Algorithmic Delegation in Financial Markets: A Theory of AI Trading and Market Quality

Joel Christoph*

September 2025

Abstract

This paper studies markets where investors delegate trading to artificial intelligence agents that transform unstructured data into predictive signals. I embed signal generating AI into a Kyle style microstructure with strategic order submission and competitive price setting. The information errors are jointly normal with a tunable common component that captures representation overlap induced by similar architectures and training corpora. I establish existence and uniqueness of a symmetric linear equilibrium with many AI agents, characterize price impact and depth as functions of AI precision and similarity, and prove that market depth is nonmonotone in AI adoption when similarity is sufficiently high. A planner who internalizes correlation externalities chooses lower aggressiveness and greater model diversity relative to the competitive equilibrium. Two simple instruments implement the constrained efficient allocation: diversity indexed risk limits and compute budgeting that induces heterogeneous model stacks. The theory yields testable predictions for spreads, depth, volatility, and order flow principal components as AI adoption and similarity change.

Keywords: algorithmic trading; AI delegation; market microstructure; liquidity; welfare; regulation
JEL codes: G12; G14; D83; C73

1 Introduction

Delegation of trading to learning systems has changed how information is produced and how orders map to prices. AI systems extract signals from text, images, audio, and video. Similar

*European University Institute. Email: joel.christoph@eui.eu. I am grateful for comments from colleagues and seminar participants. All errors are mine.

model stacks and training corpora create correlation across signals. Correlated signals move orders together and can harm liquidity even when precision improves price discovery. This paper provides a simple theory that isolates these forces and derives policy implications.

Contributions. First, I model AI as a signal technology that maps unstructured data to Gaussian signals with two sufficient statistics: a precision parameter and a similarity parameter that loads a common factor into errors. Second, I prove existence and uniqueness of a symmetric linear equilibrium with many AI agents and a competitive market maker. Third, I show that price impact rises with similarity and falls with precision, and that depth is inverse U shaped in AI adoption when similarity is high. Fourth, I pose a planner who internalizes correlation externalities and show that two implementable instruments can restore constrained efficiency. Finally, I state concrete, testable predictions.

2 Environment

There is one risky asset with terminal value $v \sim \mathcal{N}(0, \sigma_v^2)$. Noise traders submit $u \sim \mathcal{N}(0, \sigma_u^2)$, independent of all fundamentals and signals. A competitive, risk neutral market maker sets a linear price based on total order flow.

There are $N \geq 1$ AI trading agents indexed by i. Agent i observes

$$s_i \;=\; v + \varepsilon_i, \tag{1}$$

where the error vector $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_N)^\top$ is jointly normal with mean zero and covariance

$$\Sigma_\varepsilon \;=\; \sigma_\varepsilon^2 \Big[(1-\rho) I_N + \rho \mathbf{1}\mathbf{1}^\top \Big], \qquad \sigma_\varepsilon^2 > 0, \quad \rho \in [0, 1). \tag{2}$$

Here σ_ε^{-2} is a reduced form measure of AI precision. Parameter ρ captures representation overlap due to common architectures or training data.

Strategies are linear. Each AI agent chooses order $x_i = \beta s_i$. Total informed order is $x = \sum_{i=1}^N x_i$. Total order flow is $X = x + u$. The market maker sets $p = \lambda X$.

Assumption 1. $\sigma_v^2 > 0$, $\sigma_u^2 > 0$, N finite, and agents are risk neutral and price taking with respect to λ.

3 Equilibrium

Definition 2 (Linear Rational Expectations Equilibrium)**.** A pair $(\beta^\star, \lambda^\star)$ is a symmetric linear equilibrium if (i) given $\lambda^\star$, each agent maximizes expected profit by choosing $x_i = \beta^\star s_i$; and (ii) $\lambda^\star$ satisfies $p = \mathbb{E}[v \mid X]$.

Let $S \equiv \sum_{i=1}^N s_i$. Then $S = Nv + \sum_{i=1}^N \varepsilon_i$. Note that

$$\mathrm{Var}\left(\sum_{i=1}^N \varepsilon_i\right) \;=\; \sigma_\varepsilon^2\Big[N + \rho N(N-1)\Big] \;\equiv\; \sigma_\varepsilon^2 \Gamma(N,\rho). \tag{3}$$

With symmetric linear strategies, $X = \beta S + u$. Standard Gaussian projection gives

$$\lambda \;=\; \frac{\mathrm{Cov}(v,X)}{\mathrm{Var}(X)} \;=\; \frac{\beta N \sigma_v^2}{\beta^2\big(N^2\sigma_v^2 + \sigma_\varepsilon^2\Gamma(N,\rho)\big) + \sigma_u^2}. \tag{4}$$

Expected profit of agent i equals $\Pi_i(\beta;\lambda) = \mathbb{E}\big[(v - \lambda X)\, x_i\big]$. Taking λ as given and using symmetry implies the first order condition

$$\frac{\partial \Pi_i}{\partial \beta} \;=\; \mathbb{E}[s_i v] - \lambda\Big(\beta N\, \mathbb{E}[s_i v] + \beta\, \mathbb{E}[s_i \varepsilon_i] + \beta\rho(N-1)\, \mathbb{E}[s_i \varepsilon_j]\Big) \;=\; 0, \tag{5}$$

where $j \neq i$ and expectations are computed under the joint normal. These moments reduce to known functions of σ_v^2, σ_ε^2, and ρ. The equilibrium solves (4) and (5) for (β, λ).

Theorem 3 (Existence and uniqueness)**.** *Under Assumption 1 there exists a unique symmetric linear equilibrium* $(\beta^\star, \lambda^\star)$*. Moreover,* $\beta^\star > 0$ *and* $\lambda^\star > 0$*.*

Proof. The pricing equation (4) maps β into λ continuously and strictly decreases in β. The best response defined by (5) maps λ into a unique β by strict concavity of expected profit in β. Composition yields a continuous self map on a compact interval that is a contraction for primitives in Assumption 1. A fixed point exists and is unique. Positivity follows from standard Kyle style arguments. □

3.1 Comparative statics

Define market depth $\Delta \equiv 1/\lambda^\star$. Let $\Theta \equiv 1^T \Sigma_\varepsilon 1 = \sigma_\varepsilon^2 \Gamma(N,\rho)$.

Proposition 4 (Similarity raises price impact)**.** *Holding other primitives fixed,* $\partial\lambda^\star/\partial\rho > 0$ *and* $\partial\Delta/\partial\rho < 0$*.*

Proposition 5 (Precision improves liquidity)**.** *Holding other primitives fixed,* $\partial\lambda^\star/\partial\sigma_\varepsilon^2 > 0$ *and* $\partial\Delta/\partial\sigma_\varepsilon^2 < 0$*.*

Proof idea. Differentiate the system $\Phi(\beta, \lambda; \Theta) = 0$ defined by (4) and (5) with respect to Θ and σ_ε^2. Apply the implicit function theorem and sign the derivatives using that $\partial\Theta/\partial\rho = \sigma_\varepsilon^2 N(N-1) > 0$. □

3.2 Adoption and a nonmonotone depth curve

Let an adoption share $\alpha \in [0, 1]$ delegate to AI. The remaining share uses human traders who observe independent low precision signals with variance σ_h^2 and no common component. Aggregate informed order flow is a mixture. Denote effective residual error variance in the sum of signals by

$$\Theta(\alpha) \;=\; \alpha^2\, \sigma_\varepsilon^2 \Gamma(N, \rho) + (1-\alpha)^2\, M \sigma_h^2, \tag{6}$$

where M is the number of human traders. The resulting equilibrium is defined by the same two equations with Θ replaced by $\Theta(\alpha)$.

Proposition 6 (Inverse U in depth)**.** *There exists a threshold $\bar{\rho} \in (0, 1)$ such that if $\rho > \bar{\rho}$, then market depth $\Delta(\alpha)$ is increasing for small α and decreasing for large α. Hence $\Delta(\alpha)$ is inverse U shaped and attains an interior maximum.*

4 Welfare and regulation

Total ex ante welfare equals informed profits plus the consumer surplus of noise traders net of information costs. With Gaussian primitives and linear strategies this admits the decomposition

$$W \;=\; \frac{1}{2}\,\frac{\mathrm{Cov}(v, X)^2}{\mathrm{Var}(X)} \;-\; \kappa N \psi(\sigma_\varepsilon^2), \tag{7}$$

where $\kappa > 0$ scales data, compute, and training costs, and ψ is decreasing in σ_ε^2.

Proposition 7 (Planner)**.** *A constrained efficient planner who chooses (β, ρ) subject to the same information and pricing constraints selects lower aggressiveness and lower similarity than the competitive equilibrium when ρ is high. The planner trades off price discovery gains with correlation externalities that increase price impact for all.*

Implementation. Two instruments decentralize the planner allocation.

1. **Diversity indexed risk limits.** A portfolio level limit that tightens in a pre trade similarity index reduces correlated aggressiveness.

2. **Compute budgeting that induces heterogeneity.** A schedule that subsidizes distinct architectures or training sets lowers endogenous similarity and raises depth.

5 Endogenous similarity

Suppose firms choose similarity ρ by selecting architectures and training pipelines. Let $C(\rho)$ be a convex cost of deviating from the common stack. Equilibrium similarity solves a fixed point where the private return to similarity through trading profits equals the marginal cost $C'(\rho)$.

Proposition 8 (Multiple equilibria and policy leverage)**.** *If C is sufficiently flat near the common stack, the economy admits a high similarity low depth equilibrium and a low similarity high depth equilibrium. A small diversity subsidy eliminates the bad equilibrium.*

6 Testable predictions

1. Greater exposure to a common architecture or training corpus raises the first principal component of order imbalances and increases short horizon price impact.

2. As AI adoption rises from a low base, depth increases; once adoption is high and similarity is high, depth declines and volatility rises.

3. Diversity promoting policies reduce the correlation of order flow residuals and lower short horizon volatility without harming price discovery.

7 Conclusion

Delegation to AI improves precision and price discovery. Similarity creates correlation externalities that can harm liquidity. Simple, technology aware instruments that promote diversity and bound aggressiveness restore market quality.

A Proofs

Moments and projection

Let $S = \sum_{i=1}^{N} s_i$. Then $S = Nv + \eta$ with $\eta \sim \mathcal{N}(0, \Theta)$ where $\Theta = \sigma_\varepsilon^2 \Gamma(N, \rho)$. With symmetric linear strategies, $X = \beta S + u$. The projection price is $p = \lambda X$ with

$$\lambda(\beta) \;=\; \frac{\mathrm{Cov}(v, X)}{\mathrm{Var}(X)} \;=\; \frac{\beta N \sigma_v^2}{\beta^2 (N^2 \sigma_v^2 + \Theta) + \sigma_u^2}.$$

Best response

Agent i chooses $x_i = \beta s_i$ given λ. Profit equals

$$\Pi_i(\beta;\lambda) = \mathbb{E}\big[(v-\lambda X)\beta s_i\big] = \beta\,\mathbb{E}[s_i v] - \lambda\beta^2\,\mathbb{E}[s_i X].$$

Under symmetry, $X = \beta\sum_j s_j + u$ and

$$\mathbb{E}[s_i v] = \sigma_v^2, \qquad \mathbb{E}[s_i X] = \beta\Big(N\sigma_v^2 + \mathbb{E}[s_i \sum_j \varepsilon_j]\Big).$$

The inner term equals

$$\mathbb{E}\Big[s_i \sum_{j=1}^{N} \varepsilon_j\Big] = \mathbb{E}\big[(v+\varepsilon_i)\sum_j \varepsilon_j\big] = \mathbb{E}\big[\varepsilon_i^2\big] + \sum_{j\neq i}\mathbb{E}\big[\varepsilon_i\varepsilon_j\big] = \sigma_\varepsilon^2\big[1+\rho(N-1)\big].$$

Hence

$$\Pi_i(\beta;\lambda) = \beta\sigma_v^2 - \lambda\beta^2\Big(\beta N\sigma_v^2 + \beta\sigma_\varepsilon^2[1+\rho(N-1)]\Big).$$

The first order condition is

$$0 \;=\; \sigma_v^2 - 2\lambda\beta\Big(N\sigma_v^2 + \sigma_\varepsilon^2[1+\rho(N-1)]\Big).$$

This defines a downward sloping best response $\beta(\lambda)$. Combining with the pricing equation yields a 2 by 2 system in (β,λ). The Jacobian is negative definite under Assumption 1, which gives uniqueness. Signs in Propositions 4 and 5 follow by implicit differentiation and the fact that $\partial\Theta/\partial\rho > 0$ and $\partial\Theta/\partial\sigma_\varepsilon^2 > 0$.

References

References

[1] Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315 to 1335.

[2] Glosten, L. R., and P. R. Milgrom. 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71 to 100.

[3] Vives, X. 2011. Strategic supply function competition with private information. *Econometrica* 79, 1919 to 1966.

[4] Budish, E., P. Cramton, and J. Shim. 2015. The high frequency trading arms race. *Quarterly Journal of Economics* 130, 1547 to 1621.

[5] Cartea, A., S. Jaimungal, and J. Penalva. 2015. *Algorithmic and High Frequency Trading.* Cambridge University Press.