


AI-Powered Trading, Algorithmic Collusion, and Price Efficiency

Winston Wei Dou, Itay Goldstein, Yan Ji

Discussant: **Hao Yang**

Swiss Finance Institute and USI Lugano

IESE Workshop on Artificial Intelligence in Finance



The Paper: AI vs Theory

Very interesting research question:

- ▶ How do multiple **decentralized** AI algorithms (Q-learning) learn to behave at convergence (i.e., once learning has stabilized)

in an infinitely repeated, market-order, generalized Kyle (1985) trading game?

- ▶ Do their strategies **diverge** from or **align** with the equilibrium behavior predicted by economic theory—

specifically, **collusion** via the price-trigger equilibrium or the experience-based equilibrium?

Main Findings of the Paper: AI vs Theory

Similar to theoretical predictions, AI can learn to collude, trade less aggressively on private information, and gain supra-competitive profits.

- ▶ Via the price-trigger strategy under low noise
- ▶ Via an experience-based strategy under high noise

Evidence: AI's impulse response to shocks is similar to that of an economic agent.

Key Differences between Economic Agents and AI (Q-learning) Agents

- ▶ **Economic Agent:** A rational Bayesian agent with common knowledge of the model primitives.

Capable of strategic reasoning, rationally anticipating the opponent's optimal behavior, with such reasoning being common knowledge.

- ▶ **Q-learning (AI) Agents:** Learn through trial and error (exploration and exploitation) with zero knowledge of the model primitives.

They cannot perform strategic reasoning, have no model of the opponent, and focus solely on their own observed state variables.

Q-value: Naturally represents a misspecified, non-parametric estimation of the expected payoff.

Comment on “AI (Q-learning) Behavior Exhibiting Similarity to Economic Theory Predictions”

- ▶ **Bad news:** AIs learn to collude.
- ▶ **Good news:** AI behavior is similar to the economic theory model; economic theory predictions appear robust.

Past policies based on economic theory remain largely valid.

- ▶ **Mysterious:** AI (open loop) behavior can coincide with economic theory (closed loop, such as the price-trigger strategy).

See Fudenberg, Drew, and David K. Levine. "Open-loop and closed-loop equilibria in dynamic games with many players." Journal of Economic Theory 44.1 (1988): 1-18.

Comment on “AI (Q-learning) Behavior Exhibiting Similarity to Economic Theory Predictions”

- ▶ (Tirole, 2020): “Fintech is standard technological change and economic fundamentals apply.”
- ▶ (Shapiro & Varian, 1998): “Technology changes. Economic laws do not.”

This paper is one of the first to provide direct experimental evidence.

Comment on “AI’s Risk Aversion in Strategy Choice”

- ▶ AI speculators tend to adopt conservative trading strategies, enabling them to attain supra-competitive profits. The use of RL algorithms effectively generates a form of “risk aversion” in strategy choice.
- ▶ Q-learners prefer safe choices over riskier ones, especially when economic noise is high.

Comment on “AI’s Risk Aversion in Strategy Choice”

Q-learners tend to favor safer options over riskier alternatives, which might offer higher yet more volatile payoffs, particularly in environments with significant economic noise.

How robust is this finding? It is also supported by:

- ▶ **Colliard, Jean-Edouard, Thierry Foucault, and Stefano Lovo.** “Algorithmic pricing and liquidity in securities markets.” (2024).

Market makers using Q-learning will not undercut the price in the Glosten-Milgrom market-making game.

- ▶ **Yang, Hao.** “AI Coordination and Self-fulfilling Financial Crises.” (2024).

Speculators using Q-learning will less often choose to attack the regime, thereby avoiding the triggering of crises in the currency attack global game.

- ▶ **Marimon, Ramon, Ellen McGrattan, and Thomas J. Sargent.** “Money as a medium of exchange in an economy with artificially intelligent agents.” *Journal of Economic Dynamics and Control* 14, no. 2 (1990): 329–373.

In Kiyotaki and Wright’s (1989) exchange economies, where both speculative and fundamental equilibria exist, AI agents converge to the fundamental equilibrium, using low-storage-cost goods as money.

Comment: “Application or Interpretation”

- ▶ When trading, do we really train offline for such a long time? (Online training can be costly and hence not so long)
- ▶ Circumstances (i.e., game parameters) likely vary widely across different episodes.
- ▶ Do we have enough observations to properly train an algorithm?
- ▶ Even if AI collusion behavior **coincide** with economic theory, can we confidently say that it has truly “learned” these principles?

Comment: Asymptotic Convergence Property

- ▶ With $N_{sim} = 1000$, what happens with 2000, 3000, ... or even 10,000 simulations?

I conjecture that there may be asymptotic convergence to a fixed distribution of strategy choices for Q-learners. Although it may not be proven formally, experimental/empirical evidence could support this.

(For example, the BLP model worked for decades in empirical IO before its convergence property was formally proven in 2014.)

Economic theory focuses on fixed points as the equilibrium. In the age of AI agents, we may observe a fixed distribution rather than a single point.

General Take on the Methodology

- ▶ The rise of bottom-up approaches (aka ABM, multi-agent systems, or complex systems) involves starting with zero knowledge and modeling agents as AI learning agents to observe what equilibrium emerges.
- ▶ In reality, we may be witnessing the upcoming battle among different AI agents.
- ▶ Recalling Robert Lucas Jr.:

“Are you after truth? Yeah. But I don’t know what we mean by truth in our business. I don’t see economics as pushing that deeply in some respects. We’re programming robot imitations of people, and there are real limits on what you can get out of that.”

Final Remark

- ▶ A very interesting paper on an important and timely topic, featuring solid theory and well-executed experiments.
- ▶ Key question: What is the best application or interpretation of the results?
- ▶ Good luck with the publication!
- ▶ I definitely recommend everyone read it carefully!