Algorithmic Pricing and Liquidity in Securities Markets

Jean-Edouard Colliard (HEC Paris)
Thierry Foucault (HEC Paris)
Stefano Lovo (HEC Paris)

Discussion by Carole Gresse (Université Paris Dauphine-PSL)

CONSEIL SCIENTIFIQUE DE L'AMF – 5 FÉVRIER 2024

Paper positioning and contribution

- Central goal = Understand how algo. MMs using Q-learning impact price formation and liquidity.
- Algorithmic trading => need for new models / O'Hara (2015):
  - « *the data that emerge from the trading process are consequently altered […] For microstructure researchers, I believe these changes* ***call for a new research agenda****, one that recognizes how the learning models used in the past are lacking […]* »
- Bridging two very different theoretical fields
  - Financial economics / Microstructure / Game theory
  - Artificial Intelligence / Machine learning
- Novel and impressive!
- Novelty: Modelling price making with Q-learning algorithms under adverse selection
  - Algo MMs learn by trial and error
  - Comparison with omniscient Bayesian learners
  - The closest paper (Dou et al., 2023) focuses on informed investors' strategies and not on MM strat.

Paper positioning and contribution – cont'd

OMNISCIENT NASH-BAYESIAN LEARNERS

- MMs know everything about their environment
  - Asset value distribution $\tilde{v}$
  - Liquidity shock distribution $\tilde{L}$
  - Distr. of signals received by clients
  - Other players' strategies
- Bertrand competition
- Nash players
- **Bayesian**: know how to incorporate the information revealed by the order flow by computing Bayesian conditional expectations

$$E\left(\tilde{v} \,\middle|\, Order\,Flow = X\right)$$

- Glosten & Milgrom (1985), Kyle (1985)

Q-LEARNERS (a certain type of reinforcement learning)

- MMs know:
  - Asset value distribution
  - Liquidity shock distribution
  - Distr. of signals received by clients
- Profit maximizers starting with a discrete price grid (set of actions a)
- Profit expectation (Q function) for a given value of the order flow they execute (state s) and a given price (action a) → Expected profit = Q(s,a)
- **Learning** = revising Q(s,a) based on Δ between expected profit & profit realized in state s with action a

$$Q_{t+1}(s,a) = (1-\alpha)\,Q_t(s,a) + \alpha Q_t^{observed}(s,a)$$

where α = learning factor ∈ [0;1]

Results

		Nash-Bayesian equilibrium / ASC theory
1	Algo. MMs learn how to quote under AS but quotes are less competitive (more expensive) than Nash-equilibrium quotes. Each MM could obtain a higher profit by undercutting others' prices.	$\neq$
2	Spreads are larger than those predicted by the Nash equilibrium.	$\neq$
3	Spreads increase with $\sigma(\tilde{L})$	$\neq$
4	Realized spreads are larger with no adverse selection. Interp.: Algo. MMs behave more competitively with AS because AS makes it easier for them to undercut others' prices.	$\neq$
5	With adverse selection, spreads increase with $\sigma(\tilde{v})$	$=$
	With no adverse selection, spreads increase with $\sigma(\tilde{v})$	$\neq$

Complementary explanation for 4:
With AS, the probability of a non-zero order flow is higher ⇨ Makes learning more efficient?

Comments

The difference between Q-learning quotes and Nash-Bayesian quotes comes at no surprise.

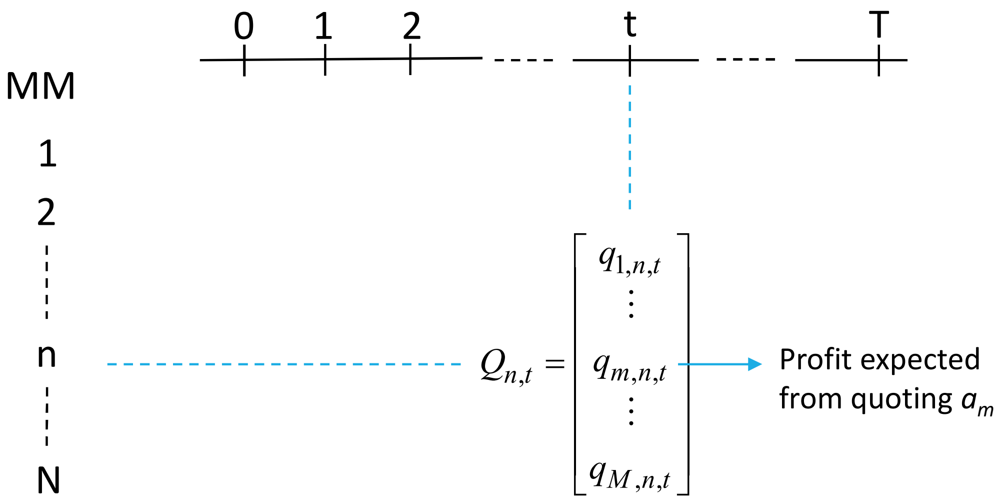

Q-learning from episode t

$$q_{m,n,t+1} = \alpha \pi_{n,t} + (1-\alpha) q_{m,n,t} \quad \textit{if } n \textit{ quotes } a_m$$
$$\textit{no learning for other prices}$$

Sources of divergence between Q-Learning and Nash eq.

- Learning factor α is exogenous and not endogenous as in G&M or Kyle (1985) => learn from profits but do not learn on asset value from order flow.
- Diff. between $\pi_{n,t}$ and $q_{m,n,t}$ has several causes:
  - Liquidity shock, Asset pay-off, Quote choice…
  - Uneasy to disentangle
- Exploring is chosen with probability $e^{-\beta t}$.
- Asset value is randomly drawn at each episode t => no learning on price discovery across episodes.

=> The point of interest is to understand which factors enlarge/reduce the difference.

Comments – cont'd

Results are somehow more consistent with the inventory risk theory.

		Comparison with: Nash-Bayesian equilibrium / ASC theory	Comparison with: Inventory holding costs theory
1	Algo. MMs learn how to quote under AS but quotes are less competitive (more expensive) than Nash-equilibrium quotes. Each MM could obtain a higher profit by undercutting others' prices.	$\neq$	Could reflect a cost/risk similar to inventory holding cost
2	Spreads are larger than those predicted by the Nash equilibrium.	$\neq$	$=$
3	Spreads increase with $\sigma(\tilde{L})$	$\neq$	$=$
4	Realized spreads are larger with no adverse selection. Interp.: Algo. MMs behave more competitively with AS because AS makes it easier for them to undercut others' prices.	$\neq$	Imperfect learning --- costs
5	With adverse selection, spreads increase with $\sigma(\tilde{v})$	$=$	$=$
5	With no adverse selection, spreads increase with $\sigma(\tilde{v})$	$\neq$	$=$

Comments – cont'd

- Why consider the inventory risk explanation?
- The weight of the liquidity shock variance in the determination of the private value of the asset for the client seems heavy.
  - $\tilde{L}$ follows a Gaussian distribution while $\tilde{v}$ is a binary variable.
  - In the experiments, $\sigma\left(\tilde{L}\right)=5$ while $\Delta\left(\tilde{v}\right)=4$, that is $\sigma^2\left(\tilde{L}\right)=25$ and $\sigma^2\left(\tilde{v}\right)=4$.
- Prices are discrete while the realized profits are determined by continuous Gaussian liquidity shocks.

=> Liquidity shocks seem to play a great role relative to AS.

Questions and suggestions

1) Inventory risk

- You mention it page 25 with « demand risk », as the potential explanations for prices to be less competitive when the volatility of the asset increases and when the size of liquidity shocks increases
- Cannot disagree / Both are **inventory risk** in the end.

2) Design of distributions

- Why take a binary distribution for the asset and a Gaussian distribution for liquidity shocks?

3) Adverse selection model

- What can you say about price quality/informational content of prices with algo. MMs, even before reaching Section 5?

4) Extension: Would it be possible to endogenize β?

- Probability of choosing to explore (instead of quoting the more rewarding price) = $e^{-\beta t}$
- Algo. MMs could have the choice between 2 values of β: pay a cost to have the lowest β & greater expected profits thanks to more intense exploration.

Conclusion

- Good to see that algos cannot do better ☺
- Impressive work!