
When Superintelligence Trades the Markets

Andrei Nikiforov, PhD, CFA

Teaching Note · Financial Data Analytics: Quantitative Trading · Rutgers School of Business · March 2026

“Noise trading . . . provides camouflage which makes it possible for the insider to make profits.”

— Albert S. Kyle, *Econometrica*, 1985

Executive Summary

On March 29, 2026, Ethan Mollick—a Wharton professor whose posts routinely reach hundreds of thousands of readers—told his audience that a superintelligent AI would make money from financial markets “almost by definition” and that no one would even know. The post reached 111,000 views. It crystallized an assumption shared by millions: intelligence equals extraction, and markets are passive pools of wealth waiting for a smarter player.

What Mollick describes as a thought experiment, the finance profession has studied as a theorem—for forty years. Kyle (1985) gave his insider the strongest possible advantage: perfect knowledge of what a stock is actually worth, plus the optimal strategy to exploit it. That is not a weaker version of superintelligence—it *is* the superintelligence scenario. And the result is known: even this best-case trader must trade slowly, hiding inside everyday market noise, because every aggressive trade moves prices and reveals the edge. Milgrom and Stokey (1982) show that counterparties who suspect a superior player simply refuse to trade—or take the same side, collapsing the price before the ASI can profit. The square-root law of market impact imposes a hard ceiling on extraction speed.

Across all global liquid markets, a realistic ASI extraction rate is **\$15–18 million per day**, or **\$4–5 billion per year**—comparable to the output of Renaissance Technologies’ legendary Medallion fund, not “unlimited wealth.” The constraint is not cognition. It is liquidity, market impact, and the willingness of counterparties to keep showing their cards.

This teaching note develops the argument in four steps: the Kyle constraint (trade slowly or reveal your edge), the equilibrium wall (counterparties adapt or withdraw), the extraction ceiling (a concrete number), and the surveillance apparatus (the watchers are already machines). It concludes with a principle that AI discourse systematically overlooks: intelligence operates inside systems that *react* to intelligence.

1. The Claim

On March 29, 2026, Ethan Mollick, a professor at the Wharton School, posted the following:

“The easiest way to make money fast from a superhuman artificial intelligence would be in the financial markets, almost by definition. So the first lab to develop one . . . would almost certainly keep it quiet for as long as they could. Beats charging for API access.”

“It is not even clear that anyone would know they were trading against an [ASI], because, also by definition, it would be better at hiding any evidence that it was a counterparty than any human could detect. Someone would just get very rich, very fast.”

— Ethan Mollick (@emollick), March 29, 2026 (111K views)

Mollick is voicing what millions assume. He is not naïve—he is reasoning from a *definition*: if something is superintelligent, then by definition it should be able to do things humans cannot, including beat financial markets. The implicit model treats markets as static puzzles. Solve the puzzle better, extract more wealth.

What Mollick describes as a thought experiment, the finance profession has studied as a theorem—for forty years. In 1985, Albert Kyle gave his insider the strongest possible advantage: perfect knowledge of what a stock is actually worth, plus the mathematically optimal strategy to exploit it. That is not a weaker version of superintelligence. It is the superintelligence scenario, applied to a single asset. An ASI cannot have more than perfect knowledge of what something is worth. It can have perfect knowledge of *more things simultaneously*, but each individual position still faces the same constraint. And the result has been known since 1985: even this best-case trader must trade slowly, hide inside random everyday trading, and accept finite profits.

The same is true of the other results in this note. Milgrom and Stokey (1982) already assumed one side has superior information. Grossman and Stiglitz (1980) already analyzed what happens when the cost of acquiring information drops toward zero. These are not approximate models being stretched to fit a new scenario. They are the *exact* scenario, formalized decades before the term “ASI” entered public discourse.

Financial markets are not puzzles waiting for a smarter solver. They are adversarial systems in which every trade requires a willing counterparty, every order moves prices, and every pattern of consistent profitability triggers adaptive responses from other participants. Adding more intelligence to one side does not produce more extraction. It triggers more defense.

This note develops four constraints that bind even an omniscient trader. None can be overcome by being smarter.

2. The Kyle Constraint: Why Perfect Information Is Not Enough

Albert Kyle's 1985 paper sets up a market with three players.¹ One *insider* knows what the stock is actually worth. *Market makers* set prices based on the flow of buy and sell orders they see. And *everyday traders*—pension funds rebalancing, index funds tracking benchmarks, retirees cashing out—buy and sell for reasons that have nothing to do with private information. The critical detail: market makers cannot see who is behind each order. They see only the total flow—a mixture of the insider's trades and everyone else's.

Kyle's result has a clean punchline. The insider's optimal trade size is proportional to the volume of everyday trading—more background noise means more camouflage and a larger safe trade. But the camouflage is finite. And every dollar traded moves the market *against* the insider: aggressive buying drives the price up before the insider has finished buying. The market learns from the scoop.

These two constraints bind simultaneously. The insider cannot trade more than the background noise allows without revealing the edge. And every trade that exceeds the camouflage worsens the insider's own execution price. Kyle's own description is precise: the insider “acts like an intertemporal monopolist,” metering private information into prices over time.² In the continuous-time version of the model, information leaks into prices at a constant rate—not because the insider chooses patience, but because speed destroys the rent.

The optimal strategy for a trader who knows the future is to trade *as if they know very little*—slowly, in small pieces, blended with noise.

Competition makes it worse. Holden and Subrahmanyam (1992) show that multiple informed traders with the same private information “compete aggressively and cause most of their common private information to be revealed very rapidly.”³ If multiple ASI systems trade, or if other quantitative traders detect the same signals, individual rents shrink further. Monopoly on information is what sustains extraction. Competition destroys it.

On the other side of the trade, market makers are not passive.⁴ Glosten and Milgrom (1985) show that the mere *suspicion* that better-informed traders are in the market forces prices to adjust. Market makers protect themselves by charging a wider spread—the gap between buy and sell prices—as insurance against being picked off by someone who knows more than they do. The market does not need to identify the ASI. It only needs to notice that someone keeps ending up on the right side of price moves.⁵ Once that pattern emerges, spreads widen, available volume shrinks, and counterparties stop offering the same terms. Empirical work confirms: on days when insider

¹A.S. Kyle, “Continuous Auctions and Insider Trading,” *Econometrica* 53(6), 1985, pp. 1315–1335. Full text.

²Kyle (1985), Theorem 2 and continuous-time limit. Kyle shows that information is incorporated into prices at a *constant rate*—not because the insider is slow, but because going fast is self-defeating.

³C. Holden and A. Subrahmanyam, “Long-Lived Private Information and Imperfect Competition,” *Journal of Finance* 47(1), 1992. Full text.

⁴L. Glosten and P. Milgrom, “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders,” *Journal of Financial Economics* 14(1), 1985.

⁵D. Easley, M. López de Prado, and M. O'Hara, “Flow Toxicity and Liquidity in a High-Frequency World,” *Review of Financial Studies* 25(5), 2012.

trading occurs, spreads widen and depth drops.⁶

Key Takeaway: Even with perfect information, the optimal trading strategy is slow and camouflaged. Every aggressive trade moves prices against the insider and reveals the edge. Market makers respond by widening the gap between buy and sell prices, making each subsequent trade more expensive. Intelligence does not overcome price impact—it is constrained by it.

3. The Equilibrium Wall

Kyle explains how an informed trader profits by hiding inside noise. Milgrom and Stokey (1982) explain why that hiding is necessary—and what happens when it fails.⁷

Their no-trade theorem asks a simple question: if you *know* the person on the other side of your trade knows more than you, why would you take the trade? You wouldn't—because the very fact that they want to trade tells you it is a bad deal for you. Milgrom and Stokey state it directly: “the receipt of private information cannot create any incentives to trade.”

Why would anyone play poker against someone who can see all the cards? They wouldn't. And in financial markets, unlike in poker, rational players can walk away.

A caveat: real markets do not literally shut down. Pension funds must rebalance. Index funds must track their benchmarks. Retirees must sell to fund living expenses. These participants trade regardless of who else is in the market—they are the everyday flow in Kyle's model, the camouflage the ASI hides inside.⁸ But the theorem's implication is still binding: **the ASI can only profit from people who have to trade regardless, and their volume is finite.**

If the ASI's presence becomes suspected, two things happen. First, speculative traders—anyone who does not *have* to trade—widen the prices they demand or withdraw entirely. Second, and more damaging, traders who figure out that the ASI is buying have every reason to buy alongside it. If the ASI is buying, the smart response is also to buy. This creates a pile-on—what Banerjee (1992) calls informational herding—that moves the price to fair value before the ASI can finish extracting profits.⁹ The sell side of the order book evaporates. Prices jump. The edge disappears.

Grossman and Stiglitz (1980) add a deeper paradox.¹⁰ If prices already reflect everything worth knowing, no one has a reason to spend time and money researching stocks. But if no one researches stocks, prices cannot reflect what is worth knowing. The market works only if prices are *wrong*

⁶See Cheng, Firth, Leung, and Rui on spread and depth changes around insider-trading events. Also Hasbrouck (1991) on the protracted lag of trade-impact effects.

⁷P. Milgrom and N. Stokey, “Information, Trade and Common Knowledge,” *Journal of Economic Theory* 26(1), 1982. Full text.

⁸S. Morris, “Trade with Heterogeneous Prior Beliefs and Asymmetric Information,” *Econometrica* 62(6), 1994. Morris emphasizes that the no-trade theorem is a benchmark; real markets sustain trade through hedging demand, liquidity needs, and disagreement.

⁹A. Banerjee, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics* 107(3), 1992. Full text. Applied to markets: traders infer direction from order flow and rationally pile on, collapsing the mispricing.

¹⁰S. Grossman and J. Stiglitz, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review* 70(3), 1980. Full text.

enough to reward the people who do the work of figuring out what things are worth.

Now suppose the ASI's cost of processing information drops to near zero while humans still face real costs—analysts, data feeds, research teams. Human traders stop investing in analysis because they cannot compete. The ASI becomes the *sole* entity doing the work of making prices accurate. But its profits are still bounded by the finite pool of everyday trading—the pension funds, the index trackers, the retirees. And the equilibrium is fragile: a market in which one entity provides all price discovery is a single point of failure for the entire price system.

Vives (2008) summarizes the budget constraint cleanly: “the expected profits of informed traders match the expected losses of liquidity traders.”¹¹ The ASI's profits are financed by, and capped by, losses that everyday traders are willing to absorb. No amount of intelligence changes that budget.

Key Takeaway: If the market suspects a superior player, counterparties refuse to trade—or take the same side, collapsing the edge instantly. Even a dominant information processor is bounded by the finite pool of everyday trading. The market does not passively surrender profits to intelligence. It adapts or withdraws.

4. The Extraction Ceiling

Kyle says the ASI must trade slowly. Milgrom-Stokey says counterparties may refuse to trade at all. The square-root law of market impact puts a number on the constraint.

The empirical regularity, documented across equities, futures, options, and cryptocurrency, is that price impact rises with the *square root* of order size relative to daily volume.¹² Double your order, and impact rises by about 40%, not 100%. The cost of trading grows faster than the trade itself.¹³

Given realistic estimates of daily price swings and the edge that even the best current trading strategies capture per trade, the profit-maximizing share of a market's daily volume lands at roughly 0.25% to 4%. Trade more than that, and impact costs eat the profits. That is not a rule of thumb. It is the square-root law at work.

Exhibit 1 applies these constraints to the world's most liquid markets.

The base case: **\$15.6–17.8 million per day**, or **\$3.9–4.5 billion per year**. Exhibit 2 stress-tests the result. Even if the ASI's edge per trade were five times better than the best current quant fund, the annual ceiling reaches only \$20–22 billion. Getting to billions *per day* would require trading roughly 1% of total global volume with an enormous edge on every trade—a footprint visible to every market maker, surveillance system, and competing algorithm on earth.

¹¹X. Vives, *Information and Learning in Markets: The Impact of Market Microstructure*, Princeton University Press, 2008.

¹²See J.-P. Bouchaud, “Price Impact,” in *Encyclopedia of Quantitative Finance*, 2010; B. Tóth et al., “The Square-Root Impact Law Also Holds for Option Markets,” CFM Working Paper, 2016, [full text](#); F. Bucci et al. confirm the same scaling for institutional metaorders using 8 million U.S. trades.

¹³R. Almgren and N. Chriss, “Optimal Execution of Portfolio Transactions,” *Journal of Risk* 3(2), 2000. [Full text](#). The Almgren-Chriss framework formalizes the tradeoff between market impact and timing risk, yielding an optimal execution rate for any given signal strength and market depth.

Table 1. The ASI Extraction Ceiling: Base Case

Market	Daily turnover	Stealth share	Edge/trade	Daily P&L
Global FX (OTC)	\$9.51T	0.20%	0.03%	\$5.8M
Rates complex (OTC + listed)	\$25.3T	0.20%	0.025%	\$12.7M
Equities + index derivatives	\$2.4T	0.15%	0.10%	\$3.6M
Commodities	\$0.15T	0.10%	0.10%	\$0.15M
U.S. corporate bonds	\$0.058T	0.10%	0.05%	\$0.03M
Crypto spot	\$0.052T	0.05%	0.20%	\$0.05M
Raw total				\$22.2M
After 20–30% correlation haircut				\$15.6–17.8M

Source: Author’s model. FX and rates: BIS Triennial Central Bank Survey, April 2025, net-net basis. Equities: World Federation of Exchanges, 2024 annualized. Equity option notional haircut reflects 98.5% APAC concentration (WFE). U.S. bonds: SIFMA, full-year 2025 ADV. Crypto: CoinGecko top 15 centralized exchanges, 2024. Correlation haircut based on Asness, Moskowitz, and Pedersen (2013) cross-asset factor structure. All figures gross of fees, infrastructure, and edge decay.

Table 2. Sensitivity: What If the ASI Is Much Better?

ASI edge vs. best quant fund	Daily ceiling	Annual ceiling	Real-world comparison
1× (base)	\$15.6–17.8M	\$3.9–4.5B	≈ Medallion at \$12B AUM
2×	\$31.1–35.6M	\$7.8–9.0B	Top-decile hedge fund
5×	\$77.8–89.0M	\$19.6–22.4B	Entire top-10 HF profits
10×	\$155.7–177.9M	\$39.2–44.8B	< 1% of global HF capital

Source: Author’s model (see Exhibit 1 notes). The 5× and 10× rows are thought experiments, not stable equilibria: at those extraction rates, counterparty adaptation, spread widening, and surveillance response would compress realized alpha well below the assumed level. Medallion comparison: *Institutional Investor* (2024), reporting 66% gross / 39% net annualized, 1988–2021, on approximately \$12B of internal capital. Global hedge fund capital: \$5.15T at year-end 2025 (HFR).

The comparison to Renaissance Technologies’ Medallion fund is the note’s most important data point. Medallion generated roughly 66% gross annualized returns over three decades—arguably the closest real-world analogue to “superhuman trading.”¹⁴ But Medallion is capped at roughly \$12 billion of internal capital. At 39% net, that is about \$4.7 billion per year, or \$18.6 million per trading day—the same neighborhood as the base-case ceiling above.

The crucial contrast: Renaissance’s scalable, outside-investor fund (RIEF) annualized just 9.8% since its 2005 launch.¹⁵ Same scientific tradition, applied at larger scale, delivered ordinary returns. The bottleneck is not intelligence. It is capacity. Pástor, Stambaugh, and Taylor find strong evidence of decreasing returns to scale in active management at the industry level—more capital compresses the very mispricing it seeks to exploit.¹⁶

¹⁴*Institutional Investor*, “Renaissance’s 2024 Rebirth.” Also G. Zuckerman, *The Man Who Solved the Market*, Penguin, 2019. Self-cap reported at approximately \$10–12B of insider capital.

¹⁵Reuters, reporting on RIEF performance. Zuckerman confirms the Medallion–RIEF performance divergence.

¹⁶L. Pástor, R. Stambaugh, and L. Taylor, “Scale and Skill in Active Management,” *Journal of Financial Economics* 116(1), 2015.

Key Takeaway: A plausible ASI extraction ceiling is \$4–5 billion per year across all global liquid markets. That is best-hedge-fund-in-history territory, not an unlimited money printer. The constraint is not cognition. It is liquidity and the square root of order size.

5. The Surveillance Stack

Mollick’s second claim—that ASI “would be better at hiding any evidence that it was a counterparty than any human could detect”—assumes the watchers are human. They are not.

FINRA’s Consolidated Audit Trail (CAT) captures every order, modification, cancellation, and execution across all U.S. equities and options markets, linked to customer identity through reversible anonymized identifiers.¹⁷ The system was designed for 58 billion records per day. By Q3 2024, actual volumes averaged **666 billion events per day**, with a single-day peak of 905 billion.¹⁸ FINRA states that it runs “many algorithms that use AI” across 26 self-regulatory organizations operating 35 exchanges.¹⁹ The Insider Trading Detection Program monitors 100% of trading around material news events every day.²⁰

The infrastructure requirements are equally constraining. To trade at any meaningful scale, an entity needs a legal identity, funded accounts, KYC/AML verification, a broker or sponsored access agreement, and—for speed—co-location contracts with named exchanges.²¹ The SEC’s Large Trader reporting threshold is triggered at \$20 million in a single day—a level the ASI would exceed immediately.²²

History confirms the theory. SAC Capital’s Mathew Martoma sold \$700 million in two stocks over seven days—more than 20% and 11% of U.S. daily volume, respectively—and the pattern was reconstructable from the tape.²³ Navinder Sarao’s spoofing algorithm left repetitive order-book signatures on more than 400 trading days.²⁴ Medallion itself, operating entirely legally, generated trails through banks, auditors, Congress, and the IRS—settling a tax dispute for approximately \$7 billion.²⁵

Insider trading is not caught every time—Patel and Putniņš estimate detection and prosecution at roughly 15% of actual incidents.²⁶ But this actually strengthens the argument. If sporadic human

¹⁷SEC Rule 613. Full CAIS Compliance Go-Live: May 31, 2024. See SEC.gov.

¹⁸FINRA Rule Filing SR-FINRA-2024-023, full text.

¹⁹FINRA, “Advancing FINRA’s Mission With AI,” 2025.

²⁰FINRA, “Insider Trading Detection: FINRA’s Vital Role in Ensuring Market Integrity,” 2024, reporting about 65 staff and more than 450 referrals to the SEC per year.

²¹FINRA Rule 4512 (customer account information); SEC Rule 15c3-5 (market access, requiring sponsoring broker pre-trade controls); FinCEN CDD framework (beneficial ownership); Nasdaq and CME co-location agreements require named legal entities and signed contracts. CME lists 10 Gbps direct connectivity at \$12,000/month.

²²SEC Rule 13h-1, Form 13H. SEC.gov. Aggregation is across all NMS securities and brokers; it is not netted.

²³DOJ, “Martoma Sentenced,” 2014. SAC ultimately pled guilty and paid a \$1.8B penalty.

²⁴CFTC, “Charges Against Sarao,” 2015. Enforcement involved CME, DOJ, FBI, the U.K. FCA, and Scotland Yard.

²⁵Senate hearing record: GovInfo. Reuters reported the settlement amount.

²⁶Reported in University of Technology Sydney summary of Patel and Putniņš, with estimates of insider trading in about 1 in 5 M&A events and at least four times prosecution cases. UTS.edu.au.

insiders already create detectable statistical footprints with a nontrivial prosecution rate, an entity extracting *billions per year continuously* would create an immensely larger and more persistent anomaly surface.

Key Takeaway: The ASI would not sneak into an unwatched bazaar. It would operate inside a market that captures 666 billion events per day, links them to reversible identities, and runs AI-powered pattern detection. Every dollar of trading capacity requires a dollar of identity exposure.

6. Intelligence and Equilibrium

The ASI-in-markets question reveals a principle that AI discourse systematically overlooks: **intelligence operates inside systems that react to intelligence.**

Hayek (1945) made the foundational point. The relevant knowledge in an economy “never exists in concentrated or integrated form.”²⁷ Prices are not descriptions of the world. They are aggregators of dispersed local knowledge produced by millions of agents who each know something small. A superintelligence entering this system does not face a warehouse of dormant profits. It faces a decentralized process already built to compress knowledge into prices.

The efficient-markets tradition, sharpened by Grossman-Stiglitz, implies that every exploitable pattern gets traded on until the profits from exploiting it shrink to the cost of finding it. Lo’s Adaptive Markets Hypothesis (2004) extends the point: markets are ecologies of competing strategies governed by “competition, adaptation, and natural selection.”²⁸ New edges appear when environments change, but they get competed away. The ASI would be the apex predator—but the ecosystem adapts.

Markets are harsher on the attacker than most adversarial systems, because the defender can refuse the game. In Glosten-Milgrom, if informed traders are too numerous or too well informed, there may be “no bid and ask prices at which trading can occur.”²⁹ The market’s answer to overwhelming intelligence is not heroic resistance. It is worse terms, then thinner liquidity, then—in the limit—silence.

This is the methodological error in Mollick’s reasoning, and it is common among technologists thinking about AI capabilities. He reasons from *definitions*: ASI is superintelligent, therefore it can extract unlimited wealth. Market microstructure reasons from *mechanisms*: here is how information enters prices, how dealers defend themselves, how profits attract imitation, and how the equilibrium deforms when one side becomes too informed. The mechanisms impose constraints that definitions cannot wish away.

Superintelligence could run a formidable hedge fund. It could not turn a double auction into an ATM.

²⁷F.A. Hayek, “The Use of Knowledge in Society,” *American Economic Review* 35(4), 1945. Full text.

²⁸A.W. Lo, “The Adaptive Markets Hypothesis,” *Journal of Portfolio Management* 30(5), 2004. ResearchGate.

²⁹Glosten and Milgrom (1985), op. cit.

7. Discussion Questions

1. **The Camouflage Problem.** Kyle shows that an informed trader can only hide inside noise-trader volume. If the ASI wanted to trade U.S. equities (\$1.1 trillion daily), global FX (\$9.5 trillion), and Bitcoin (\$50 billion), which market offers the most camouflage? Which offers the least? Is the “deepest” market necessarily the most profitable one for the ASI?
2. **The No-Trade Cascade.** A credible rumor circulates that a major AI lab has deployed an ASI trading system. What happens to the cost of trading? To the volume of orders available to trade against? To the direction of remaining trades? Would a market dominated by hedge funds react differently than one dominated by retail investors? Why?
3. **The Ceiling in Context.** The note estimates an ASI extraction ceiling of \$4–5 billion per year. Global hedge fund capital is \$5.15 trillion. Berkshire Hathaway compounded at 19.9% annually for 60 years. Is \$4–5 billion per year “a lot” or “not much”—and for whom? What would change if two competing AI labs each deployed an ASI trading system simultaneously?
4. **The Medallion Puzzle.** Medallion generated 66% gross annualized returns for decades but reportedly capped itself at roughly \$12 billion. Its sibling fund RIEF, using similar scientific infrastructure at larger scale, annualized just 9.8%. Why the divergence? What does this imply for an ASI that wanted to trade at $10\times$ Medallion’s scale?
5. **Beyond Markets.** The note argues that adversarial equilibrium systems absorb unilateral intelligence advantages. Identify another domain—cybersecurity, biological immune systems, military strategy, or regulatory enforcement—where this principle applies. What plays the role of “liquidity” (the finite resource the intelligent agent needs)? What plays the role of “spread widening” (the system’s adaptive defense)?

© 2026 Andrei Nikiforov. Financial Data Analytics: Quantitative Trading, Rutgers School of Business, Camden. This teaching note draws on publicly available academic research, market data, and regulatory filings current as of March 2026. It does not constitute investment advice.