

# Topics in Economic Theory: AI-Driven Markets

---

Naples Ph.D. in Economics: Spring 2026

Niccolò Lomys\*

## 1 Overview

**Introduction.** Artificial intelligence (AI) and algorithmic decision-making are changing how markets work. Algorithms now set prices, place bids, allocate attention, recommend products, and mediate interactions between firms and consumers. In some settings, they are tools chosen by firms; in others, they become strategic agents, market interfaces, or synthetic decision-makers. These developments raise fundamental questions about market efficiency, competition policy, consumer welfare, and the design of rules for markets populated by autonomous or semi-autonomous agents.

This course studies these questions using tools and ideas from economic theory, industrial organization, econometrics, computer science, and artificial intelligence. We start from learning models and market-design background, then develop a core analysis of algorithmic collusion through simulations, theory, IO models, empirical evidence, and structural econometric methods. The final topics turn to adjacent frontiers: digital advertising and auto-bidding, LLMs as synthetic economic agents, LLM-mediated conversational markets, and the design and evaluation of algorithms when objectives such as accuracy, fairness, incentives, and welfare may conflict.

Throughout the course, we will ask how economic models and empirical tools should change when agents are algorithms, when platforms structure the informational environment, and when AI systems mediate strategic interaction. The course is intended both as an entry point to a rapidly growing literature and as a map of open research directions.

**Prerequisites.** This advanced Ph.D. course builds on first-year microeconomic theory and econometric sequences. Students should be comfortable with game theory, industrial organization, and empirical methods. Although not strictly required, familiarity with computer science concepts—particularly algorithmic game theory, complexity theory, reinforcement

---

\*CSEF and Università degli Studi di Napoli Federico II; [niccolomys@gmail.com](mailto:niccolomys@gmail.com).

learning, and basic machine learning—will be valuable for engaging with the technical material. The course combines theoretical analysis with applications, simulations, empirical evidence, and policy-oriented questions.

**Requirements.** I will distribute slides before each meeting. The reading list below is deliberately broader than the material covered in detail in class: it is intended both as a guide for the course and as a reference for future readings. Students are expected to:

- Attend all meetings.
- Select one paper from topics 3–11 to read and present in class.
- Select one paper from topics 3–11 to read and write a referee report on.

The presentation paper and the referee-report paper should be different. I will help you select the two papers from the references below based on your research interests.

## 2 Topics and Readings

Topics 1 and 2 provide the conceptual and technical background. Topics 3–8 form the core of the course, moving from algorithmic collusion to platform competition, empirical evidence, and structural econometric approaches. Topics 9–11 collect adjacent and frontier directions that naturally extend the agenda toward LLMs as synthetic economic agents, LLM-mediated conversational markets, AI-native advertising, algorithm design, fairness, and statistical evaluation.

### 1. Background and Surveys

This section introduces the main economic and policy questions raised by algorithmic competition. The readings provide broad overviews of autonomous pricing, algorithmic collusion, market design for AI agents, and the legal and regulatory challenges created by autonomous decision-makers. They are useful entry points for understanding why standard IO questions become more subtle when the strategic actors are learning algorithms.

- Abada, Harrington, Lambin, and Meylahn (2025). “Algorithmic Collusion: Where Are We and Where Should We Be Going?” *Working Paper*.
- Assad, Calvano, Calzolari, Clark, Denicolò, Ershov, Johnson, Pastorello, Rhodes, Xu, and Wildenbeest (2021). “Autonomous Algorithmic Collusion: Economic Research and Policy Implications.” *Oxford Review of Economic Policy*.
- Banchio and Skrzypacz (2022). “Market Design for AI Algorithms.” *ACM SIGecom Exchanges*.
- Bichler, Durmann, and Oberlechner (2025). “Algorithmic Pricing and Algorithmic Collusion.” *Business & Information Systems Engineering*.
- Chassang and Ortner (2023). “Regulating Collusion.” *Annual Review of Economics*.

- Harrington (2018). “Developing Competition Law for Collusion by Autonomous Artificial Agents.” *Journal of Competition Law & Economics*.

## 2. Learning, Regret, and Games

This section introduces the learning concepts used later to model algorithmic agents. Reinforcement learning studies sequential decision problems in which agents choose actions over time, observe payoff-relevant feedback, and learn policies when rewards and transition rules may be unknown. Online learning and regret minimization study adaptive decision rules evaluated by their performance relative to benchmark actions, policies, or deviations after observing feedback, such as full-information, bandit, or gradient feedback. Calibration and approachability provide related consistency notions for forecasts and play, and connect learning dynamics to equilibrium concepts such as correlated equilibrium and coarse correlated equilibrium. These references provide the technical language used later to describe pricing algorithms, auto-bidding agents, and empirical restrictions based on no-regret behavior; students are not expected to read all books cover to cover.

- Albrecht, Christianos, and Schäfer (2024). “Multi-Agent Reinforcement Learning: Foundations and Modern Approaches.” *MIT Press*.
- Bertsekas (2025). “A Course in Reinforcement Learning.” *Athena Scientific*.
- Cesa-Bianchi and Lugosi (2006). “Prediction, Learning, and Games.” *Cambridge University Press*.
- Kroer (2026). “Games, Markets, and Online Learning.” *Cambridge University Press*.
- Nisan, Roughgarden, Tardos, and Vazirani (2007). “Algorithmic Game Theory.” *Cambridge University Press*.
- Roughgarden (2016). “Twenty Lectures on Algorithmic Game Theory.” *Cambridge University Press*.
- Shoham and Leyton-Brown (2009). “Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.” *Cambridge University Press*.
- Sutton and Barto (2018). “Reinforcement Learning: An Introduction.” *MIT Press*.
- Online resources:
  - Simons Institute for the Theory of Computing at UC Berkeley;
  - Program on Learning and Games (Spring 2022);
  - Video lectures and slides at <https://simons.berkeley.edu/programs/games2022>.

## 3. Simulation Evidence

This section studies what happens when artificial agents repeatedly compete in stylized market environments. The main question is whether standard learning algorithms, especially

variants of  $Q$ -learning, can sustain prices above competitive levels without explicit communication or agreement. The readings also show how design choices—such as the state space, action space, timing, exploration, and reward feedback—affect collusive-looking outcomes.

- Abada and Lambin (2023). “Artificial Intelligence: Can Seemingly Collusive Outcomes be Avoided?” *Management Science*.
- Asker, Fershtman, and Pakes (2024). “The Impact of Artificial Intelligence Design on Pricing.” *Journal of Economics & Management Strategy*.
- Calvano, Calzolari, Denicolò, and Pastorello (2020). “Artificial Intelligence, Algorithmic Pricing, and Collusion.” *American Economic Review*.
- Klein (2021). “Autonomous Algorithmic Collusion:  $Q$ -learning under Sequential Pricing.” *The RAND Journal of Economics*.
- Tesauro and Kephart (2002). “Pricing in Agent Economies Using Multi-Agent  $Q$ -Learning.” *Autonomous Agents and Multi-Agent Systems*.
- Waltman and Kaymak (2008). “ $Q$ -learning Agents in a Cournot Oligopoly Model.” *Journal of Economic Dynamics and Control*.

#### 4. Theory of Algorithmic Collusion

This section develops analytical explanations for when and why algorithms may support collusive, supra-competitive, or otherwise non-standard outcomes. The readings ask whether algorithmic behavior can be connected to repeated-game logic, whether learning dynamics implement folk-theorem-like outcomes, and how monitoring, market primitives, and algorithm design affect stability. They also clarify the distinction between explicit collusion, tacit coordination, and independent algorithmic behavior that nevertheless leads to high prices.

- Askenazi-Golan, Mergoni Cecchelli, Plumb, and Possnig (2026). “The Bounds of Algorithmic Collusion;  $Q$ -learning, Gradient Learning, and the Folk Theorem.” *Working Paper*.
- Banchio and Mantegazza (2026). “Artificial Intelligence and Spontaneous Collusion.” *American Economic Review*, forthcoming.
- Cho and Williams (2026). “Collusive Outcomes Without Collusion: Algorithmic Pricing in a Duopoly Model.” *Working Paper*.
- Hansen, Misra, and Pai (2021). “Frontiers: Algorithmic Collusion: Supra-competitive Prices via Independent Algorithms.” *Marketing Science*.
- Lamba and Zhuk (2026). “Pricing with Algorithms.” *American Economic Review: Insights*, forthcoming.
- Possnig (2026). “Monitoring, Market Primitives, and the Stability of Algorithmic Collusion.” *Working Paper*.

- Salcedo (2015). “Pricing Algorithms and Tacit Collusion.” *Working Paper*.

## 5. IO, Platforms, and Personalization

This section broadens the analysis from learning dynamics to IO environments in which firms, platforms, consumers, and algorithm providers interact strategically. The readings study outsourcing of pricing algorithms, platform design when sellers use algorithms, competition between humans and algorithms, personalized pricing, and algorithmic recommendations. The goal is to understand how market structure, data, platform rules, and algorithmic delegation shape competitive outcomes.

- Brown and MacKay (2023). “Competition in Pricing Algorithms.” *American Economic Journal: Microeconomics*.
- Dubé and Misra (2023). “Personalized Pricing and Consumer Welfare.” *Journal of Political Economy*.
- Harrington (2022). “The Effect of Outsourcing Pricing Algorithms on Market Competition.” *Management Science*.
- Ichihashi and Smolin (2025). “Buyer-Optimal Algorithmic Recommendations.” *Working Paper*.
- Johnson, Rhodes, and Wildenbeest (2023). “Platform Design when Sellers Use Pricing Algorithms.” *Econometrica*.
- Leisten (2024). “Algorithmic Competition, with Humans.” *Working Paper*.
- Miklós-Thal and Tucker (2019). “Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?” *Management Science*.
- Rhodes and Zhou (2024). “Personalized Pricing and Competition.” *American Economic Review*.

## 6. Advertising, Auto-bidding, and Data

This section studies markets in which algorithmic agents bid, allocate attention, and use data at scale. Online advertising is a central example: platforms design auctions, advertisers rely on auto-bidding tools, and data affects both ad allocation and product-market competition. These readings connect algorithmic competition to auction design, budget constraints, platform data, and downstream product prices.

- Aggarwal et al. (2024). “Auto-bidding and Auctions in Online Advertising: A Survey.” *ACM SIGecom Exchanges*.
- Balseiro and Gur (2019). “Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium.” *Management Science*.
- Bergemann and Bonatti (2024). “Data, Competition, and Digital Platforms.” *American Economic Review*.

- Bergemann, Bonatti, and Wu (2025). “Bidding with Budgets: Data-Driven Bid Algorithms in Digital Advertising.” *International Journal of Industrial Organization*.
- Bergemann, Bonatti, and Wu (2025). “How Do Digital Advertising Auctions Impact Product Prices?” *The Review of Economic Studies*.

## 7. Empirical Evidence

This section turns to evidence from real markets. The readings examine how algorithmic pricing affects competition in gasoline, rental housing, and e-commerce markets, and they illustrate the empirical challenges of measuring algorithm adoption, identifying competitive effects, and distinguishing price coordination from efficiency gains or ordinary strategic interaction. These papers also show how institutional detail and market-specific data shape empirical work on AI-driven competition.

- Assad, Clark, Ershov, and Xu (2024). “Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market.” *Journal of Political Economy*.
- Calder-Wang and Kim (2026). “Algorithmic Pricing in Multifamily Rentals: Efficiency Gains or Price Coordination?” *Working Paper*.
- Musolf (2025). “Algorithmic Pricing, Price Wars and Tacit Collusion: Evidence from E-Commerce.” *Management Science*, forthcoming.

## 8. Econometrics for AI and Dynamic Games

This section examines how economists can identify, estimate, and test models of strategic behavior when agents may be learning algorithms rather than fully rational equilibrium players. The readings connect structural methods for dynamic games to newer approaches based on no-regret learning, behavioral restrictions, and tests of algorithmic collusion. The broader goal is to develop empirical IO and econometric tools for markets in which AI systems make repeated strategic decisions.

- Cozzolino, Gualdani, Gufler, Lomys, and Magnolfi (2026). “Robust Identification in Repeated Games: An Empirical Approach to Algorithmic Competition.” *Working Paper*.
- Hartline, Long, and Zhang (2024). “Regulation of Algorithmic Collusion.” *Proceedings of the 2024 Symposium on Computer Science and Law*.
- Hartline, Wang, and Zhang (2025). “Regulation of Algorithmic Collusion, Refined: Testing Pessimistic Calibrated Regret.” *Proceedings of the 2025 Symposium on Computer Science and Law*.
- Lomys and Magnolfi (2024). “Estimation of Games under No Regret: Structural Econometrics for AI.” *Working Paper*.
- Nekipelov, Syrgkanis, and Tardos (2015). “Econometrics for Learning Agents.” *Proceedings of the Sixteenth ACM Conference on Economics and Computation: EC 2015*.

## 9. LLMs as Economic Agents and Synthetic Subjects

This section studies whether LLMs can serve as synthetic subjects, simulated economic agents, and experimental infrastructure for economic theory and market design. The readings distinguish several roles: LLMs as substitutes or complements for human subjects in behavioral experiments; LLM agents as models of boundedly rational economic behavior; LLM systems as tools for generating, testing, and refining social-scientific hypotheses; and synthetic agents as low-cost laboratories for auctions, bargaining, search, and market design. The goal is not to treat simulated evidence as a substitute for field or laboratory evidence, but to understand when synthetic agents can generate useful hypotheses, stress-test mechanisms, and discipline models of strategic behavior.

- Aher, Arriaga, and Kalai (2023). “Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies.” *Proceedings of the 40th International Conference on Machine Learning*.
- Horton, Filippas, and Manning (2026). “Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?” *Working Paper*.
- Manning, Zhu, and Horton (2024). “Automated Social Science: Language Models as Scientist and Subjects.” *Working Paper*.
- Manning and Horton (2026). “General Social Agents.” *Working Paper*.
- Shah, Zhu, Jiang, Wang, Dayi, Horton, and Parkes (2025). “Learning from Synthetic Labs: Language Models as Auction Participants.” *Working Paper*.
- Shahidi, Rusak, Manning, Fradkin, and Horton (2025). “The Coasean Singularity? Demand, Supply, and Market Design with AI Agents.” *Working Paper*.

## 10. LLMs and Conversational Markets

This section considers markets in which LLMs act as strategic agents, intermediaries, or conversational interfaces. The readings study LLMs as pricing, bargaining, recommendation, and advertising technologies, as well as mechanisms for AI-generated content and conversational product recommendations. The central question is how market design and competition change when consumers and firms interact through generative AI systems rather than through standard search, posted prices, or static advertising formats.

- Banchio, Mehta, and Perloth (2025). “Ads in Conversations.” *Working Paper*.
- Bergemann, Bojko, Dütting, Paes Leme, Xu, and Zuo (2026). “From Conversations to Mechanisms: Aligning Advertiser Incentives in AI-Powered Product Recommendations.” *Working Paper*.
- Bergemann, Bonatti, and Smolin (2026). “Menu Pricing of Large Language Models.” *Working Paper*.

- Bergemann, Ghili, Hu, Li, and Yang (2026). “Training Language Models for Bilateral Trade with Private Information.” *Working Paper*.
- Dütting, Mirrokni, Paes Leme, Xu, and Zuo (2024). “Mechanism Design for Large Language Models.” *Proceedings of the ACM Web Conference 2024*.
- Fish, Gonczarowski, and Shorrer (2026). “Algorithmic Collusion by Large Language Models.” *Working Paper*.
- Fudenberg and Liang (2026). “Friend or Foe: Delegating to an AI Whose Alignment is Unknown.” *Working Paper*.
- Gao, Han, and Liang (2026). “How Well Do LLMs Predict Human Behavior? A Measure of Their Pretrained Knowledge.” *Working Paper*.
- Hajiaghayi, Lahaie, Rezaei, and Shin (2024). “Ad Auctions for LLMs via Retrieval Augmented Generation.” *Advances in Neural Information Processing Systems*.
- Liang (2026). “Artificial Intelligence Clones.” *Working Paper*.

## 11. Algorithm Design and Evaluation

This final section broadens the course from market outcomes to the design and evaluation of algorithms themselves. The readings study how to formalize trade-offs among accuracy, fairness, and other objectives, and how to conduct statistical inference to determine whether an algorithm can be improved across multiple dimensions. These papers connect AI-driven market dynamics to the broader economics of algorithmic governance, regulation, and evaluation.

- Auerbach, Liang, Okumura, and Tabord-Meehan (2025). “Testing the Fairness-Accuracy Improvability of Algorithms.” *Working Paper*.
- Liang, Lu, Mu, and Okumura (2026). “Algorithm Design: A Fairness-Accuracy Frontier.” *Journal of Political Economy*, forthcoming.
- Liu and Molinari (2025). “Inference for an Algorithmic Fairness-Accuracy Frontier.” *Working Paper*.