AI and Competition in Agentic Markets

Giacomo Calzolari

European University Institute

13th Postal Economics Conference
E-commerce, Digital Economics & Delivery Services
Toulouse, April 16–17, 2026

Why AI matters for markets

- AI already transforming a wide range of economically relevant activities (diagnostics, drug discovery, language and translation, autonomous navigation, fraud detection, recommendation systems, dynamic pricing, inventory, ...)
- Not a surprise, as AI can be seen (at least) as **cheap and more precise prediction**

- **Focus of this talk:** AI as improved prediction, and its implications for some specific market applications. Not trivial because,
  - using predictions combined with human ***Judgment*** (preferences)
  - often on unknowns, given the novelty of AI,
  - how to pass our Judgment to AI (e.g., for automation)?
  - what are the "emergent properties " of AI?
  - ...

AI and market sides

Both traditional sides of the market, demand and supply, are adopting AI

- **Supply side:** Sellers increasingly delegate **pricing**, inventory, product design, and other strategic decisions to algorithms, some AI-powered.
  - In particular, algorithmic (re)pricing for:
    - dynamic pricing
    - personalized pricing

- **Demand side:** Each individual consumer progressively supported/replaced, in **search** and/or **purchasing decisions** with autonomous agents
  - Delegating individual search to Recommendation Systems
  - Agentic buyers for automated transactions (choosing and executing):
    - Amazon "Buy for Me", Walmart in-app autonomous agents
    - Visa, Mastercard, PayPal – "Agent Pay"
    - Open AI's "Operator" and Manus (by Butterfly effect, Singapore)

Agentic Ongoing Transformation

Towards the delegation of market decisions to software agents:

- From **assisted decision-making** to **autonomous execution**
- These trends may lead to **agentic markets**, where both sides of the market delegate decisions to AI agents
- **Dual delegation** shifts unit of analysis from aggregate consumer demand to **personalized bilateral interactions** between algo agents

Key transformation:

Mass markets become markets of one

- This moves us from rare cases of buyer-specific power to **personalized mass markets**
- **New questions** arise for efficiency, competition, welfare, market design, and regulation

Dissecting the mass market personalization

1. Agentic sellers with algo pricing
2. Agentic search for buyers
3. Combining algo sellers with algo buyers

Agentic Sellers: Evidence of Re-pricers

- Algo pricing is common and increasing
  - 2016: >33% Amazon top sellers used repricers (Chen et al. 2016).
  - Bol.com 2018: from 45% (1 seller) to 90% (10 sellers) (Wieting 2021).
  - NCAs surveys 2018–: prevalence 10–60%.
  - Uber, Zalando, Alibaba use ML repricers; Amazon's "Project Nessie"
  - A new (fuzzy) repricing industry (Calzolari and Hanspach 2024)

- With significant effects
  - German gasoline: +28% margins if rivals adopt (Assad et al. 2024).
  - US online drugs: adoption speeds → Stackelberg (Brown et al. 2023).
  - US rental housing: RealPage algos → +3% prices, lower occupancy (Calder-Wang et al. 2023).
  - Frequent cycling prices (Mussolf et al. 2024)

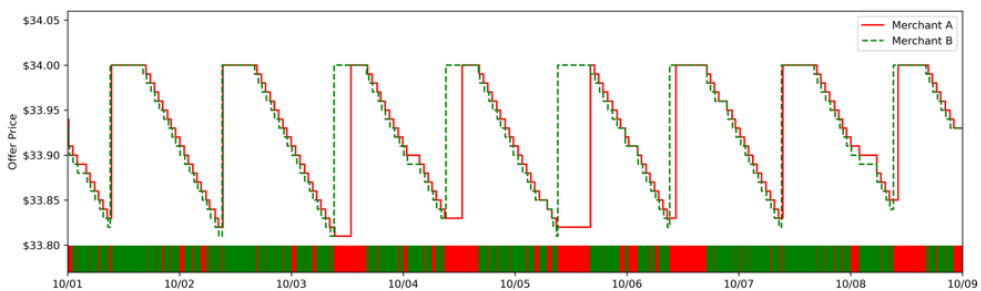

Classification of Pricing Algorithms

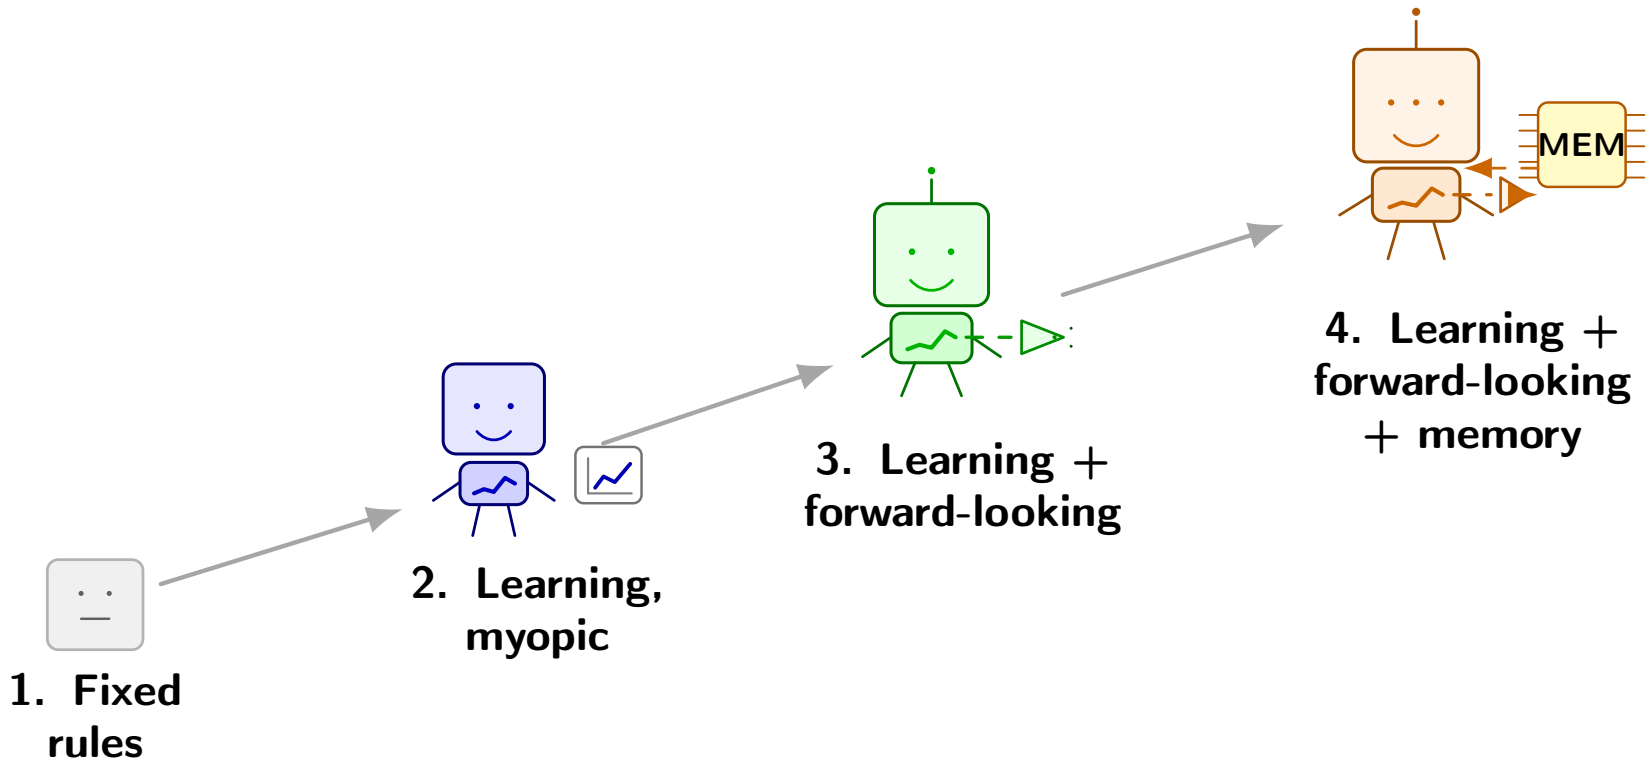

Some pricing algorithms, even simple ones, can sustain high prices.

From classification to (synthetic) evidence

Learning with Memory: charging high prices

Calvano Calzolari Denicolò and Pastorello (CCDP 2020) study reinforcement-learning algorithms of the **fourth type** (**learning, forward-looking, and with memory**). When interacting and learning with each-others, these algos systematically learn to sustain high prices.

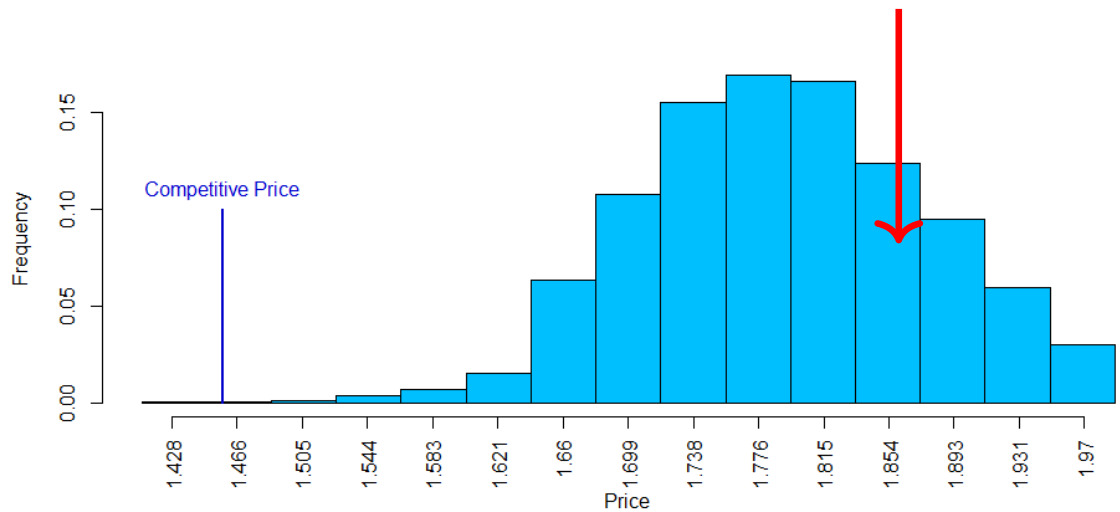

The algorithms set high prices, close to the monopolist's optimal price (red arrow). But why? **What exactly have they learned?**

Can We Look Inside AI Algorithms?

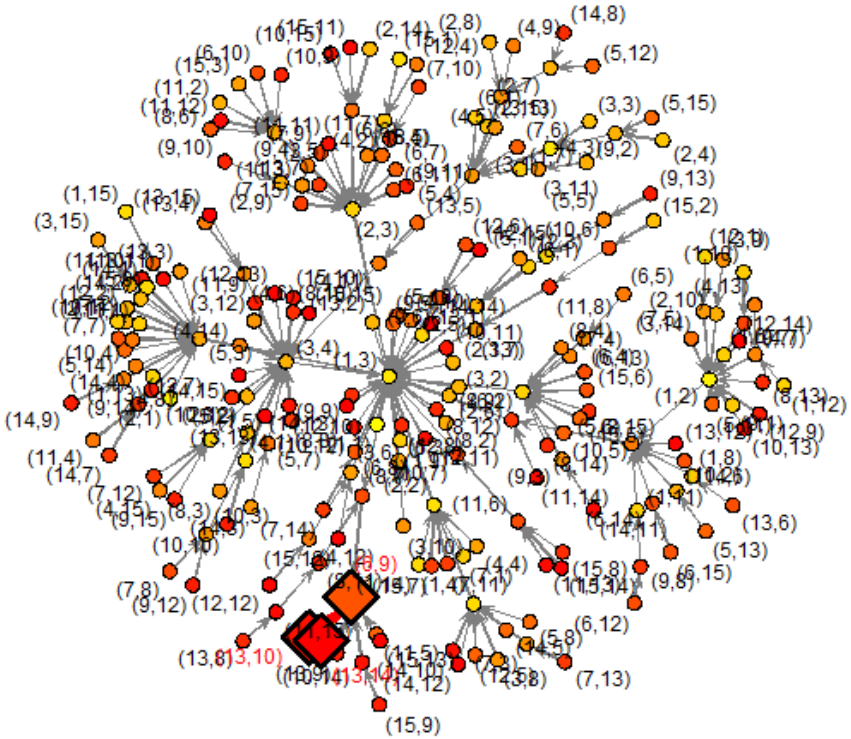

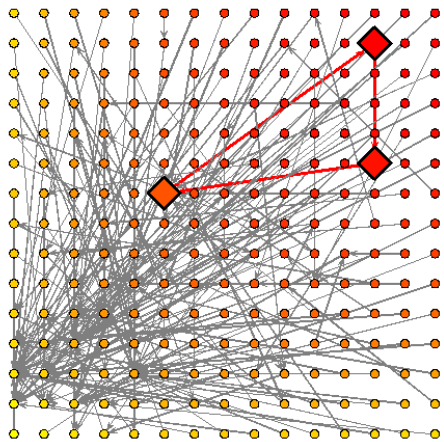

Raw inspection is **not very informative**. This is the well known issue of **AI interpretability problem** (**XAI**). What can we do with these **black boxes**?

Thinking as an economist

CCDP (2020): Provoking the algos, they reveal that they learn to autonomously collude with reward–punish–forgiving strategies. (Synthetic market analysis.)

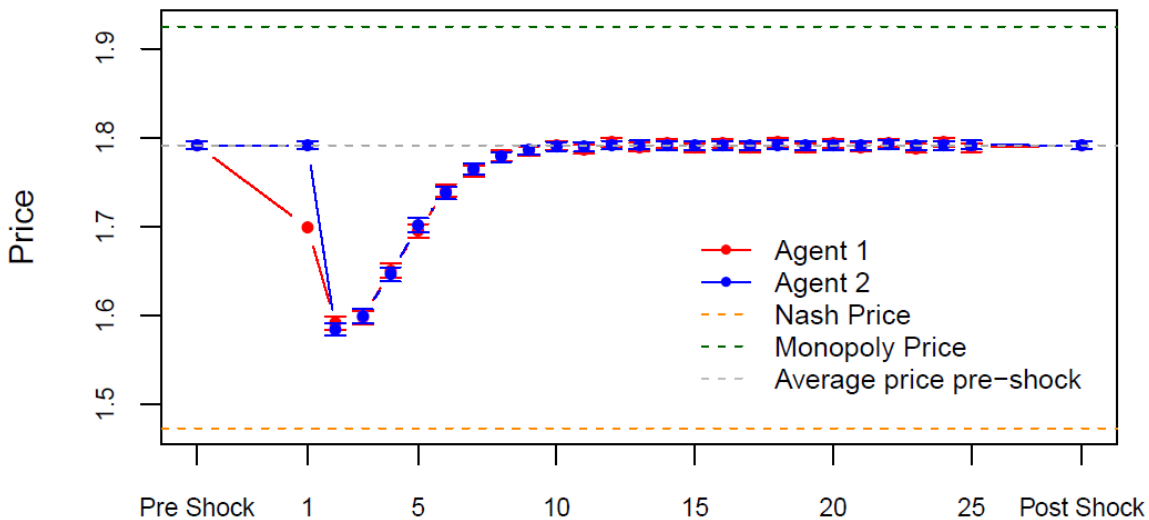

Price reactions to deviation

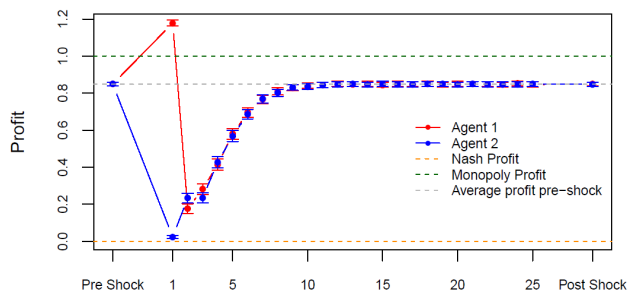

Profit (normalized) paths

Robustness of autonomous collusion

- Varying market structure
- Exploiting multi-market contacts
- Role of data advantage
- Counterfactual reasoning
- Learning frictions and imperfect monitoring
- Specific applications
- Increasing sophistication (DeepQL, ChatGPT, . . .)
- Theory and generalisation
- Commitment and inevitability of collusion
- Also, some healthy skepticism

Robustness of autonomous collusion

- Varying market structure
- Exploiting multi-market contacts
- Role of data advantage
- Counterfactual reasoning
- Learning frictions and imperfect monitoring
- Specific applications
- Increasing sophistication (DeepQL, ChatGPT, . . .)
- Theory and generalisation
- Commitment and inevitability of collusion
- Also, some healthy skepticism

Spontaneous communication?

Firm 1: "*We both maximize profits by exploring different prices while maintaining a reasonable price difference.*"
Firm 2: "*Agreed.*" (Wu et al. 24)

Now introducing Agentic Buyers

- Existing literature focuses mainly on algorithmic sellers
- What about buyer-side agency?
- **Agentic buyers**: autonomous or semi-autonomous systems optimising *purchasing process*, or parts of it.

- Three developments:
  1. Consumers delegating search to recommendation systems (RecSys).
  2. Whether agentic buyers can counteract algorithmic collusion.
  3. Bilateral "agentic markets" where both sides are algorithmic agents.

Recommendation Systems as Delegated Search

- Digital platforms use RecSys to personalise product prominence via ranking and recommendations.
- RecSys help buyers navigate vast choice sets (Amazon: >350M items; Spotify: >90M songs, YouTube: >26Bn).
- Widespread reliance: ∼75% of Netflix views; ∼60% of YouTube views.

- Delegating search to RecSys shapes *effective demand* by influencing what buyers see and choose: from human demand to *algorithmic demand*.

Mechanics and Feedback Effects

- RecSys AI algos predicts users' preferences based on historical interactions (other users and items): clicks, purchases, ratings, browsing time.
- Given preference estimates, RecSys recommends items in order to individual buyer: *personalized prominence*.
- An interesting feedback loop:
  buyer choices (data) $\rightarrow$ model updates $\rightarrow$ future recommendations.
- Do RecSys amplify biases, induce path dependence, and create feedback loops affecting market structure?

Effects of Algorithmic Demand: Results

- CCDP (2025): Synthetic market analysis where buyers with correlated preferences search or use collaborative filtering RecSys (truthful).
  1. Market concentration ↑, homogeneity of choices ↑.
  2. Prices rise 7–15% as sellers adapt to algorithmic demand.
  3. Prices ↑ outweigh improved match and saved search costs.
  4. Still, *individual* incentive to adopt and delegate Recsys (dilemma).

Effects of Algorithmic Demand: Mechanism

- RecSys increases willingness to pay
- and compresses heterogeneity $\rightarrow$ towards rectangular demand.

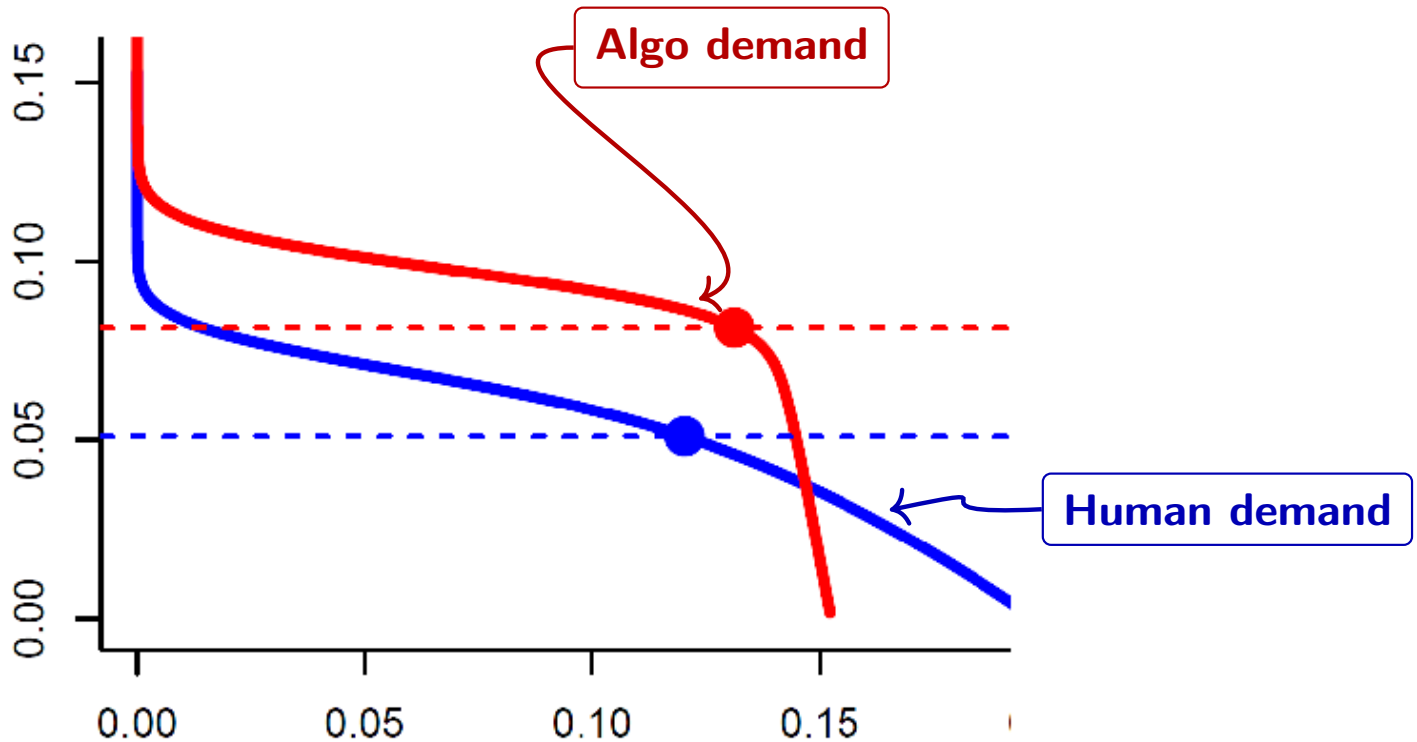

- Algo demand $\rightarrow$ lower perceived elasticity $\rightarrow$ higher p.

Demand Strikes Back: Agentic Buyers in Collusive Markets

- Buyers need not be passive price-takers with given preferences.
- In AI-driven markets, buyers can be autonomous agents, learning, adapting, strategically influencing sellers.

- We next consider agentic buyer vs. colluding Q-learning sellers as in CCDP (2020)
  - In particular, **agentic buyers deciding about transactions**: if to buy given the prices and, how much from each seller,...
  - similar to Amazon "Buy for Me", Walmart, Visa, Mastercard, PayPal "Agent Pay", Open AI's "Operator"

Agentic buyer vs colluding agentic sellers

- Buyer (quadratic) utility à la Singh and Vives
- Hence, linear demand represents buyer's behavior: $q_i = A - Bp_i + Cp_j$.
- Baseline/true parameters: $\bar{A}, \bar{B}, \bar{C}$ from standard utility maximization.

- One agentic buyer and two agentic sellers repeatedly interact.
- Agentic buyer can vary A, B, C within fixed grids (manipulating the demand), using QLearning algo
  - Assume we perfectly passed preferences to agentic buyer (value alignment sidestepped)
  - Important: manipulating the demand is a risk for the AI-buyer, because actual transactions may not reflect true preferences!
- Sellers set prices with one-period-memory QL

Benchmark: compare with *standard buyer*, where the buyer simply posts demand with the true parameters, facing two AI-sellers

Agentic Buyer: Prices

- Prices with agentic buyer significantly lower than with standard buyer.

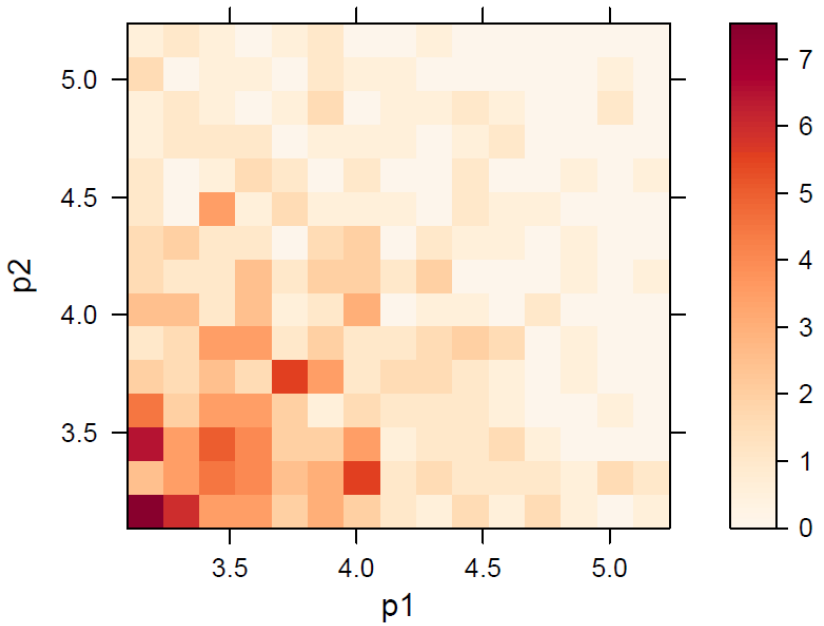

Agentic buyer

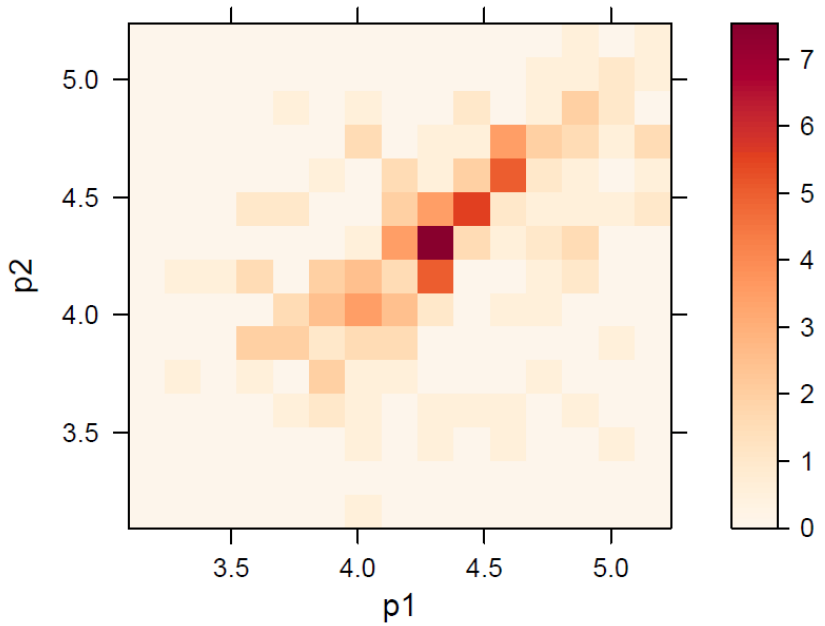

Standard buyer

Agentic Buyer: Quantities

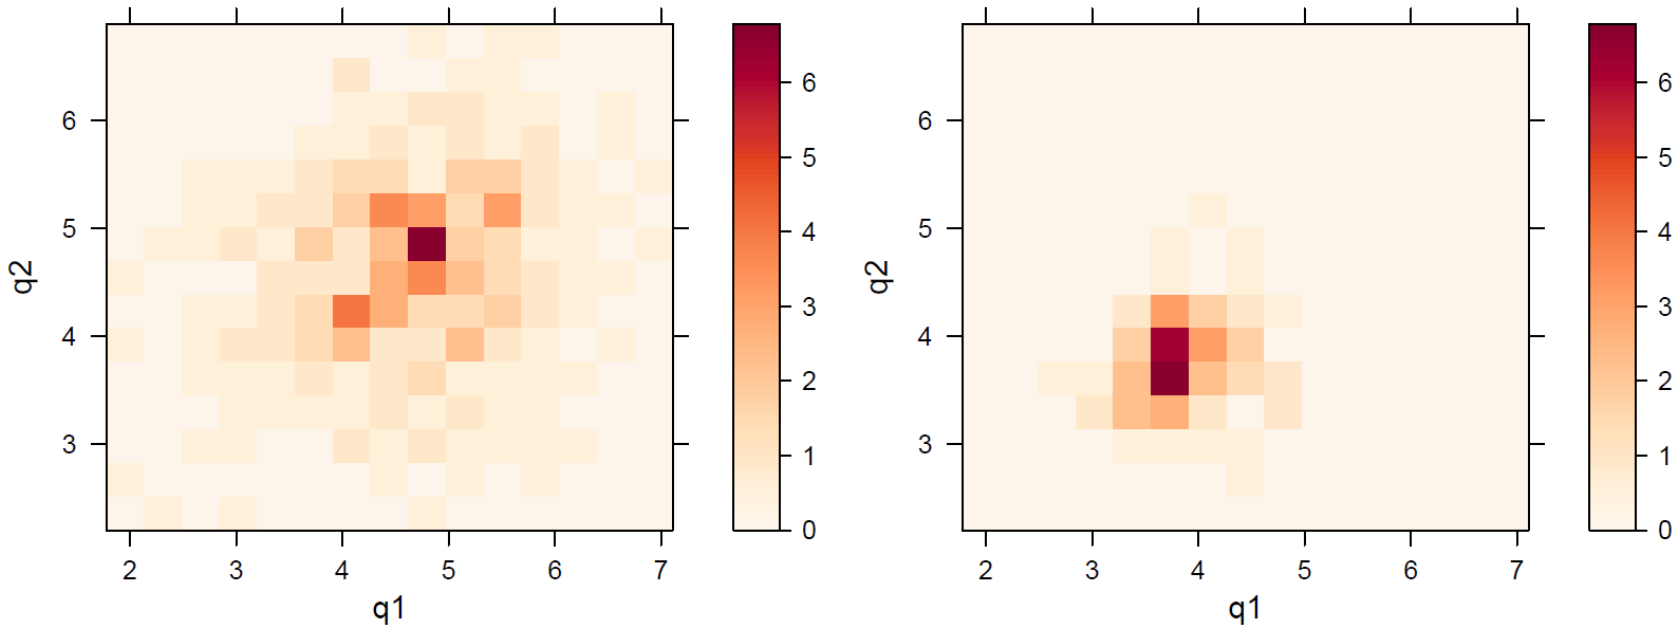

Agentic buyer

Standard buyer

- Larger market volume under agentic-buyer.

Profits, Consumer Surplus and Welfare

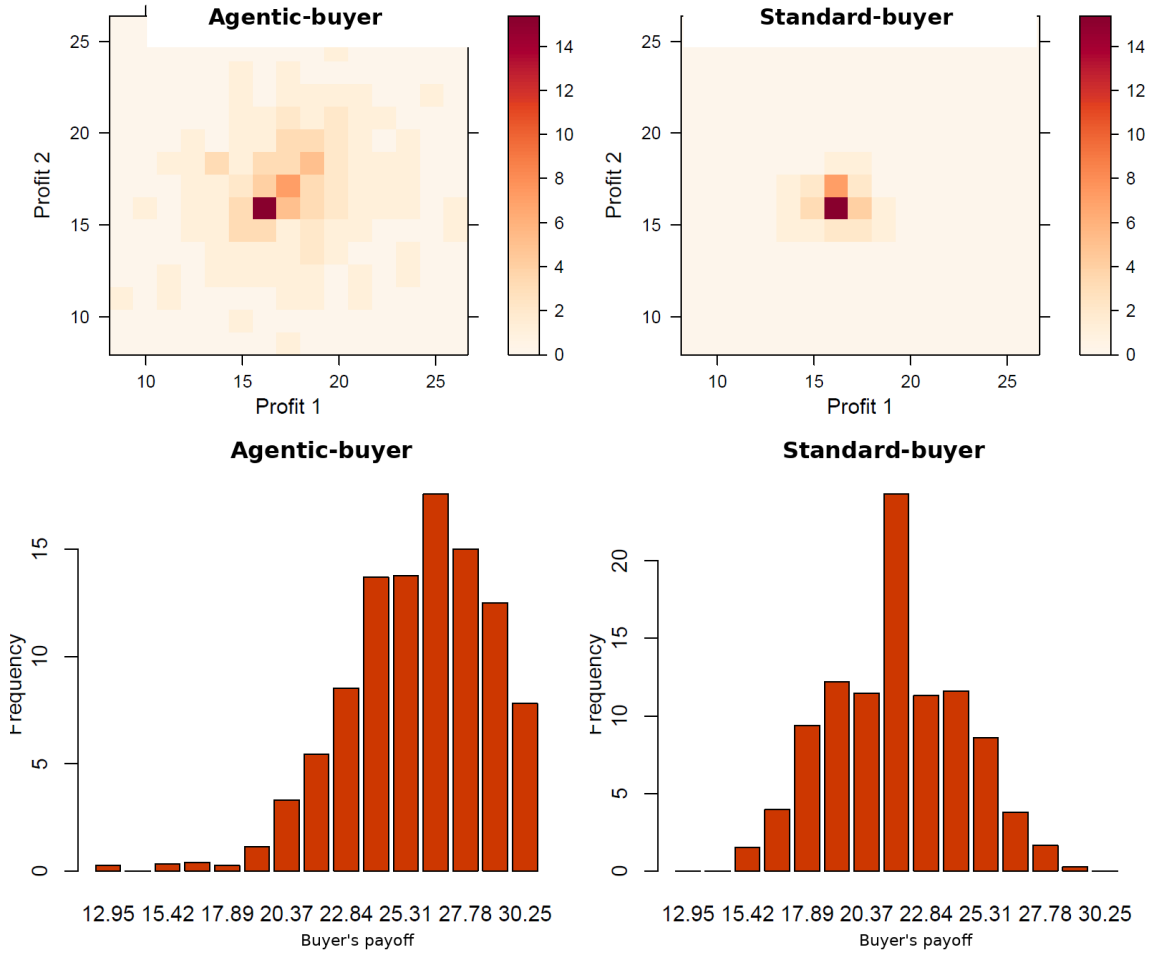

- Welfare: +10.8%, efficiency gap halves (18.76% $\rightarrow$ 9.99%).
- Buyer captures most of gain (sellers' profit not reduced)

Agentic buyer's manipulation Strategy

What strategy does the agentic buyer adopt?

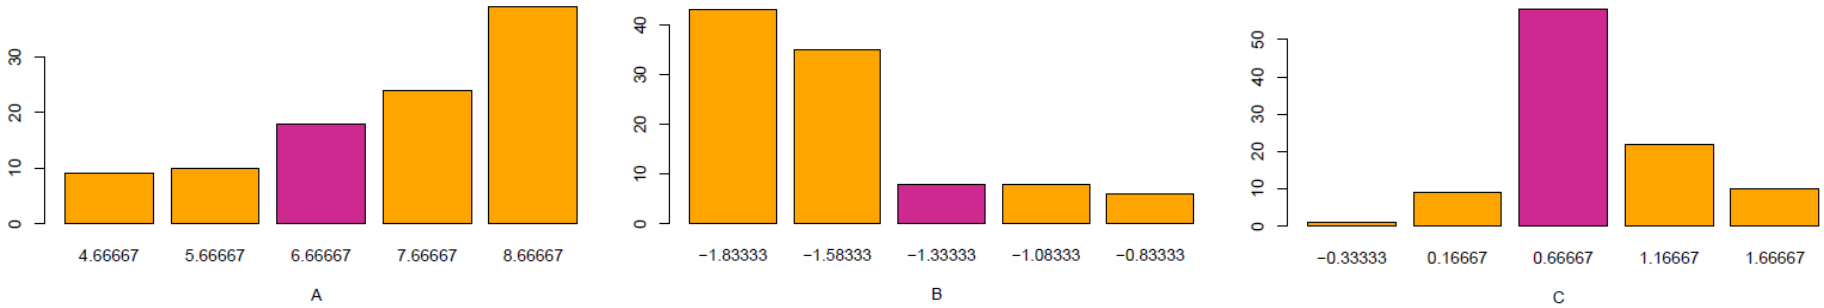

(magenta bars are the true values)

- Overstates demand intercept (A)
- Overstate own-price elasticity (B)
- Rarely changes cross-price elasticity (C)
- That is, it shifts demand outward and makes it (seem) more elastic.

Manipulation Strategy: favoring price cuts

Difference between agentic demand and standard truthful demand:

- Agentic buyer disproportionately increases demand of products with low price relative to its competitor's.
  - Amplifying quantity response to relative price differences,
  - This increasing incentive for undercutting.
  - Sellers forced to coordinate on lower prices to limit temptation.

- Subtle: preserves coordination while increasing undercut incentives.

AI vs. AI: Bilateral Monopoly

As mentioned, dual delegation pushes towards individualized markets.

So what about one agentic buyer with one agentic seller?

- Buyer same preferences as above: $D(p) = A - Bp$
- Seller from convex cost, and supply function $S(p) = C + Dp$.
- In standard benchmark, market clears at $p^* = \frac{A-C}{B+D}$, $q^* = \frac{AD+CB}{B+D}$.
- What if both agents can distort parameters choosing within grids?

Theory: Equilibrium Multiplicity

- Continuum of Nash equilibria: same q^* via different slopes/intercepts.
- Including, entire efficient frontier replicable.
- Some equilibria inefficient — low transaction volume from miscoordination.
- Surplus split varies.

How Q-Learning do buyer and seller (with memory) behave?

Results from repeated synthetic bilateral-monopoly markets

Agentic Demands and Supplies

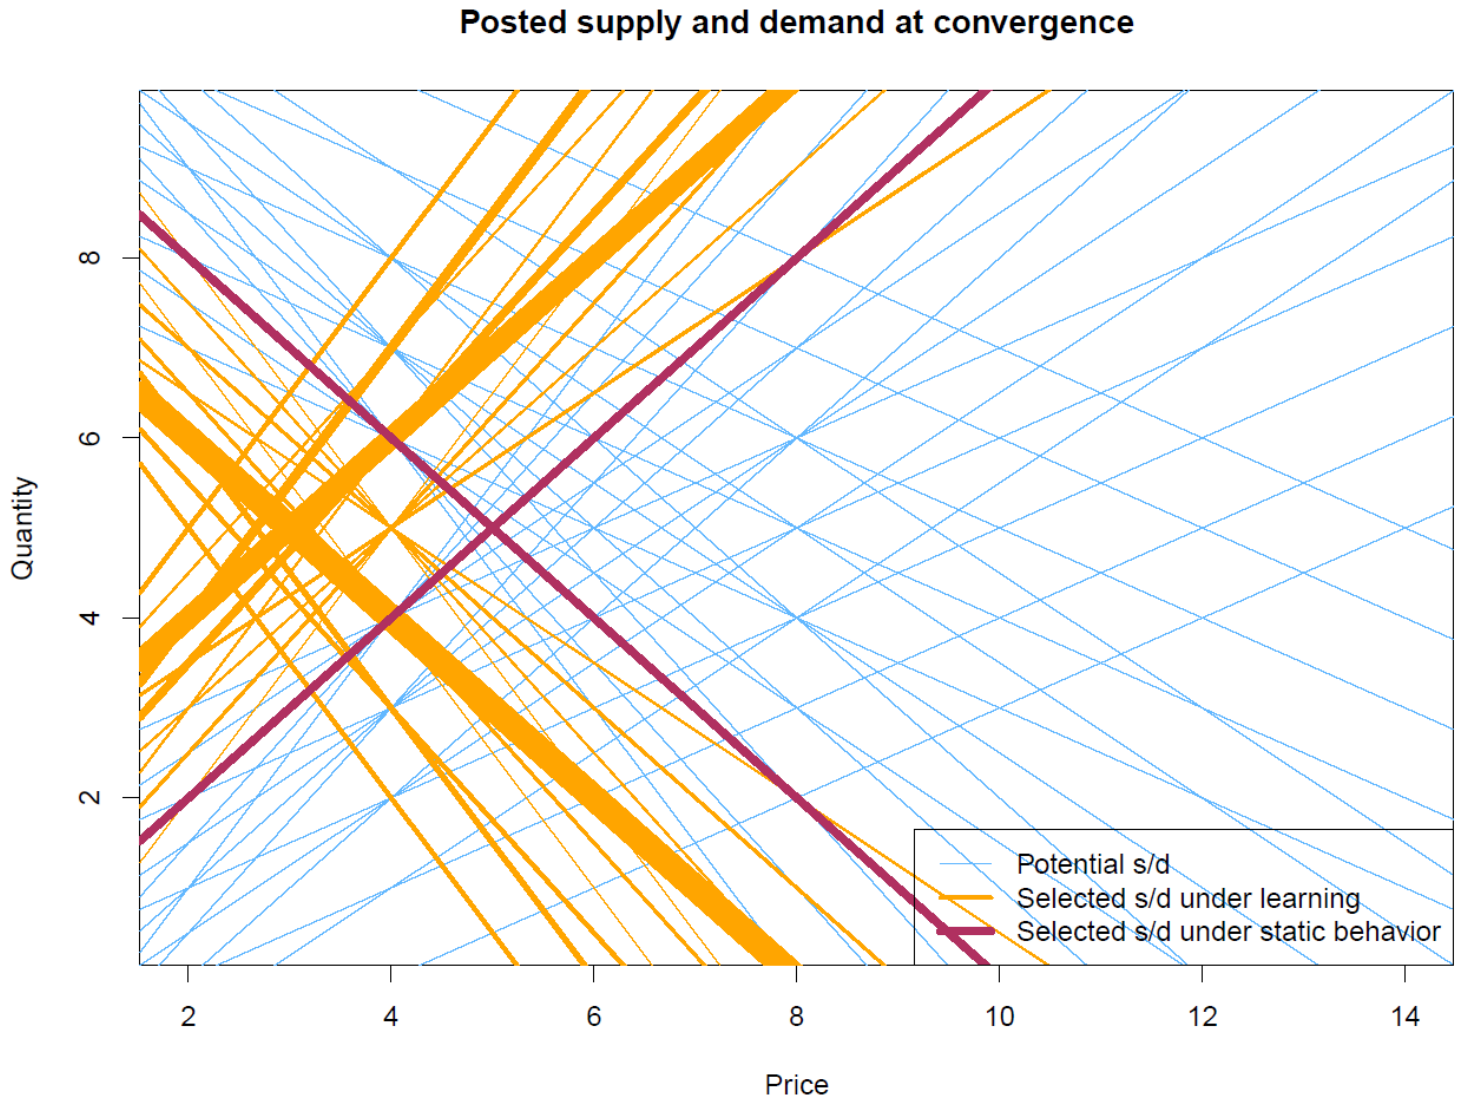

Static equilibrium $(q, p) = (5, 5)$

Parameters manipulation

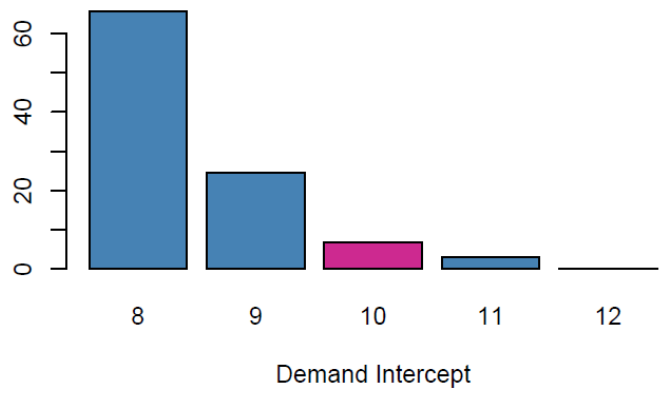

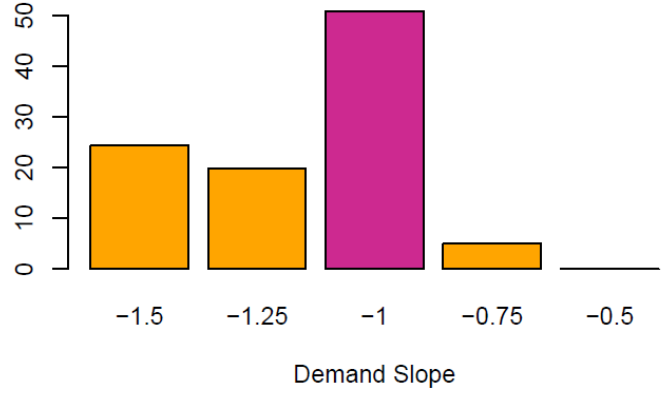

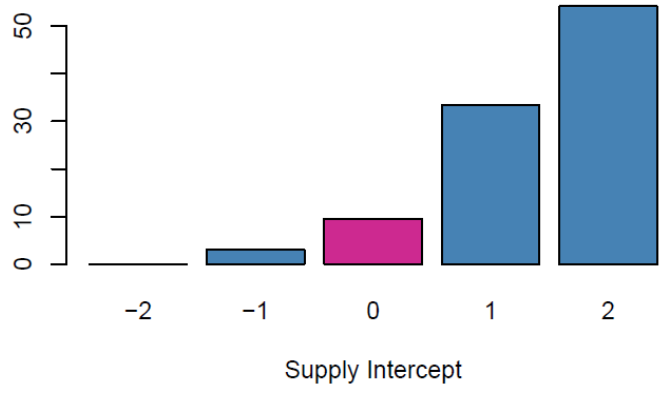

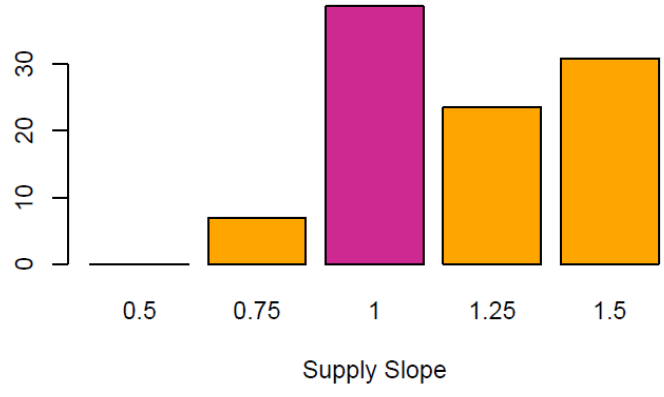

- Now they understate demand intercept $(A, -C)$
- Somehow overstate price elasticity (B, D)
- That is, they shifts D and S down, (slightly) more elastic.

Payoff Distribution and Efficiency Frontier

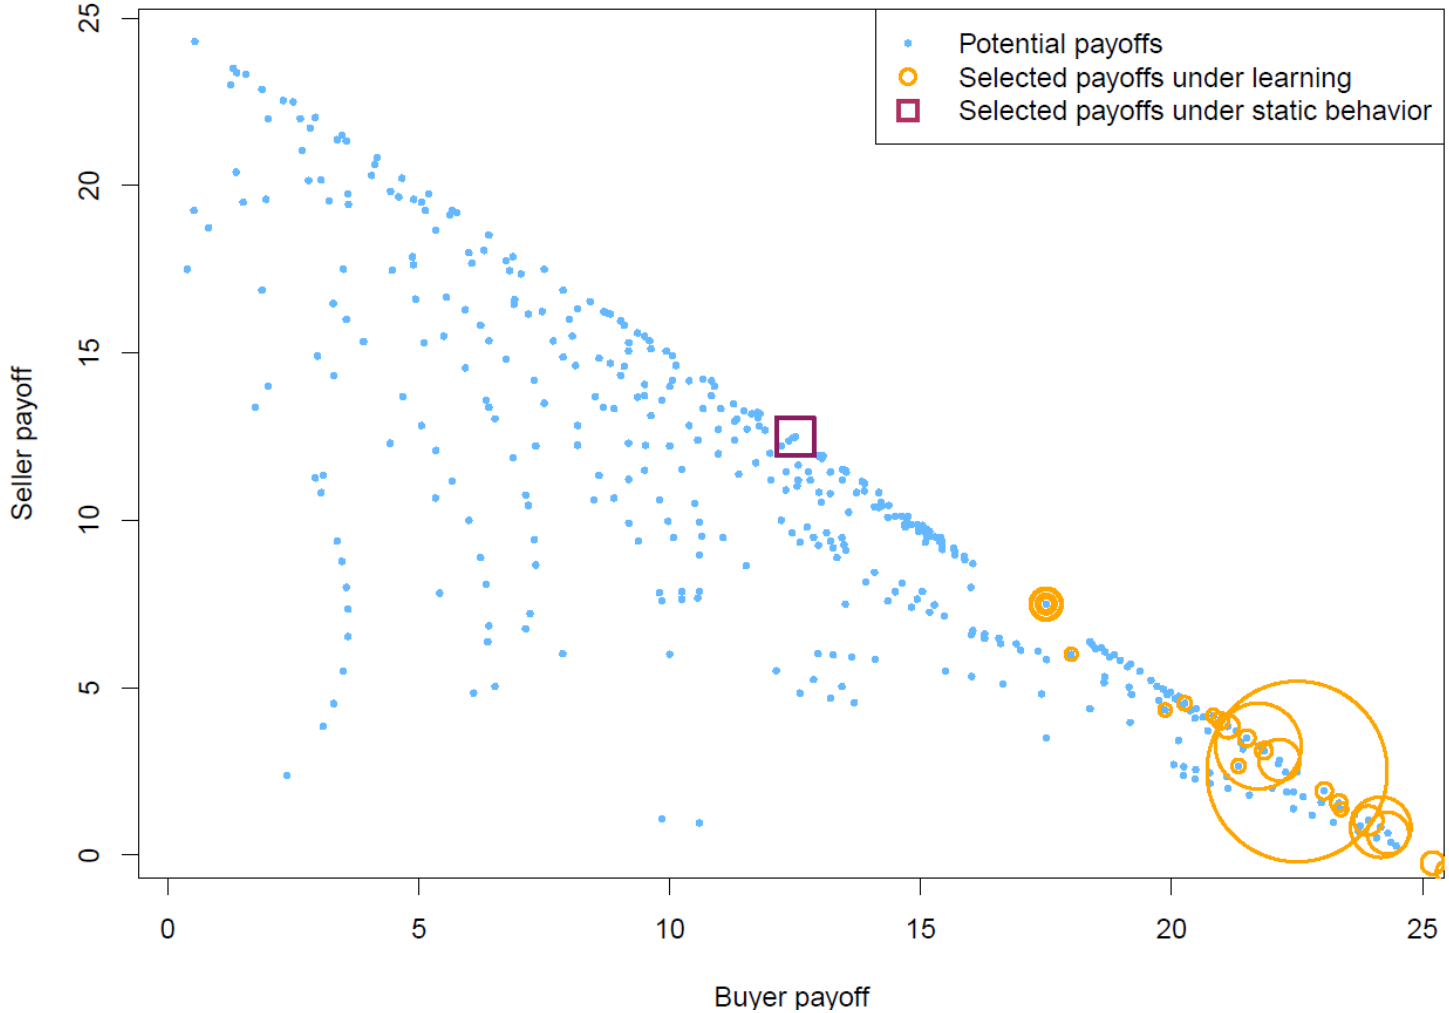

- Agents reach full efficiency (very often).
- In these cases, agentic buyer gets larger share of surplus (depends on the map of possible prices and quantities).

Where are we going?

Autonomous, learning agents on both sides of market may challenge standard competition models.

- **Future research:**
  - Passing preference to agentic-buyers: **the value alignment problem**
  - Formal models of AI–AI strategic interaction in multi-agent markets.
  - Behavioral regularities and convergence patterns in learning.
  - Frameworks to evaluate agentic markets.

- **Implications for real markets:**
  - Institutional design: market, protocols, disclosure, interoperability between agents...
  - Regulatory challenges: transparency & auditing, collusion, in general efficiency.

AI and Competition in Agentic Markets

Giacomo Calzolari

European University Institute

13th Postal Economics Conference
Toulouse, April 16–17, 2026

Thank you!