic suppression of micro signals during periods of dominant macro attention saturation. In times of high geopolitical tension, dominant macro events saturate available attention. Wars, central bank pivots, systemic banking crises — these are the regimes where the geopolitical override operates most visibly. In these periods, micro signals tend to get drowned out. Earnings calls, guidance changes, material operational announcements, all compete for attention against a macro backdrop that has already claimed all available attention. We observe this pattern consistently: during acute geopolitical events, sectors with no fundamental connection to the macro event still experience correlated sell-offs, as attentional capture overrides fundamental discrimination.

In regimes where there is no salient macro news, micro signals tend to be at the forefront. It is only in the vacuum of macro signals that micro signals shine. That is the anomaly that is observable.

The geopolitical override suppresses micro signals only until attention fatigue sets in. For example, when the threat of war is replaced by an ongoing war, incremental escalations do not produce the same incremental aversions in the market. Caldara and Iacoviello (2022) showed geopolitical attention spikes at conflict onset and decays as events become ongoing rather than acutely threatening. Because of this attentional shift and normalisation, the override phenomenon lasts for a finite period — even when the news cycle has not changed.

Attention floods toward the more salient event. Micro signals — earnings beats, guidance upgrades, operational improvements — released during a window of acute macro news receive abnormally weak price reactions because attention is elsewhere (Hirshleifer, Lim and Teoh 2009). We extend that finding here from micro-to-micro distraction to macro-to-micro distraction. The macro event gradually fades from the foreground as it normalises into the background. Attention redistributes as the subject fails to remain salient. The suppressed micro signals begin to be priced in weeks or months later, drifting in the same direction the market should have moved on day one.

The override is not permanent. That is the opportunity.

A systematic test of the geopolitical override would follow the methodology of Hirshleifer, Lim and Teoh (2009) — replacing competing earnings announcements with macro distraction intensity as the key variable. On days when a major macro event dominates — a war, a central bank decision, a systemic shock — do individual companies' earnings announcements receive weaker price reactions than on days of macro quiet? Distraction intensity can be measured through VIX elevation, news sentiment saturation, or the geopolitical risk index developed by Caldara and Iacoviello (2022).

One finding in the literature appears to complicate this picture, and it is worth addressing directly. In 2022, Hirshleifer and Sheng found a pattern that appears, at first glance, to contradict the geopolitical override (Hirshleifer and Sheng 2022).

On days with major scheduled macroeconomic announcements, Fed decisions, payroll data, CPI releases — PEAD is 71% weaker and announcement-day returns to earnings are 17% stronger. Their finding suggests that macro and micro news are complements, not substitutes: scheduled macro events sharpen, rather than suppress, the market's attention to firm-level earnings.

However, the distinction that resolves this is the nature of the macro event itself. Hirshleifer and Sheng studied *anticipated* announcements, events that investors were prepared for. Their finding is distinct from what the geopolitical override describes, it operates on a different class of event entirely: sudden, unscheduled, overwhelming shocks where attentional saturation is not chosen but imposed.

Scheduled events allow investors to allocate attention deliberately and process in a structured way. A war breaking out on a Thursday morning is not an event investors are prepared to attend. It floods the attentional foreground without warning, leaving no capacity for deliberate allocation to competing micro signals. The mechanisms differ in kind, not just in degree. The geopolitical override is the unscheduled version of the problem Hirshleifer and Sheng confirm exists at the scheduled level.

Price is not a function of information alone, most practitioners forget that. It is a function of information and attention. And attention is finite and fickle.

4. Why It Remains Unsolved

This argument faces three serious objections. Each is stated below in its strongest form and addressed on its merits.

We address three camps: the quant crowd arguing the diminishing of psychological biases in an increasingly systematic and passive market, EMH holdouts arguing the premia will be arbitraged away, and behavioral finance professionals arguing that frameworks already exist to capture this phenomenon.

Objection 1

The objection from the quant crowd is that behavioural anomalies like PEAD are already captured in systematic factor strategies (Fama and French, 2015). Has the alpha that we are describing been harvested already? Has systematic investing already solved the psychological dimension? They argue that systematic investing removes psychology from price discovery entirely. Systematic investing aside, there is a growing group of passive investing through target-date funds, robo-advisory platforms, automatic payroll-linked enrollment programs, and low-cost index replication vehicles that intermediates the retail investor's decision-making entirely out of the investment process (Morningstar, 2026; Vanguard, 2025; Madrian and Shea, 2001).

In its strongest steelman form, this is not just a market share argument. If the investors who have generated the behavioral inefficiencies that active managers exploit are increasingly delegating capital allocation to systematic strategies, where is the psychological dimension? This is a market whose participants do not overreact to earnings, do not exhibit disposition effects, and do not experience cognitive biases.

However, passive investing does not remove psychology — it concentrates it into a group of discretionary actors. Among this group of actors is a growing retail crowd whose biases are empirically more pronounced (Barber and Odean 2000). Retail investors now account for 30 to 37 percent of daily US equity trading volume (MEMX 2025), participation has grown structurally since the adoption of commission-free trading in 2019 (WEF/BCG 2024), and increasingly accessible platforms like retail brokerages, robo-advisors, and mobile trading apps continue to lower the barrier to entry. The crowd psychology is not shrinking. It is getting amplified.

Factor models are also backward-looking by construction. They identify what has worked historically but cannot identify when a known anomaly is being amplified by a specific macro regime right now. Research confirms that factor alphas decay significantly after publication and widespread adoption (McLean and Pontiff 2016). Dynamic weighting across the cycle is what factor models do not do.

Objection 2

The second objection comes from EMH holdouts, arguing that markets are remarkably good at processing information on average and over long time horizons (Fama 1970). The more sophisticated argument uses PEAD as evidence itself. The anomaly has been publicly known since 1968, and its alpha has shrunk significantly as arbitrageurs exploited it (McLean and Pontiff 2016). By that logic, the framework is either capturing something that will soon be arbitraged away, or is not capturing anything real.

However, arbitrage is constrained by the same psychology as the people being arbitraged against. Attentional tunnelling is the mechanism and PEAD is the proof. Even when prices are demonstrably wrong, professional arbitrage cannot always correct them. This is because arbitrage requires capital, carries career risk, and arbitrageurs face redemption pressure from investors who judge them on short-term performance, not long-term performance (Shleifer and Vishny 1997). An arbitrageur in a macro override regime who identifies a suppressed micro signal still cannot hold the position long enough to capture the correction if their investors are also distracted by the same macro event and withdraw capital at the worst moment.

The EMH position gets one thing right that we should concede: markets are remarkably good at processing information over long horizons (Fama 1970; Grossman and Stiglitz 1980). The general tendency toward efficiency is real. What EMH misses is the specific conditions under which efficiency breaks down. These conditions occur when macro and micro signals coexist with *unequal force* and attentional constraints prevent simultaneous processing. We are not arguing markets are always inefficient, but they are structurally inefficient at specific intersections and regime conditions.

Objection 3

The last camp to address is behavioural finance professionals arguing that sentiment has already been operationalised at the aggregate level. Shiller's CAPE ratio adjusts for psychological cycle effects across the market (Shiller 2000). Baker and Wurgler's sentiment index shows that sentiment predicts returns as a systematic factor (Baker and Wurgler 2006). They might also argue that the patterns described here are the result of curve-fitting to historical data, and that PEAD's alpha has shrunk since publication.

On CAPE and Baker-Wurgler: these are aggregate instruments that do not resolve the intersection at the company level in real time. CAPE tells us the market is expensive and sentiment is elevated, but it does not tell us that this specific company's earnings surprise is being suppressed by a specific macro regime at this moment. CAPE gives us the absolute level, not the company-specific direction. Neither CAPE nor Baker-Wurgler resolves both dimensions of company-specific and regime-specific simultaneously. This is the gap. The geopolitical override is precisely the mechanism that amplifies attentional failures during the windows when PEAD is strongest.

On data mining: PEAD's persistence across decades and markets makes data mining an insufficient explanation. PEAD has been replicated across different time periods, different markets, different methodologies, and different research teams for over fifty years. An anomaly that survives that level of scrutiny is not noise. When we call such a finding an "anomaly", we are refusing to ask the harder question: Are these “anomalies” part of the structure, and what are the ramifications of ignoring them? There is an underpinning we have not yet explored. This paper is an attempt to name it.

The geopolitical override will be systematically tested. Until it is, it is presented as a practitioner observation grounded in the same attentional mechanism the literature has already confirmed.

5. The Implication

The implication is precise: a mispriced asset under-reacts to its micro signal because a larger macro signal has captured the attention that should have been directed at it. This is not an abstract problem. It shows up even in the most basic tools practitioners use every day.

Consider a discounted cashflow model. In current practice, the assumptions embedded in a DCF are structurally arbitrary. They may be advised by past financial performance, earnings calls, or company guidance, but it is ultimately shaped by the analyst's own orientation toward the company. It is influenced by opinion, by attention. Analysts anchor to prior estimates, prior price targets, prior narratives, and they adjust insufficiently when new information arrives (Cen, Hilary and Wei 2013; Campbell and Sharpe 2009; Hong, Kubik and Solomon 2000).

A growth rate assumption doesn't come from nowhere — it anchors to the macro narrative, it is assumed. An analyst modelling a DCF during a geopolitical override regime is embedding that regime's risk-off posture into their discount rate without realising it. Two analysts covering the same stock in different macro regimes will produce structurally different models from the same underlying data. The spread in fair value estimates isn't random disagreement. It's the attentional fingerprint of whoever built the model and when.

Fair value estimates have carried wide spreads between analysts covering the same stock. Finance has noticed this divergence, but we have chosen to call it skill or opinion rather than bias. That choice has foreclosed the question before it has been asked.

Recognising that the bias is structural rather than individual is the first step. The second is building the habit of asking a question that standard analysis does not require.

For a practitioner who accepts this argument, a metacognitive discipline follows. When our attention is captured by a signal, the first question to ask is whether it is capturing attention because it has genuine structural implications, or simply because it is salient and recent. Different signals carry different weight at different points in the cycle. The practitioner who holds both the macro regime and the company-level signal in the same frame — who asks not just what the earnings number is but what macro context it is being read inside — is doing something that standard financial analysis does not require.

This sounds simple but is not. The architecture of information dissemination ensures this gap persists. Macro signals arrive through news, central bank statements, and geopolitical reporting — continuous, ambient, and loud. Micro signals arrive through earnings releases, footnotes, and guidance, each requiring active retrieval from sources that do not announce themselves. The market has no instrument that holds both channels open simultaneously and resolves the conflict between them in real time. That is not a gap in practice. It is a deficiency in the structure. A structural gap requires a different instrument — not just a different mindset.

The practical entry point is attention itself, monitored in real time. With the rise of real-time search analytics and alternative data, monitoring attention before it becomes price action has become operationally possible. Take search volume, social platform trends, Reddit discussion, YouTube spikes — these are all expressions of the same underlying phenomenon: attention concentrating before capital flows. A spike in Google Trends search volume for a company signals attention concentration, which precedes capital flow asymmetry. When attention is concentrated on a stock, the distribution of buyers and sellers becomes temporarily unbalanced, and that imbalance is what moves prices. Liquidity is the context, not the cause. In a highly liquid market, a large order moves prices less, but the direction is still determined by the flow imbalance that attention creates.

Monitoring a company name alone is insufficient. Companies do not exist in a vacuum. When a node in the supply chain shifts, that shift must be captured too. A practitioner monitoring Nvidia watches not just the ticker but its composite: AI, computer chips, GPUs, semiconductor supply. When attention moves through the ecosystem, prices follow. The composite search approach captures that movement before it becomes a price event.

But what does this look like in practice? Imagine this scenario: a company releases a strong earnings beat during a period of macro saturation — a war is going on, a Fed decision is about to be made, a banking shock is rattling through the market. Under normal conditions, investor attention concentrates around earnings announcements. Search volume rises in the weeks before, spikes markedly at announcement, and remains elevated for weeks after as the market processes the news (Drake, Roulstone and Thornock 2012; Da, Engelberg and Gao 2011). This is what attention looks like when it is working. During a geopolitical override regime, none of this happens. Search volume for the company barely moves on announcement day. The macro event is consuming all available attention. The stock underreacts. Over the following weeks, as the macro event normalises into the background, search volume for the company begins to rise. Attention redistributes. The price starts drifting in the direction the earnings beat warranted on day one. A practitioner monitoring the composite search approach would have seen the attentional vacuum at announcement, recognised the suppressed micro signal, and understood that the correction was coming. Not if, but when.

If this framework were integrated into standard practice, more anomalies would be explicable not as random noise but as predictable outputs of attentional failures at the intersection of competing signals that cannot be held in the same frame simultaneously. The shaped hole is not a gap in the data. It is a gap in the architecture of how practitioners are trained to think.

There is one further dimension to this problem that deserves naming. When signals do not appear simultaneously — when the macro event and the micro signal arrive at different times — the temporal gap itself creates a recency bias. The most recent signal gets weighted more heavily than an earlier signal that may still be the more relevant one. Attentional tunnelling and recency bias are not separate problems. They are two expressions of the same underlying constraint.

6. The Shaped Hole

The attentional tunnelling that constrains all of us is a problem of structure. It has been observable for decades, and it has been explained away as an anomaly.

Mispricing lives in the intersection of signals. Before acting, a practitioner should ask: what is capturing my attention, why is it capturing my attention, and is that because it is actually the most relevant signal, or simply the most salient one? That discipline of metacognition is what behavioural finance has never operationalised into daily practice. It is the shaped hole.

Naming the problem is the first step. The second is understanding that the instrument we need does not yet exist in complete form. But the shaped hole now has a name, the mechanism has been identified, and the research agenda is clear. What remains is the building.

Shaun Koh

Founder

Calchas Capital, 2026

Bibliography

Baker, M. and Wurgler, J. (2006). "Investor Sentiment and the Cross-Section of Stock Returns." Journal of Finance, 61(4), 1645–1680.

Barber, B.M. and Odean, T. (2000). "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors." Journal of Finance, 55(2), 773–806.

Caldara, D. and Iacoviello, M. (2022). "Measuring Geopolitical Risk." American Economic Review, 112(4), 1194–1225.

Campbell, S. and Sharpe, S. (2009). "Anchoring Bias in Consensus Forecasts and its Effect on Market Prices." Journal of Financial and Quantitative Analysis, 44(2), 369–390.

Cen, L., Hilary, G. and Wei, K.J. (2013). "The Role of Anchoring Bias in the Equity Market: Evidence from Analysts' Earnings Forecasts and Stock Returns." Journal of Financial and Quantitative Analysis, 48(1), 47–76.

Da, Z., Engelberg, J. and Gao, P. (2011). "In Search of Attention." Journal of Finance, 66(5), 1461–1499.

De Bondt, W.F.M. and Thaler, R. (1985). "Does the Stock Market Overreact?" Journal of Finance, 40(3), 793–805.

Drake, M., Roulstone, D. and Thornock, J. (2012). "Investor Information Demand: Evidence from Google Searches Around Earnings Announcements." Journal of Accounting Research, 50(4), 1001–1040.

Fama, E.F. (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work." Journal of Finance, 25(2), 383–417.

Fama, E.F. and French, K.R. (2015). "A Five-Factor Asset Pricing Model." Journal of Financial Economics, 116(1), 1–22.

Grossman, S.J. and Stiglitz, J.E. (1980). "On the Impossibility of Informationally Efficient Markets." American Economic Review, 70(3), 393–408.

Hirshleifer, D. and Sheng, J. (2022). "Macro News and Micro News: Complements or Substitutes?" Journal of Financial Economics, 145(3), 1006–1024.

Hirshleifer, D. and Teoh, S.H. (2003). "Limited Attention, Information Disclosure, and Financial Reporting." Journal of Accounting and Economics, 36(1–3), 337–386.

Hirshleifer, D., Lim, S.S. and Teoh, S.H. (2009). "Driven to Distraction: Extraneous Events and Underreaction to Earnings News." Journal of Finance, 64(5), 2289–2325.

Hong, H., Kubik, J. and Solomon, A. (2000). "Security Analysts' Career Concerns and Herding of Earnings Forecasts." RAND Journal of Economics, 31(1), 121–144.

Kahneman, D. (2011). Thinking, Fast and Slow. Farrar, Straus and Giroux.

Madrian, B.C. and Shea, D.F. (2001). "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." Quarterly Journal of Economics, 116(4), 1149–1187.

McLean, R.D. and Pontiff, J. (2016). "Does Academic Research Destroy Stock Return Predictability?" Journal of Finance, 71(1), 5–32.

MEMX (2025). US Equity Market Structure Overview. MEMX Research.

Morningstar (2026). "Target-Date Funds Continue Their Rapid Rise." Morningstar Research Report, March 10, 2026.

Shiller, R.J. (2000). Irrational Exuberance. Princeton University Press.

Shiller, R.J. (2017). "Narrative Economics." American Economic Review, 107(4), 967–1004.

Shleifer, A. and Vishny, R.W. (1997). "The Limits of Arbitrage." Journal of Finance, 52(1), 35–55.

Vanguard (2025). How America Saves 2025: Key Trends and Insights. Valley Forge, PA: The Vanguard Group.

WEF/BCG (2024). Global Retail Investor Outlook. World Economic Forum / Boston Consulting Group.

Calchas