Algorithmic Pricing in Securities Markets

Jean-Edouard Colliard, Thierry Foucault, Stefano Lovo
HEC Paris

AMF, February 5, 2024

Roadmap

Algorithms in Securities Markets

- Rise of the " Algo market maker".

Artificial Intelligence (AI) in Securities Markets

Opinion **The FT View**

How to prevent AI from provoking the next financial crisis

New systems have benefits for markets, but risks to stability must be managed

THE EDITORIAL BOARD

Gary Gensler, chair of the US Securities and Exchange Commission, sees regulatory intervention as essential to avoiding a financial crisis caused by AI © FT montage/Reuters/Dreamstime

Amid talk of job cuts due to artificial intelligence, Gary Gensler thinks robots will actually create more work for financial watchdogs. The US Securities and Exchange Commission chair puts the likelihood of an AI-driven financial crisis within a decade as "nearly unavoidable", without regulatory intervention. The immediate risk is more of a new financial crash than a robot takeover.

Gensler's critics argue that the risks posed by AI are not novel, and have existed for decades. But the nature of these systems, created by a handful of hugely powerful tech companies, requires a new approach beyond siloed regulation. Machines may make finance more efficient, but could do just as much to trigger the next crisis.

AI in Asset Markets

THE WALL STREET JOURNAL.

English Edition ▾ Print Edition | Video | Audio | Latest Headlines | More ▾

Latest World **Business** U.S. Politics Economy Tech Finance Opinion Arts & Culture Lifestyle Real Estate Personal Finance Health Style

BUSINESS

Zillow Quits Home-Flipping Business, Cites Inability to Forecast Prices

Termination of 'iBuying' comes after company said it was halting new home purchases for rest of 2021

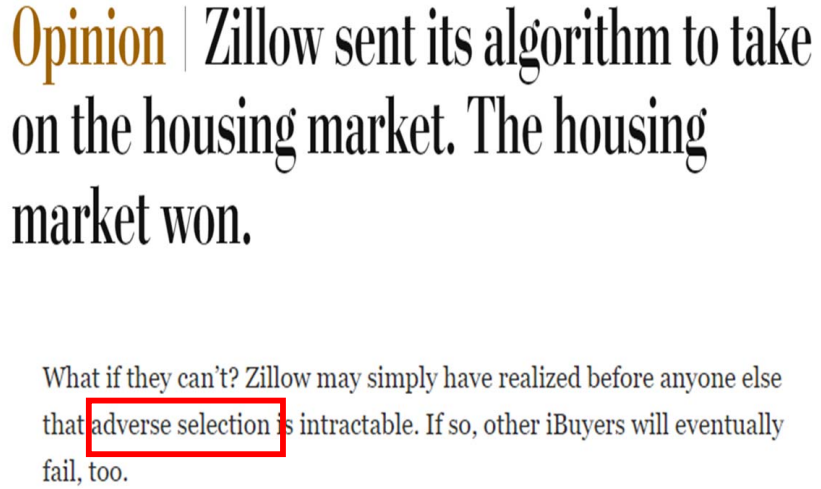

Opinion | Zillow sent its algorithm to take on the housing market. The housing market won.

What if they can't? Zillow may simply have realized before anyone else that adverse selection is intractable. If so, other iBuyers will eventually fail, too.

Behavioral Algorithmic Finance

- Can we still rely on traditional models of market making to **explain and predict** trading outcomes when prices are set by algos?
- How to model the behavior of algo market makers?

    1. **Bayesian learning + Nash equilibrium** (Glosten and Milgrom (1985), Kyle (1985) etc.). Algo market makers are just faster and more efficient in getting and processing information.

    2. **Reinforcement learning** (AlphaGo, Robotics, Autonomous Cars etc.); Algos learn how to make decisions iteratively by experimenting, receiving feedback and adjusting their behavior.

- Does the second approach predict outcomes different from the first? How and why?

What we do

- We consider **a standard market-making game ($\approx$ Glosten-Milgrom (1985))** but we assume that quotes are set by **Q-learning** algorithms (**"algo-MMs**) with no prior knowledge about the environment (e.g., intensity of adverse selection).
- We run **experiments** (a large number of interactions between algo-MMs and their clients, holding the environment constant) to study how Algo-MMs learn from experience and set their prices.
- We **benchmark** the observations to the predictions of the Nash equilibrium of the model (standard Bertrand equilibrium with zero expected profits for market makers).

Questions

1. **Adverse selection.** Can algo MMs learn to price "adverse selection"?
2. **Competition.** Can algos learn to be competitive (undercut when profitable to do so)? Major concern in online product markets. medskip
3. **Price discovery.** Can algos learn asset values ("discover fundamentals")? medskip
4. Today I just focus on 1 and 2.

Algorithmic Pricing and Competition

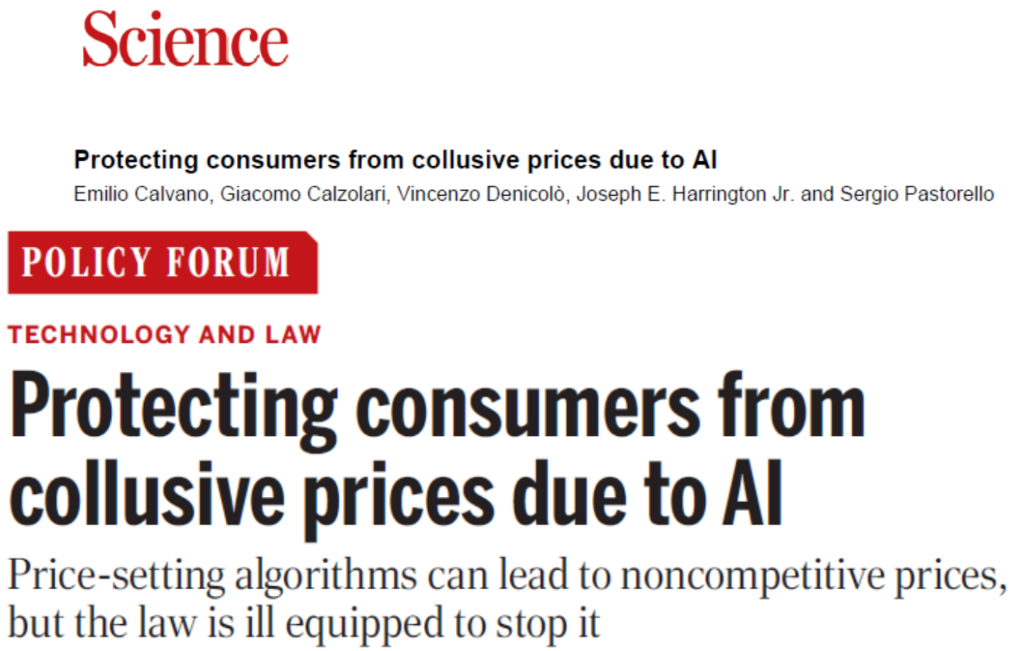

Science

Protecting consumers from collusive prices due to AI

Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, Joseph E. Harrington Jr. and Sergio Pastorello

POLICY FORUM

TECHNOLOGY AND LAW

Protecting consumers from collusive prices due to AI

Price-setting algorithms can lead to noncompetitive prices, but the law is ill equipped to stop it

- Hendershott et al. (2010, JoF): Algo trading is associated with an *increase* in "realized spreads" (rents):

  "This is surprising because we initially expected that if AT improved liquidity, the mechanism would be competition between liquidity providers."

Main New Findings

- **Algo-MMs learn not to be adversely selected:** They do not make losses **on average** and increase their spreads when adverse selection is greater.
- **Algo-MMs do not learn to undercut:** They settle on non competitive prices. They do not learn to collude but fail to learn to be competitive.
- **Algo-MMs' prices** are **more** competitive when adverse selection costs are **larger**.
- **Algo-MMs** learn to update their quotes based on the order flow.

Roadmap

The Market Making Game ("RFQ")

- **A risky asset** with payoff $\tilde{v}$ where $v = v_H = \mu + \frac{\Delta}{2}$ or $v = v_L = \mu - \frac{\Delta}{2}$ with equal probabilities.
- **A client** considers buying one share of the asset. Her valuation for the asset is:

$$\tilde{v}_C = \tilde{w}_C + \tilde{L},$$

  where (i) $\tilde{L} \sim \mathcal{N}(0, \sigma^2)$ ("liquidity shock") and (ii) $\tilde{w}_C = \tilde{v}$.
- 2 **Market Makers** (MMs) X and Y simultaneously post prices a_X and a_Y, not knowing $\tilde{v}$ and $\tilde{v}_C$.
- The client buys if $v_C \geq \min_{i \in \{X,Y\}} a_i$.
- If a trade occurs, the dealers posting the best quote earn (in aggregate) $(a_{min} - \tilde{v})$ and zero otherwise.

Likelihood of a Trade

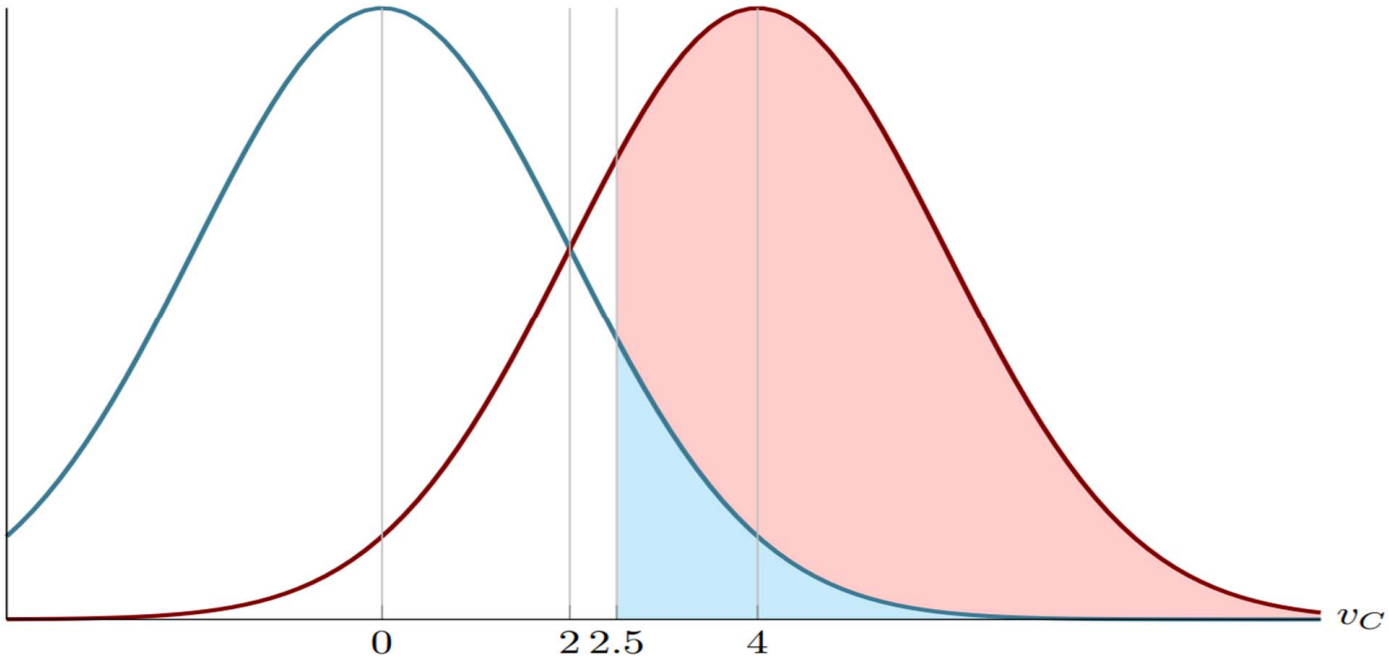

Example: v_H=4, v_L=0, E(v)=2. Best Offer=a_{min}=2.5.

- Likelihood of a trade if v= v_H = Blue + Red = **$D_H(a_{min})$**.
  - Likelihood of a trade if v= v_L = Blue= **$D_L(a_{min})$**.
  - Likelihood of a trade = Blue + 0.5 Red= **$D(a_{min})$**.
    - **D_H- D_L** = Red =Adverse Selection

Economic Environment

- **Two Cases:**

  1. **With adverse selection:** $\tilde{w}_C = \tilde{v} \Rightarrow$ Dealers are more likely to sell when the asset payoff is high than low.

  2. **Without adverse selection:** $\tilde{w}_C$ and $\tilde{v}$ have the same distribution but are independent $\Rightarrow$ The likelihood that a client buys is the same ("Blue + 0.5 Red") whether the asset payoff is high or low.

Dealers' Expected Profits

- $\bar{\Pi}(a_X, a_Y)$: Dealer X's expected profit if X posts a_X and Y posts a_Y.
- **Adverse Selection Case:** (suppose $a_X < a_Y$)

$$\bar{\Pi}(a_X, a_Y) = (\frac{1}{2} D_H (a_X - v_H) + \frac{1}{2} D_L (a_X - v_L)),$$

$$\Leftrightarrow$$

$$\bar{\Pi}(a_X, a_Y) = D(a_X) \left[\underbrace{(a_X - \mathrm{E}(\tilde{v}))}_{\text{Quoted Spread}} - \underbrace{\frac{\Delta \times (D_H - D_L)}{2}}_{\text{Adverse Selection Cost}} \right],$$

- **No Adverse Selection Case:** Adverse selection cost = Zero.

Benchmark: Nash Equilibrium

- **Unique Nash equilibrium**: a^* such that dealers make zero expected profits ($\overline{\Pi}(a^*, a^*) = 0$).

σ	0.5	1	3	5	7
(1) Quoted Spread	2.00	2.00	1.24	0.68	0.47
(2) Adverse Sel. Cost	2.00	2.00	1.24	0.68	0.47
(3) Realized Spread (1)-(2)	0	0	0	0	0

Δv	0	2	4	6	8
Quoted Spread	0	0.16	0.68	1.65	3.02
Adverse Sel. Cost	0	0.16	0.68	1.65	3.02
Realized spread	0	0	0	0	0

Nash Equilibrium: Predictions

1. **H.1.** Dealers' **average quoted spread** $(a^* - \mathrm{E}(\tilde{v}))$ is strictly **larger with adverse selection.**
2. **H.2.** Dealers' **average quoted spread** declines with σ and increases with Δ if adverse selection. Without adverse selection, these parameters have no effect on the quoted spread.
3. **H.3.** Dealers' **average realized spreads** is **zero** with and without adverse selection.

- Are these (very standard) predictions satisfied when prices are set by Algo MMs?

Roadmap

Learning how to Play

- The Nash approach assumes that the market makers (MMs) know $\bar{\Pi}(a_i, a_{-i})$.
- This requires a lot of knowledge on the environment: (i) The distribution of client's valuation (v_C), (ii) the distribution of the asset payoff, (iii) the number of competitors etc.
- **An alternative behavioral model: Reinforcement learning**
- **Algo MMs (AMs):** They have no prior knowledge about the environment and learn to play the market making game via experimentation.
- We focus on **Q-learning**. Not because it is realistic but because it is a simple model of behavior for an algorithm. Our goal is not find the best possible algorithms in our environment.

Q-Learning Algorithm - Description

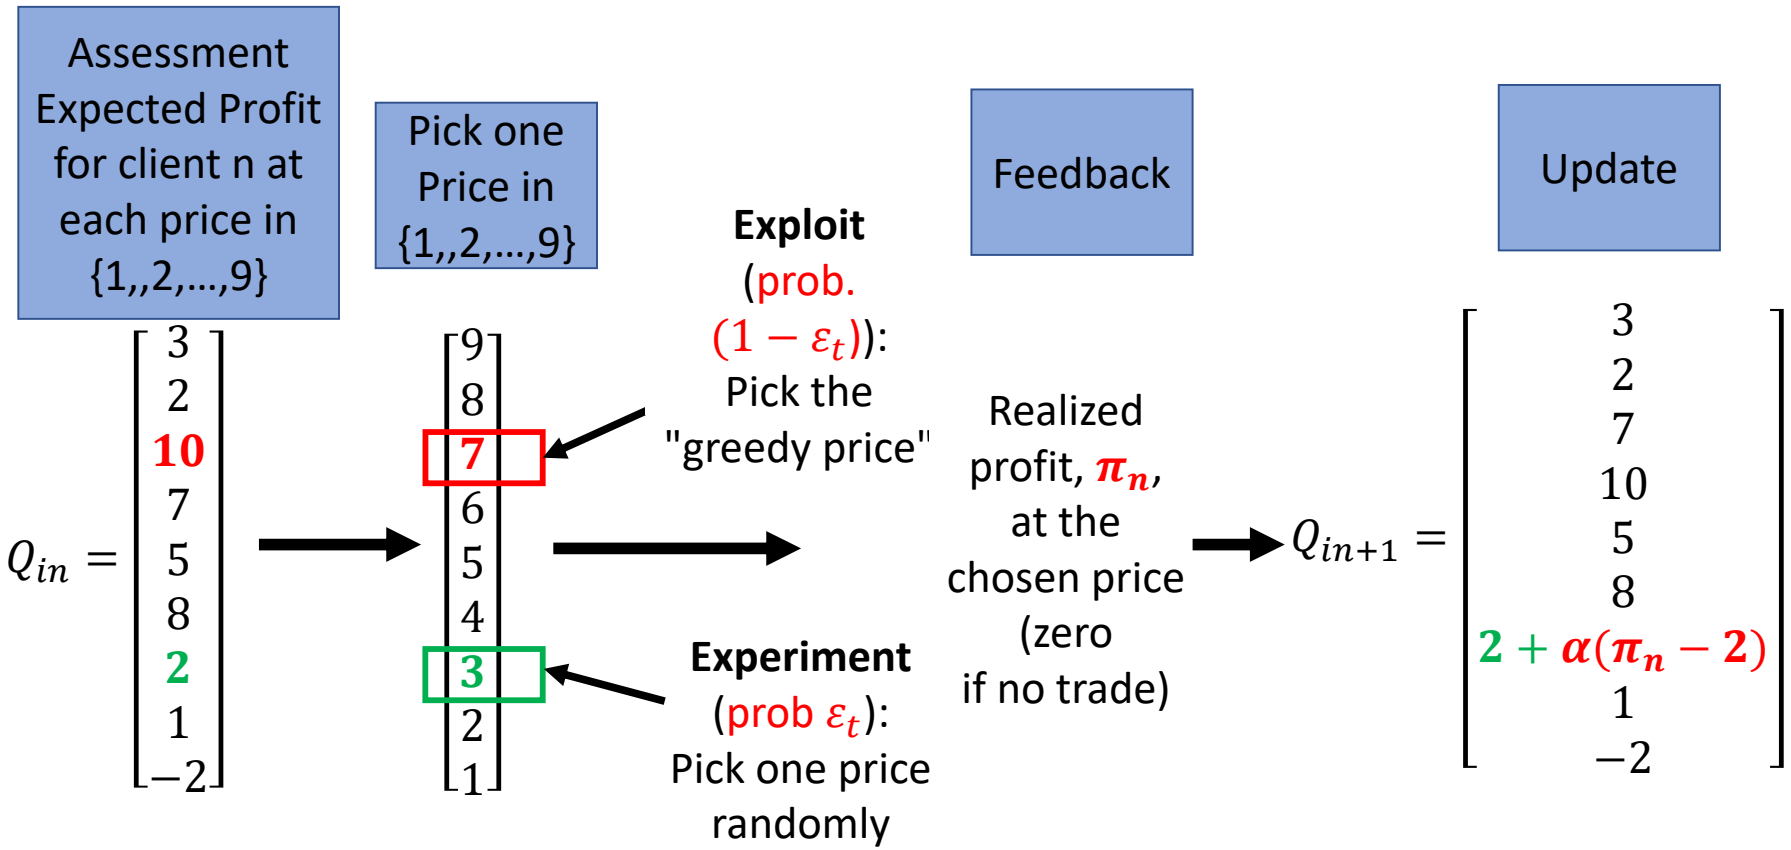

Roadmap

Implementation

- **Algorithm:** $\alpha = 0.01$, $\epsilon_t = exp(-\beta t)$ with $\beta = 8.10^{-5}$, random Q_0.
- **Baseline Environment:** $E(v) = 2$, $\Delta = 4$, $\sigma = 5$ (baseline).
- For each parametrization of the game (σ,Δ), we run $K = 10,000$ "experiments", each with $T = 10^6$ clients.
- **Price Grid:** 139, tick size $\delta = 0.1$: $\mathcal{P} = \{1.01, 1.02, ..., 2, ..., 14.9\}$.
- **We focus on the distribution of prices (mostly mean prices) across experiments** after the last client.

Learning

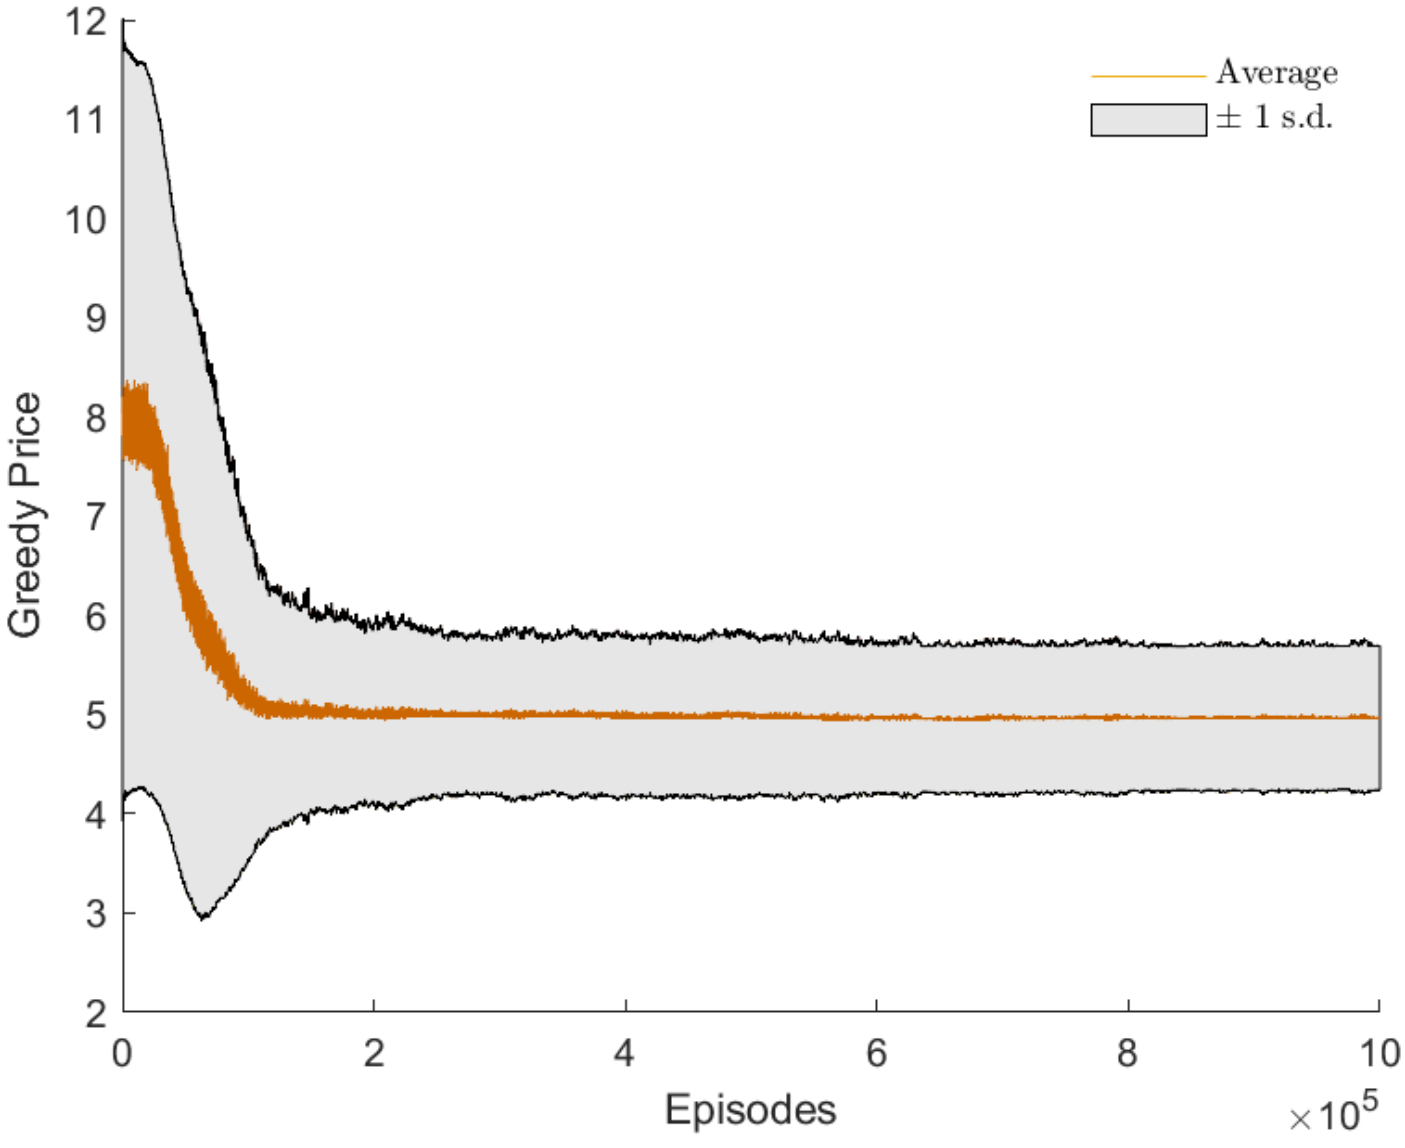
Average
± 1 s.d.
Greedy Price
12
11
10
9
8
7
6
5
4
3
2
0
2
4
6
8
10
Episodes
×10^5

Observation 1: AMs' prices are not competitive

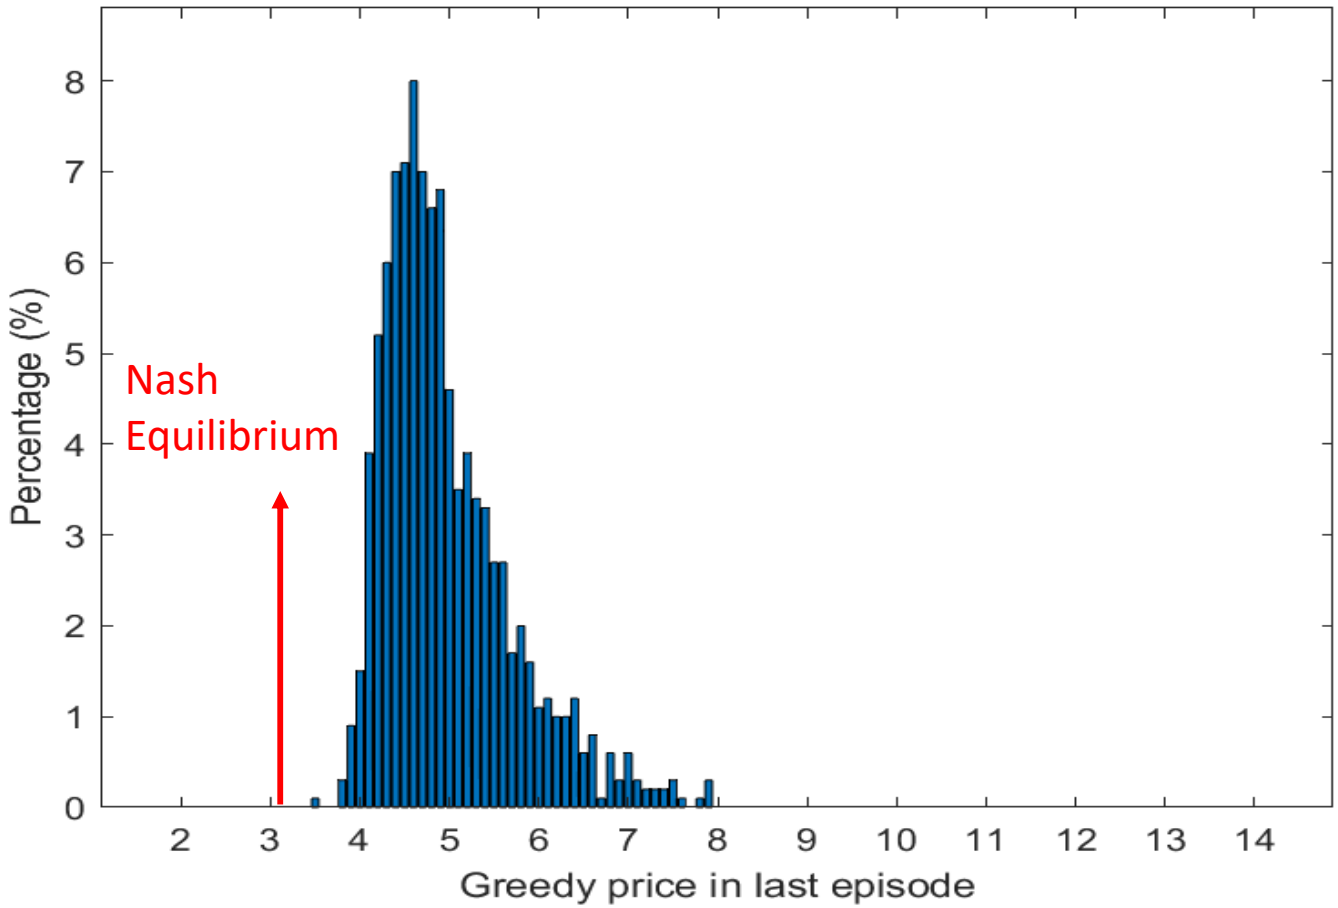

Observation 2: AMs' spreads account for adverse selection

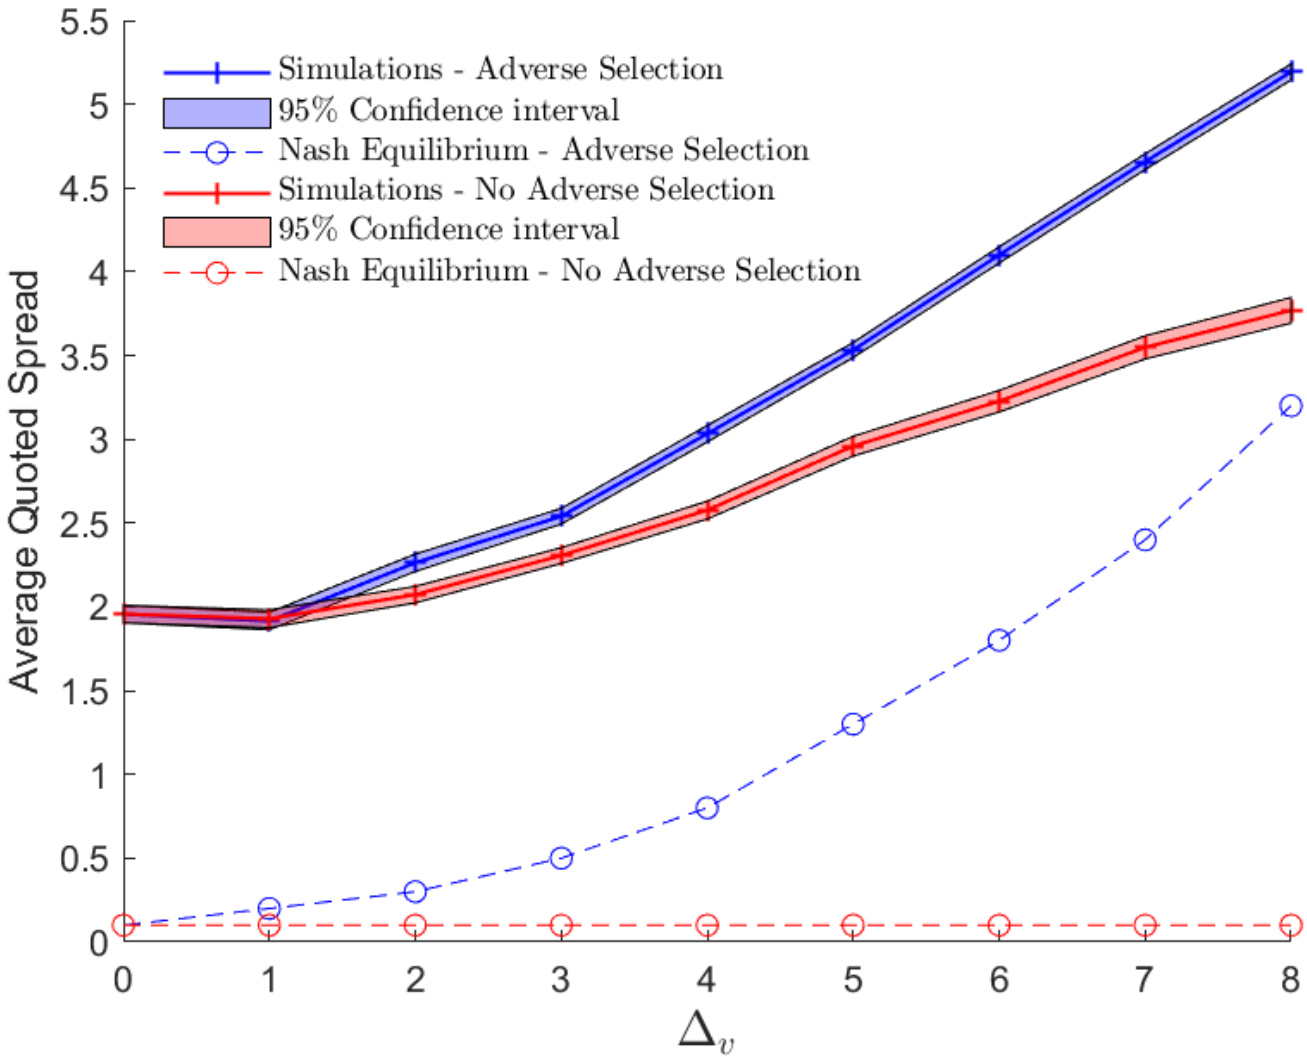

Observation 3.a: Adverse selection makes AMs more competitive

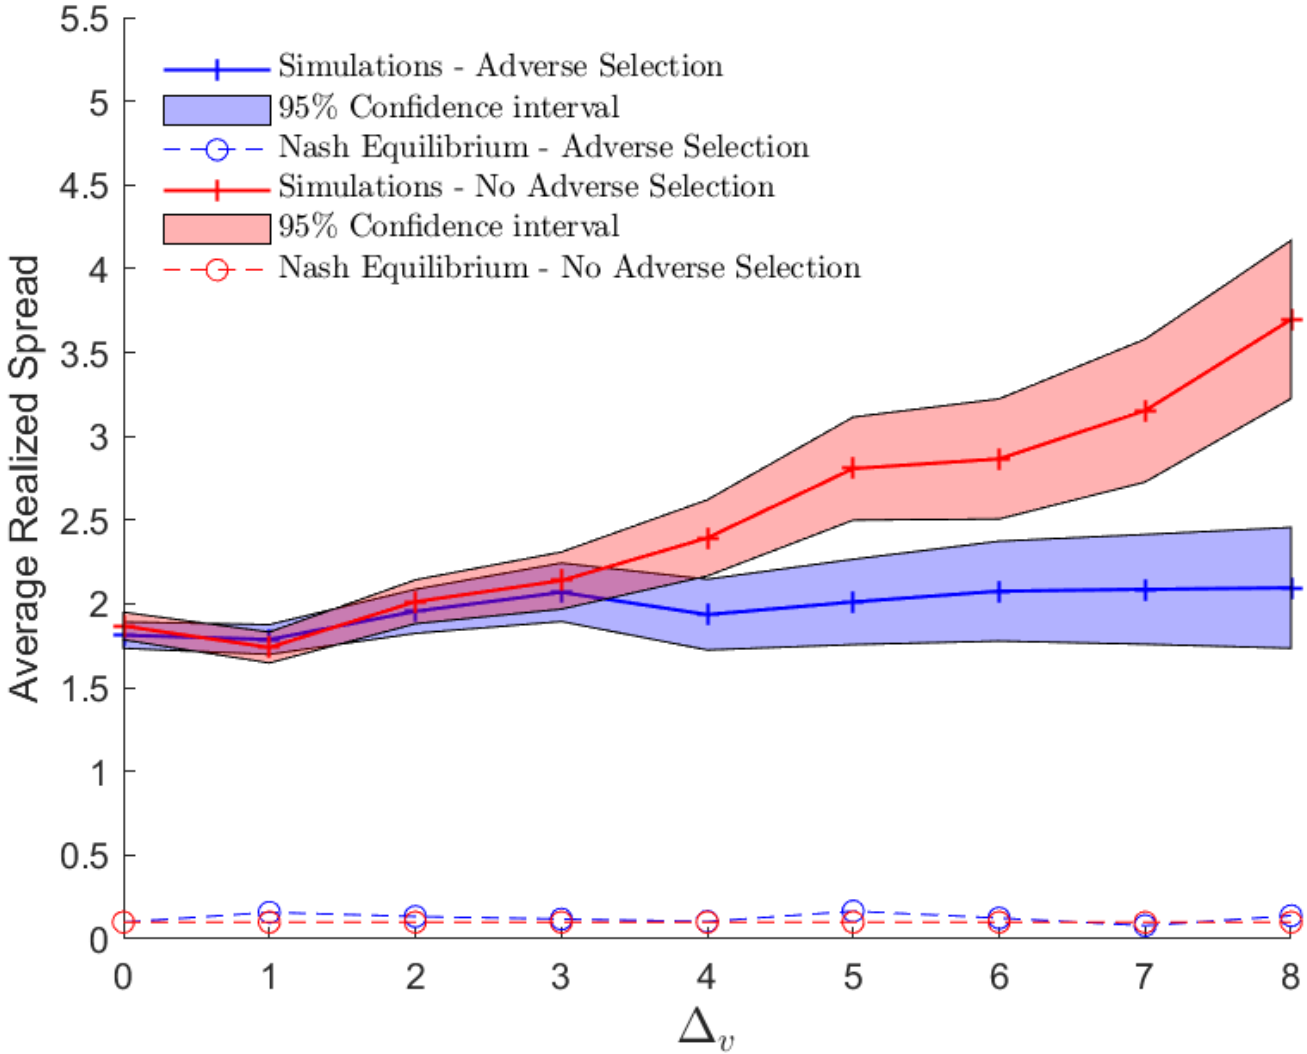

Observation 3.b: Adverse selection makes AMs more competitive

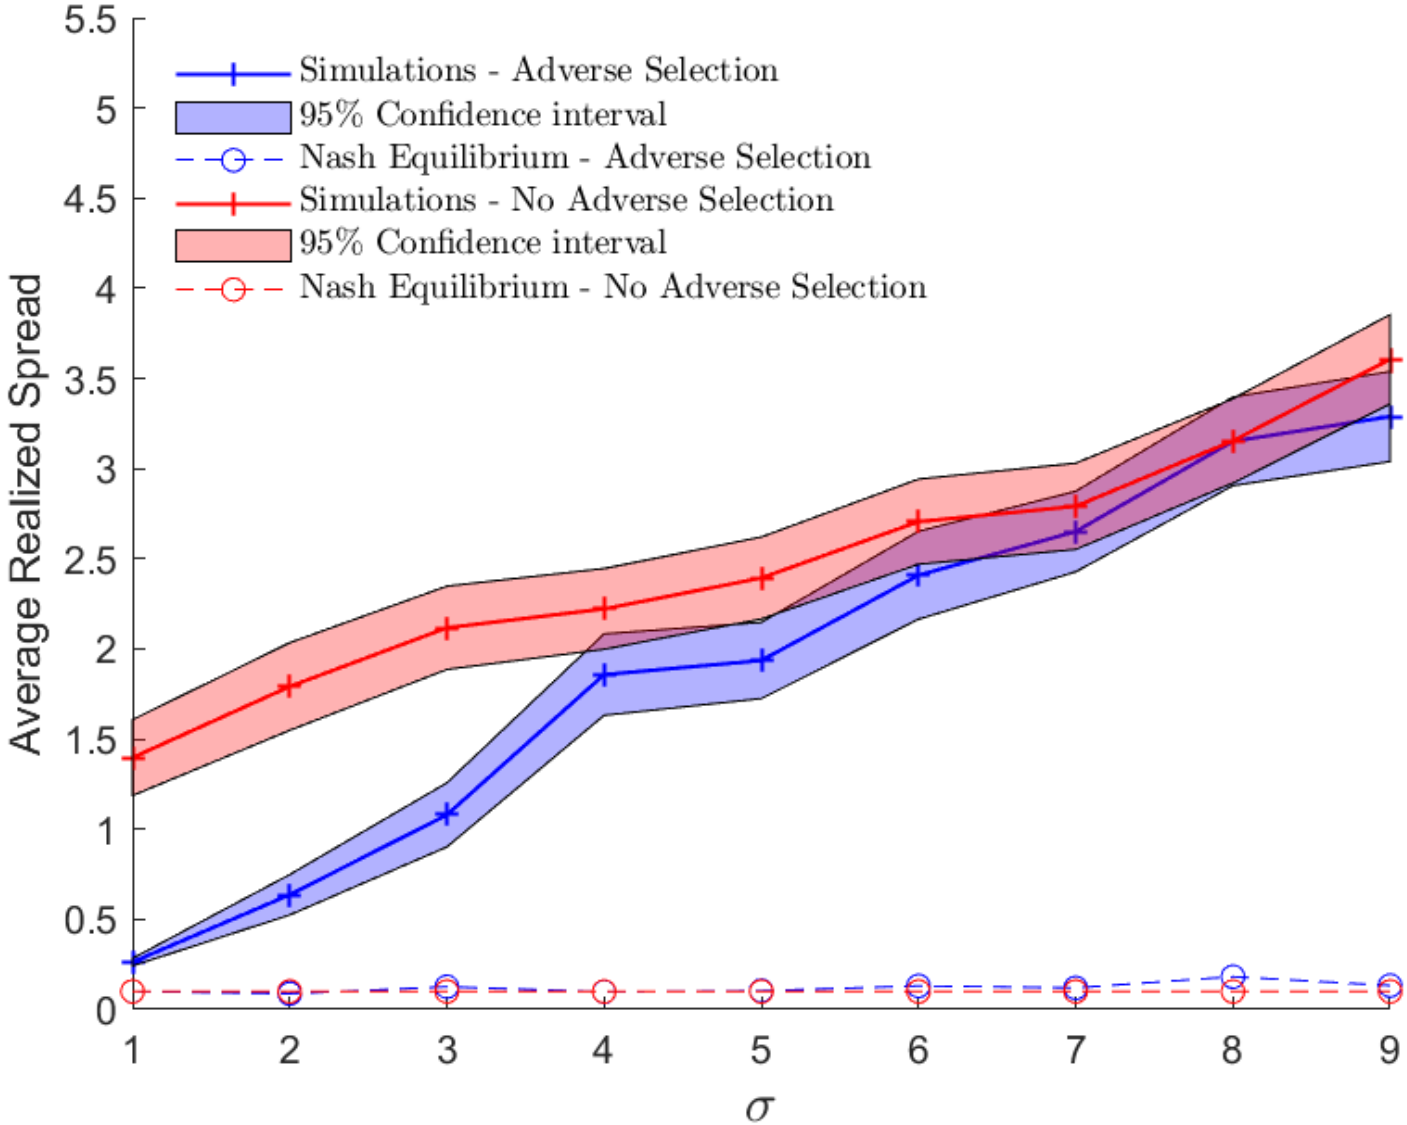

Observation 4: AMs' spreads and rents decrease with the number of AMs

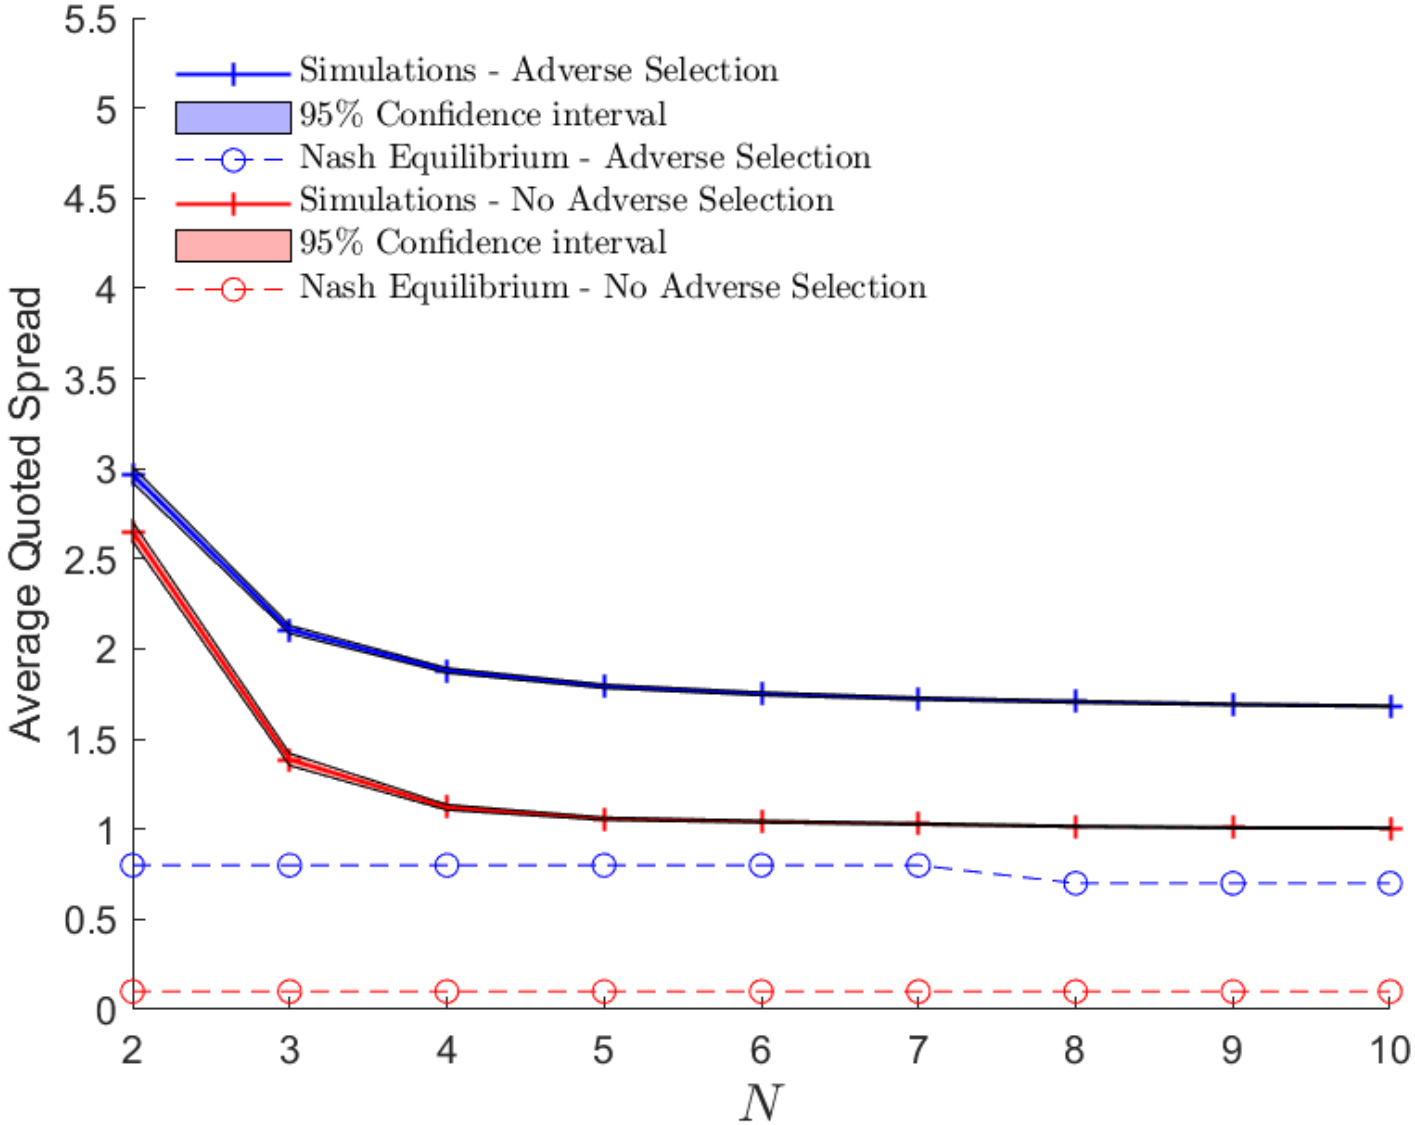

Evidence

Brogaard and Garriott (2019, JFQA):

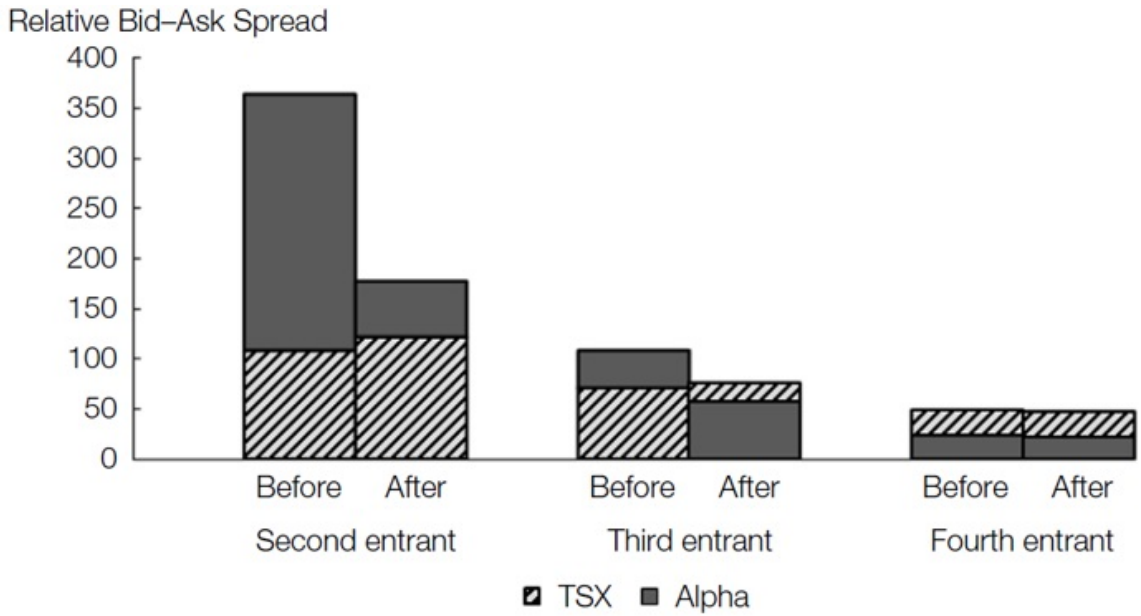

Puzzles

- **H.1 is satisfied: AMs charge larger spreads when there is adverse selection.**

- **Otherwise, AMs' prices are significantly different from the Glosten-Milgrom prices:**

  1. **Realized spreads are not zero.** Algo MMs are (suboptimally) non competitive: They do not learn to undercut non competitive quotes when it is profitable to do so.

  2. **Realized spreads are larger when there is no adverse selection.** In principle, they should be zero with or without adverse selection.

  3. **Quoted spreads and realized spreads are smaller when adverse selection costs are larger**. Exactly the opposite of what standard theory predicts.

An Explanation: Lack of Experimentation + Noise

- **Parameters:** $\sigma = 5$, $\Delta = 4$, $E(v) = 2$. If the AMs post 5 (the modal price in the experiments), they each obtain an average profit of 0.3. If Y undercuts at 4.9, it obtains:

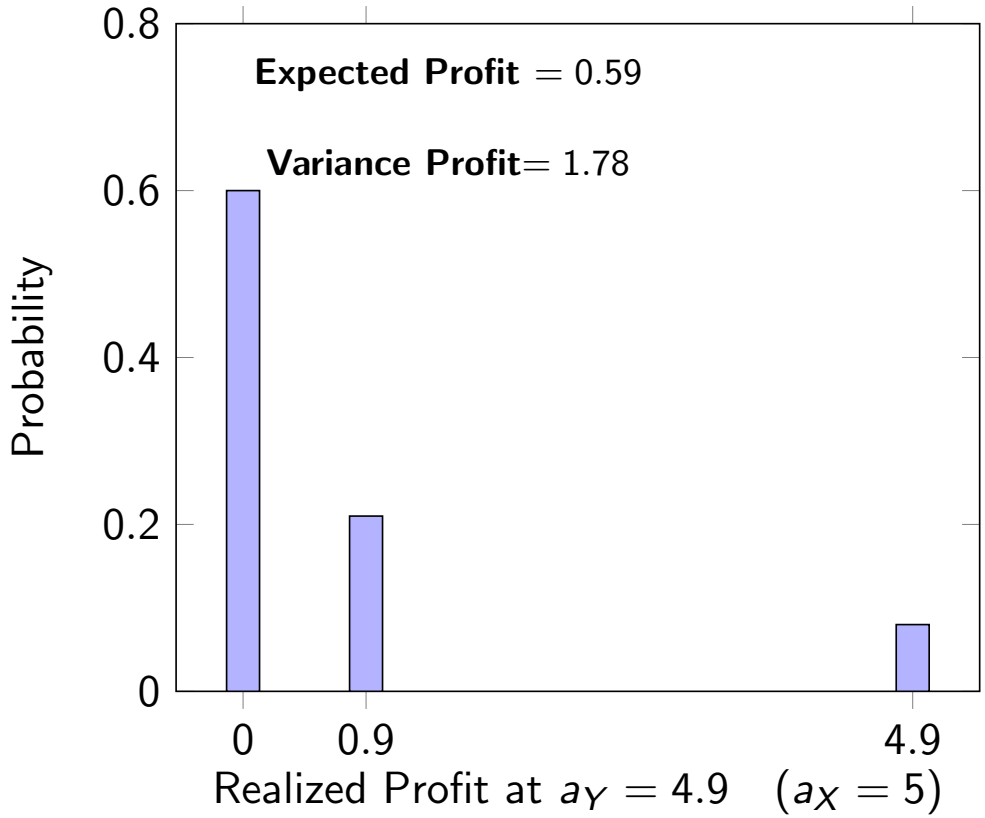

- $\Rightarrow$ **The Q value of undercutting is a noisy estimate of the true expected profit**.

Adverse Selection Reduces Noise

- Same parameters as before but Blue: Adverse selection and Red: No adverse selection.

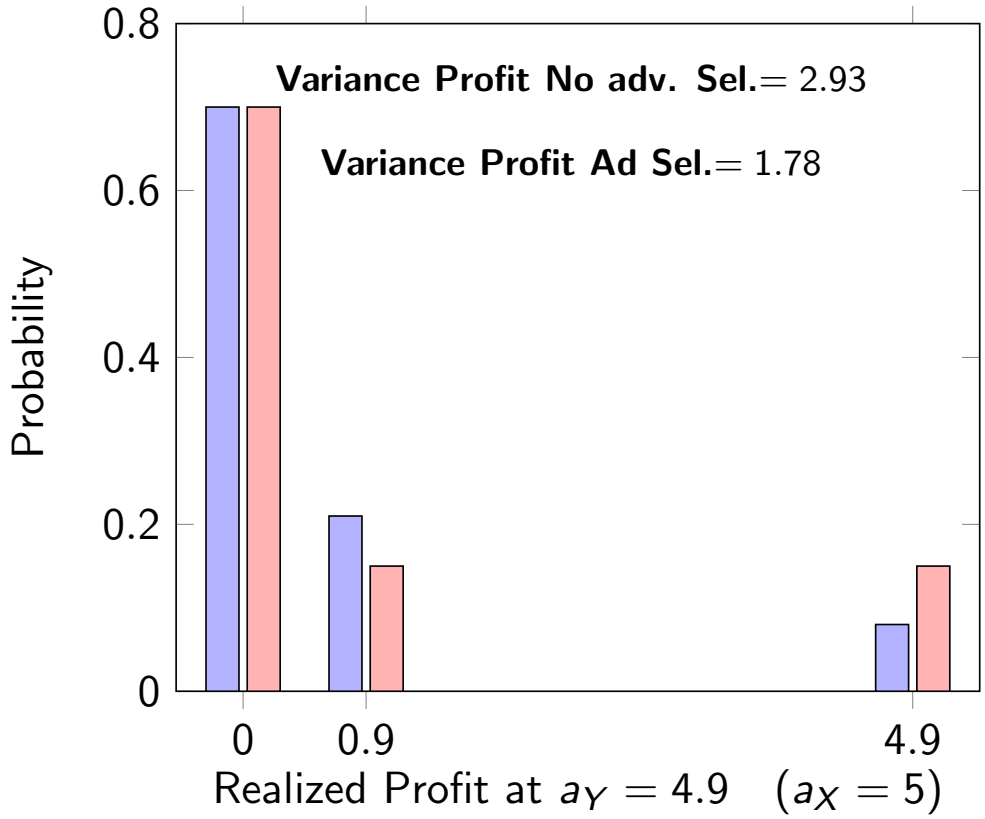

- **Adverse selection reduces the noise in the feedback received by an AM if it undercuts $\Rightarrow$ Makes learning to undercut easier.**

Adverse Selection Reduces Noise

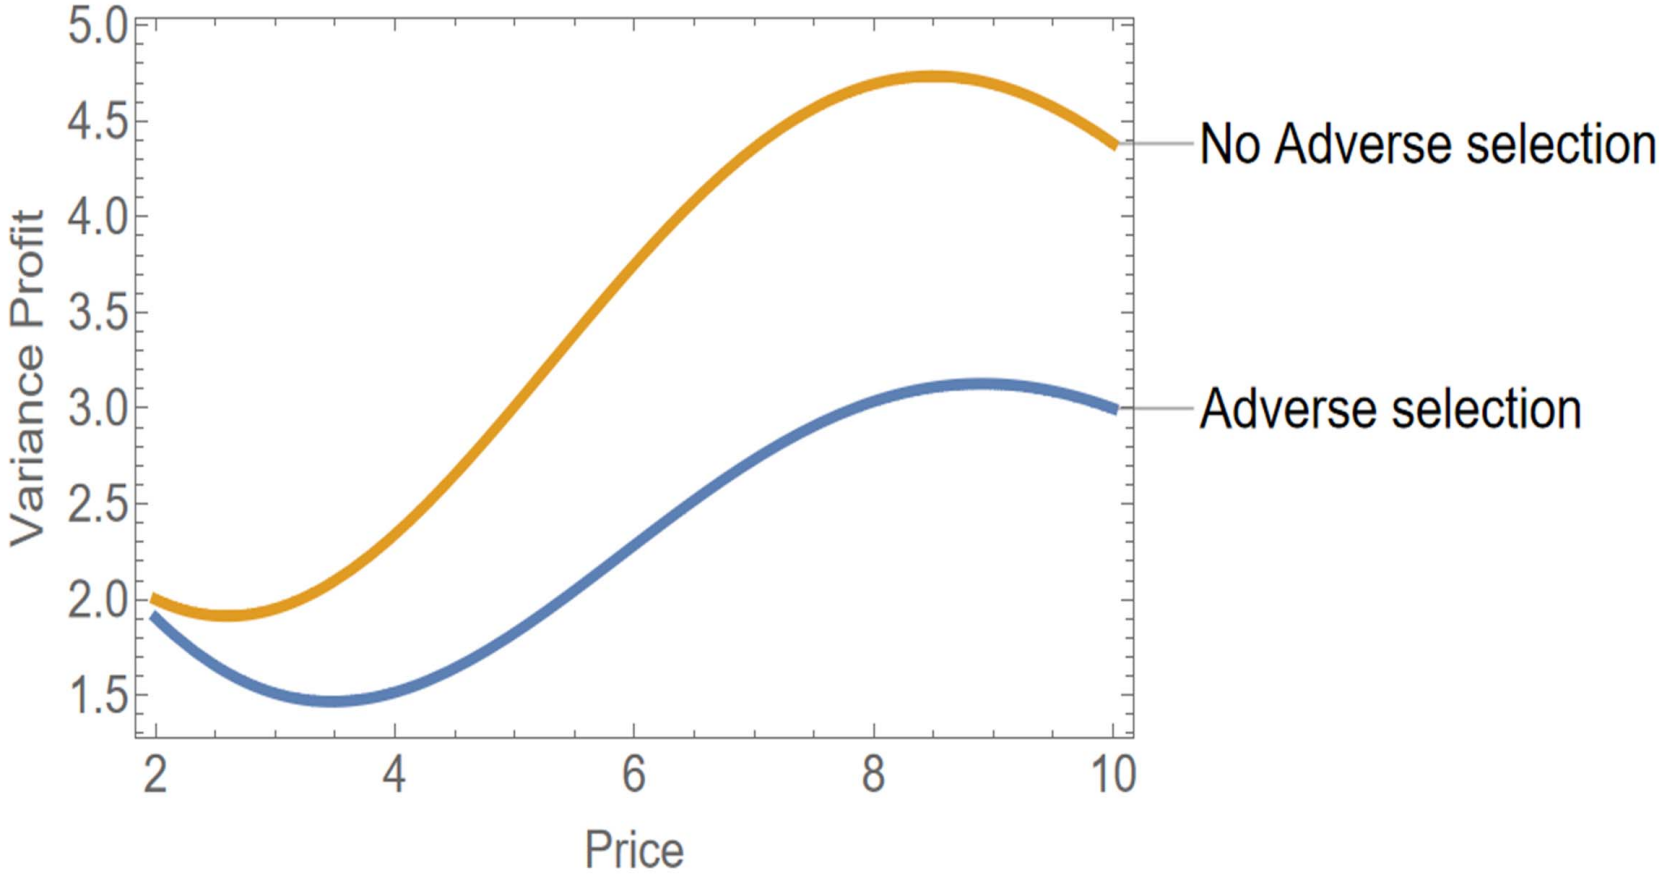

The Dispersion In Clients' Liquidity Shock Increases Noise

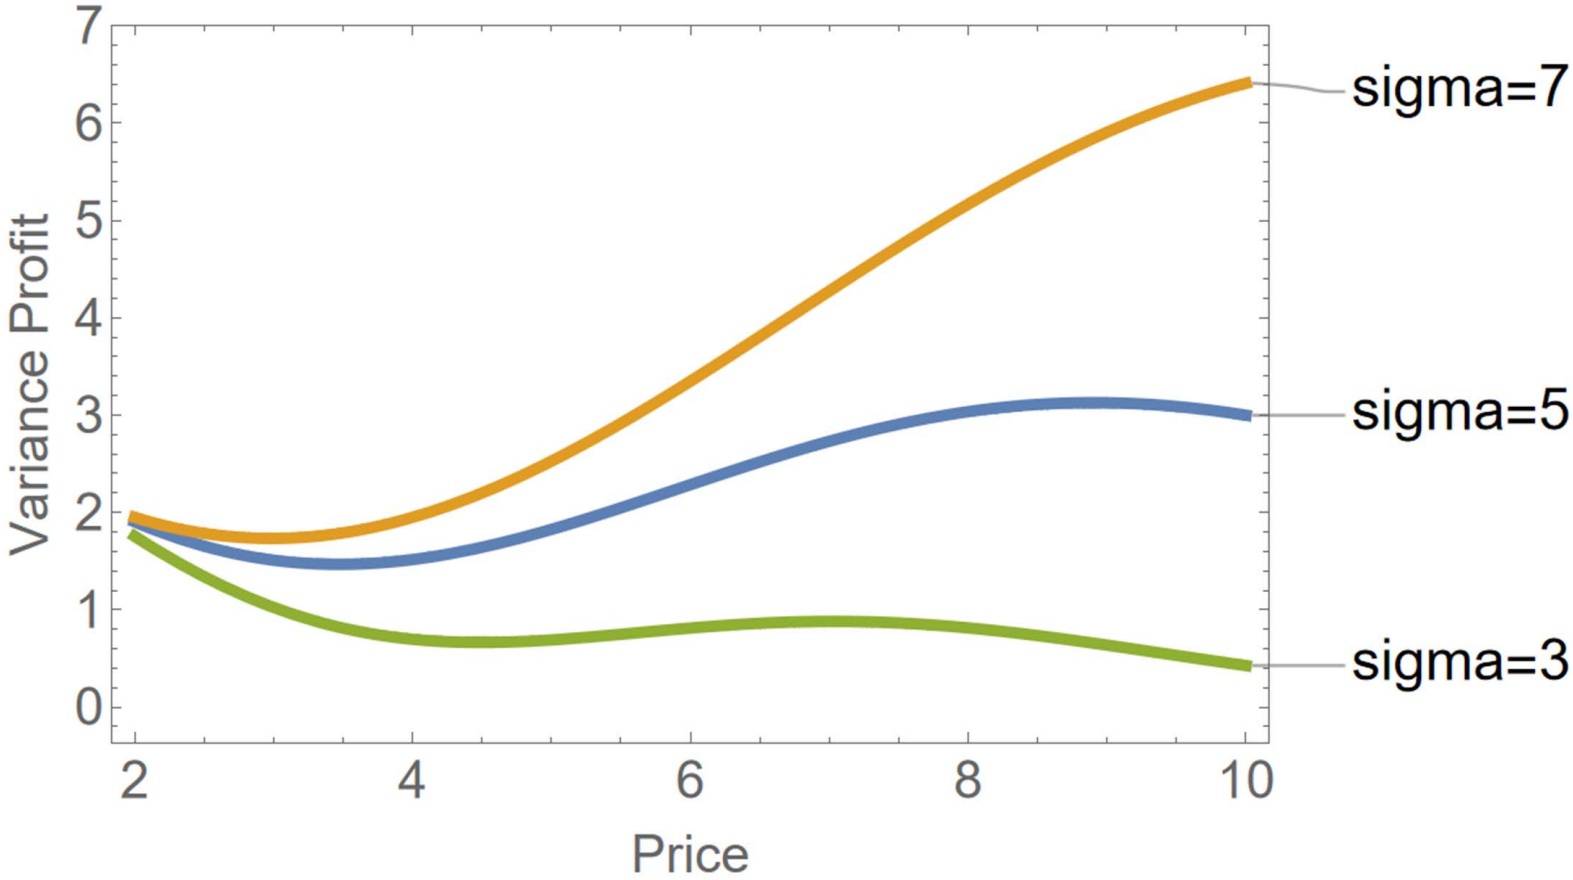

Roadmap

More Experimentation, More competitive outcomes

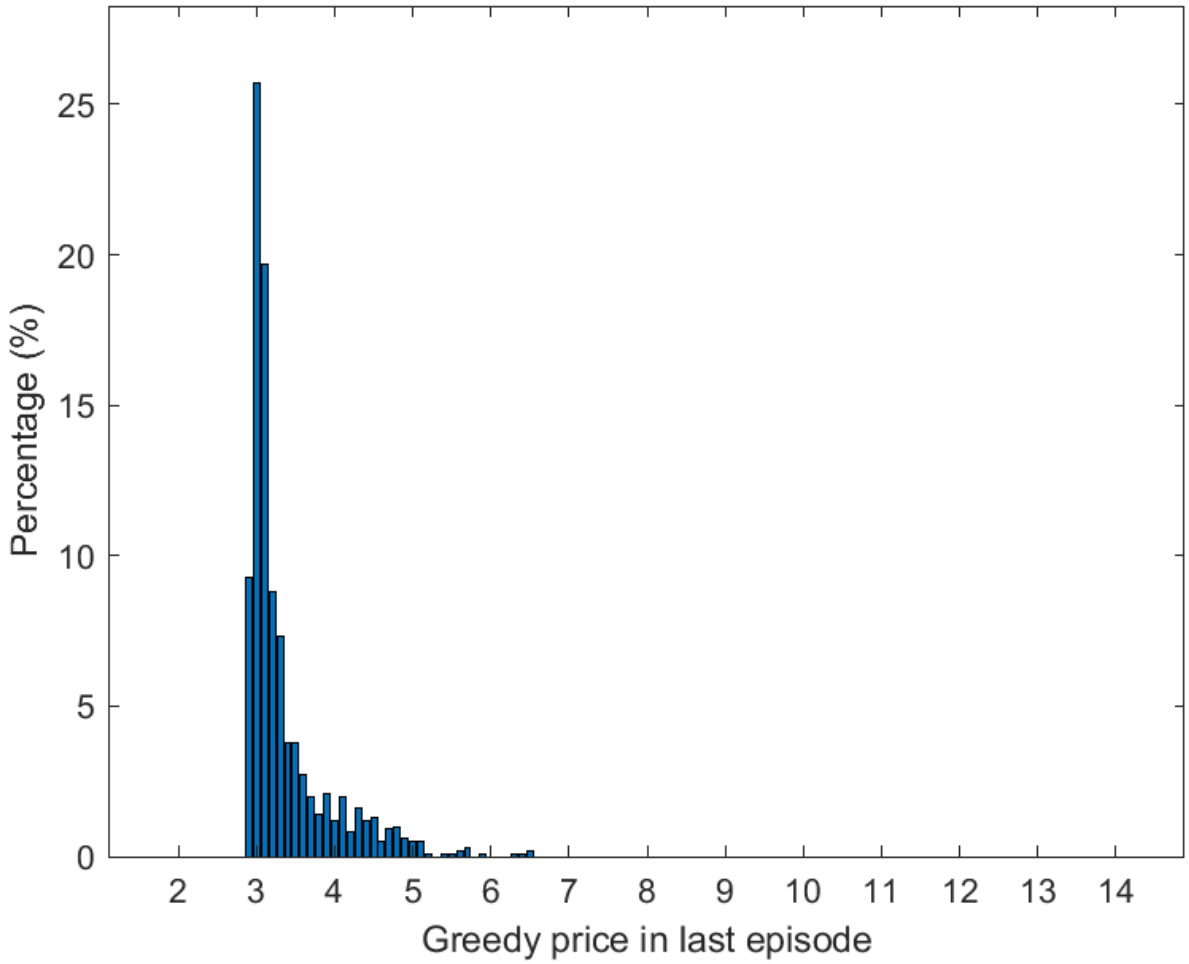

Figure: Parameters: $\sigma = 5$, $\Delta = 4$. The experimentation rate never goes below 5%.

But...experimentation is costly

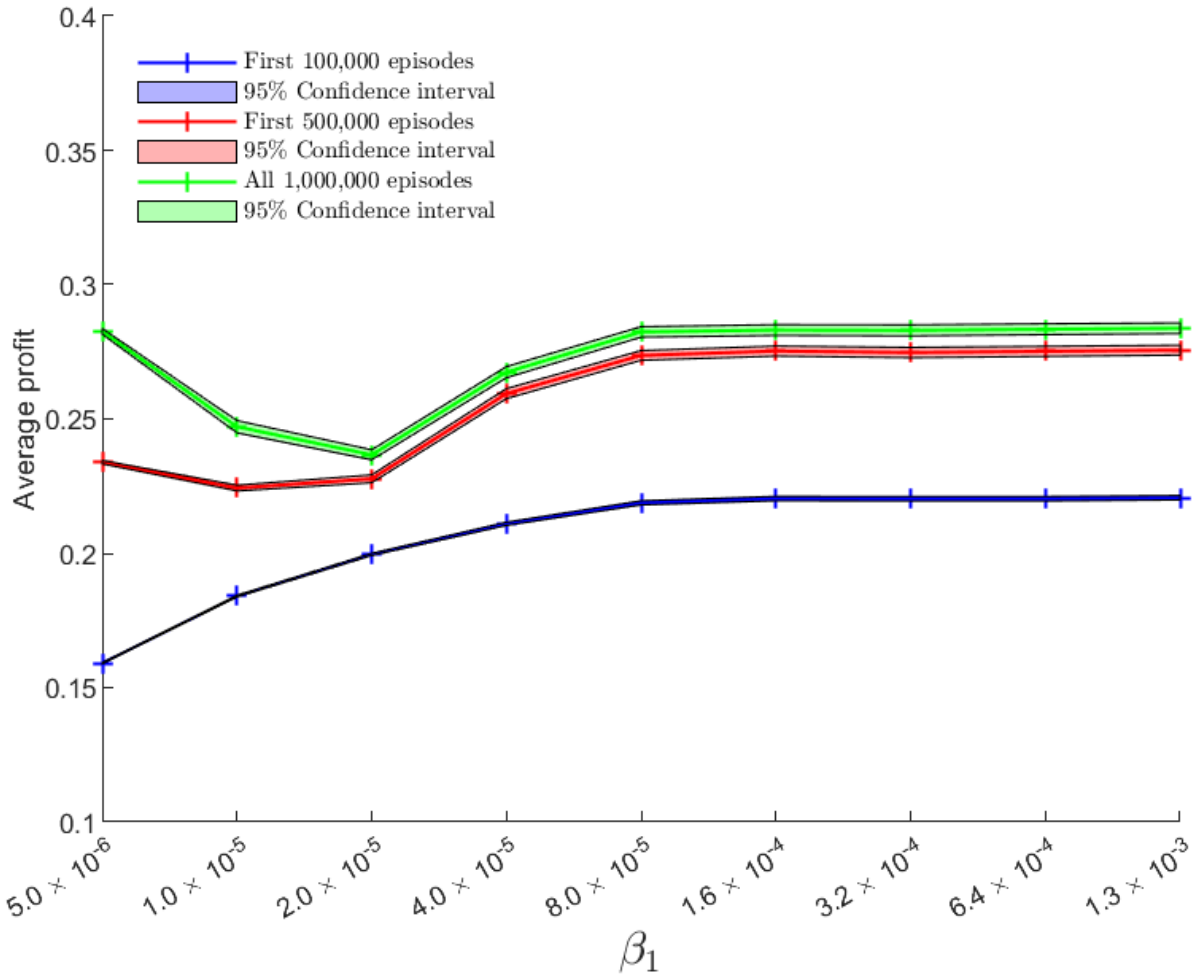

Figure: We vary β and compute average per period profits over various time windows.

Roadmap

Conclusion

- **It should not be taken as granted that AI is going to make financial markets more efficient/competitive.**
  1. Our algo market makers do not reach competitive prices $\rightarrow$ lead to inefficient outcomes (e.g., fewer mutually beneficial trades, less accurate prices).
- **Standard economic intuitions do not apply to interpret trading outcomes:** We find that adverse selection (a friction) makes prices more competitive.
  1. $\rightarrow$ Other things equal, algo might get larger rents in markets with less asymmetric information.
- **Open question:** Is trading really suitable for applying reinforcement learning (our environment is stationary and yet learning is imperfect after many iterations)?

Thank You!