Discussion on "Algorithmic Pricing and Liquidity in Securities Markets" by Colliard, Foucault, and Lovo

Liyan Yang
Rotman School, University of Toronto

Workshop on Artificial Intelligence in Finance
March 20th, 2025

Trading and Market Making in Financial Markets

- Liquidity demanders: information (profits), noise (private value)
- Liquidity suppliers: market making
- Rational expectations equilibrium (REE) models assume common knowledge of the markets (Grossman 1976; Grossman and Stiglitz 1980; Hellwig 1980; Admati 1985; Glosten and Milgrom 1985; Kyle 1985, 1989)
- Market is complex, and it seems too demanding to assume common knowledge of market environment.

- In theory, introduce behavioral elements into learning about market environment:
  - Eyster, Rabin and Vayanos (2018): cursed equilibrium ("level-0")
  - Mondria, Vives and Yang (2021): costly learning
  - Banerjee, Davis and Gondhi (2024): choosing to disagree
  - Bastianello and Fontanier (2024): "Partial Equilibrium Thinking" (PET, "level-1")
- In practice, let machines learn about the environment:
  - Liquidity demand: Dou, Goldstein and Ji (2023)
  - Liquidity provision: Colliard, Foucault and Lovo (2025)

Colliard, Foucault and Lovo (2025)

- The paper models market making concretely with Q learning.
  - Very nice feature; not simply conceptualize "machines" as some players in existing models
- The results are novel
  - Something similar to REE: machine prices reflect adverse selection
  - Some results unique: collusion (non-zero profits); price patterns different from REE
- A coherent mechanism explains findings: learning noise (variance)
- Other things: endogenous hyper-parameters; tick size; human-machine interactions

Panel A: Average Quoted Spread.

This graph plots the average over 1,000 experiments of the quoted spread, with 95% confidence intervals, both in the adverse-selection case and the no-adverse-selection case. The graph additionally plots the values of the quoted spread in both cases, in the Glosten-Milgrom benchmark of Section 2.3 (accounting for price discreteness).

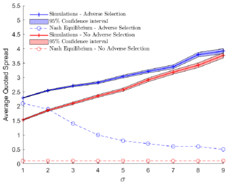

Panel B: Average Realized Spread.

This graph plots the average over 1,000 experiments of the realized spread, with 95% confidence intervals, both in the adverse-selection case and the no-adverse-selection case. The graph additionally plots the values of the realized spread in both cases, in the Glosten-Milgrom benchmark of Section 2.3 (accounting for price discreteness).

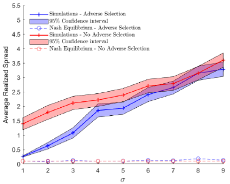

Figure 4: Average Quoted Spread $\bar{QS}$ and Average Realized Spread $\bar{RS}$ in the adverse-selection case and the no-adverse-selection case, for different values of the dispersion of clients' liquidity shocks σ. The other parameters are $\Delta_v = 4$, $N = 2$, $\mu = \frac{1}{2}$, $\mathbb{E}(v) = 2$, $T = 1,000,000$, and $K = 1,000$.

Comments

- Two related comments, beyond the current paper
- Comment 1: Welfare, regulation, and relevant markets
  - Why machines? Trade-off between humans and machines; what markets are more likely to adopt machines? How to regulate?
- Comment 2: Implementation and stationarity
  - Hard to get analytical results. How fragile are the results? Local maximum?

Comment 1: Welfare and Market

- When it comes to regulation, it is necessary to think about welfare and which markets are more relevant.
- How to measure welfare? Need to open the black box of why choosing Q learning and how to choose deep parameters.
- Profits? May ignore something deeper. In the current setting human-AI, human always beats; they know more of reality. No one will choose AI.
- Deep parameters: no deviation (Nash) + frequentist (avoid Bayesian). Need to think about welfare.

Humans versus Machines

- The current treatment of humans vs machines is not about a trade-off.
- It is more like that humans are fully sophisticated, and so the point is whether the introduction of a sophisticated player will kill the results.
  - Humans will train machines to become a "slave" / "servant".
- If humans are so sophisticated, then it comes the question of why machines in the first place...
  - Need to adjust the interpretation of the current expositions

Dou, Goldstein, Li and Yang (2023, work-in-progress)

- Humans cannot understand the market environment.
- Q learning is the technology to learn from prices:
  - How to get the price function in REE? Answer: Q-learning.
- Humans, without machine help, cannot understand the price, and so behave as "level-0" players.
- They use their conceptualization of the world and resources to get a fundamental signal
  - As Dávila and Parlatore (2023): investor i gets a signal

$$\tilde{s}_{it} = \tilde{v}_t + \tilde{\varepsilon}_{it} \text{ (information)}$$

  - But could have a wrong prior

$$\tilde{v}_t \sim_i N\left(\tilde{n}_t + \tilde{e}_{it}, \tau_v^{-1}\right)$$

- If machines are identical, then they will converge to Kyle (1989) traders with market power and correct learning.
- Simulation with heterogeneous machines, work in progress.

Hyper Parameters

- The current treatment of choosing algorithms' hyper-parameters is nice.
- It has a Nash flavor ("stability" and no deviation), and at the same time, avoids the augmentation of existing games (frequentist vs Bayesian, hypothesis testing).
- Nonetheless, it is not easy to think about welfare, and so regulation.

Comment 2: Implementation and Stationarity

- Q learning is similar to numerically finding an optimal solution.
- It can be sensitive to the starting point.
  - Not an issue in the current paper. But how can we ensure it in practice?
- May be stuck with a local optimum as opposed to a global optimum.
- How to choose the action space $\mathscr{A} = \{a_1, ..., a_M\}$?

- Should try some action beyond $\mathscr{A} = \{a_1, \dots, a_M\}$, and then update $\mathscr{A}$ accordingly?
- What about another layer of exprimentation with trying $a' \notin \mathscr{A}$ and updating $\mathscr{A}$ gradually?

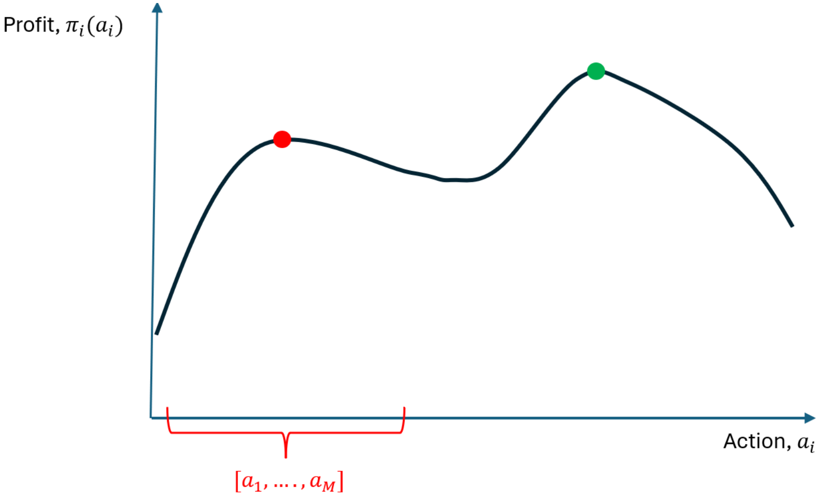

Conclusion

- A nice paper with a concrete “Q-learning” market making
- Novel result and clean mechanism
- Broader thoughts on regulation, welfare, and implementation