THE SILENCE OF THE SIGNALS

On the Upstream Migration of Edge in the Age of Agentic Finance

Timothy Oluwatobi Ojebiyi

Move 37 • Essay Track • 4 June 2026

≈ 870 words

Companies announce good news loudly. Press releases, earnings calls, investor days: the machinery of positive disclosure is highly visible, deliberately so. But a sustained empirical anomaly suggests the alpha was never in what companies chose to say. In a study of SEC Form 10-K filings spanning 1995 to 2014, Cohen, Malloy, and Nguyen (2020) found that more than 86% of year-over-year textual changes in annual reports conveyed negative sentiment. A portfolio strategy that shorted firms making significant textual changes and held firms that did not earned abnormal returns of up to 188 basis points a month (over 22% a year), or 4–7% a year on a value-weighted basis, and the anomaly persisted for 18 months without reversing. The market was systematically failing to price something hiding in plain sight. The signal was always there. The edge lay in reading the silence.

That edge is now under structural threat, but not for the reason most people assume.

* * *

The conventional worry about AI in finance is that it will make markets more efficient by processing public information faster. This is true, and it matters less than it appears. The deeper consequence is not speed; it is homogenisation. When every participant's agent reads the same earnings transcripts, the same 10-Ks, and the same news feeds using the same foundation models, the alpha in faster processing does not merely shrink. It is arbitraged toward extinction.

Meng and Chen (2026), in a preprint, model this. In their framework, signal half-life is convex-decreasing in the share of AI adoption: at current adoption they put it at about 18 months, against five to seven years before AI. At monoculture, which is the limit this essay has in mind (identical models, identical training data, identical prompts), their net alpha is zero. Their convergence results are simulations calibrated to 13F filings, so this is a model's argument, not a measurement. Universal processing does not democratise alpha. In their model, it destroys it.

Compounding this is a second, slower-moving structural shift. Dyer, Lang, and Stice-Lawrence (2017) document a steady increase in 10-K length between 1996 and 2013: the median filing more than doubled, from a little over 23,000 words to over 49,000, driven largely by new accounting and SEC disclosure requirements, alongside rising boilerplate and redundancy and falling specificity. The signal-to-noise ratio of public financial text has been falling for years. Universal AI processing does not solve this problem. It accelerates it. Every agent now reads more noise, faster, and arrives at the same conclusion.

Processing speed is no longer where the edge lives. The question is where it went.

* * *

The answer is upstream. If the advantage no longer belongs to the agent that reads the fastest, it belongs to the agent that asks the question nobody else thought to ask. Edge migrates from execution to origination: from processing known signals to surfacing the signal no one has yet defined.

The Lazy Prices anomaly illustrates this precisely. On my reading, the anomaly exists not because investors lack access to 10-K filings, but because they are optimised to react to what is explicitly stated, not to what is structurally removed. Corporate legal teams understand this. They rephrase, consolidate, and quietly reorganise. Standard comparison tools catch the rephrase. Only semantic equivalence detection, the capacity to distinguish “they reworded this” from “they deleted this”, catches the deletion. That capability, until recently, required institutional infrastructure. The question was known. The access was gatekept by cost.

The cost has collapsed. The question has not changed.

* * *

The strongest objection to this argument is not that it is wrong; it is that the alpha does not disappear but morphs. When a monoculture of models interprets identical text and executes simultaneously, it generates algorithmic herding: correlated positioning, structural volatility, and price dislocations that a sufficiently contrarian operator can exploit by predicting how the herd will misread the text during regime breaks. This is a legitimate form of edge, and it will exist as long as model homogeneity persists.

But it is second-order alpha. It requires predicting the herd’s miscalculation: a reaction to a reaction. Question-origination alpha is structurally different. It requires identifying a signal the herd has no mechanism to ask about, because the herd is optimised for speed on known signals and has no prompt for the question itself. One type of edge chases the herd. The other exits the race entirely.

They are not equivalent.

* * *

In 2016, an AI played a move that, by its own estimate, a human would play about once in ten thousand times. One commentator thought it was a mistake. Lee Sedol, one of the strongest Go players alive, left the room and took about fifteen minutes to reply. The machine had found something that human intuition, and the conventional logic of the game, had largely missed.

Finance is now where Go was that day. Not because AI predicts prices better. Because the advantage has quietly migrated to a place most participants are not looking. The next Move 37 in finance is not a faster model or a larger training corpus. It is a better question. Specifically: *what did this company decide to stop telling you?*

Postscript, November 2026. I have since tested a version of this idea. Unsaid flags removed and softened risk language in 10-Ks; run on 321 consecutive filing pairs from 57 US banks (FY2019–2025), its return signal did not hold up. It failed a multiplicity correction, reversed sign on a held-out year, and shrank once text-extraction errors were fixed, and the plain word-frequency similarity that Cohen et al. use beat my LLM classifier on every cut. Those banks amount to roughly six independent bets, so this does not refute Lazy Prices. But it means the essay’s central mechanism, that semantic detection of deletions is where the edge lies, remains a hypothesis I have not been able to confirm.

References

Cohen, L., Malloy, C., & Nguyen, Q. (2020). Lazy Prices. *Journal of Finance*, 75(3), 1371–1415.

Dyer, T., Lang, M., & Stice-Lawrence, L. (2017). The Evolution of 10-K Textual Disclosure: Evidence from Latent Dirichlet Allocation. *Journal of Accounting and Economics*, 64(2–3), 221–245.

Meng, S., & Chen, X. (2026). AI-Driven Alpha Decay: Algorithmic Homogenization, Reflexive Signal Erosion, and the Paradox of Intelligent Markets. *arXiv:2605.23905*.