

MKT927: INTRO TO QUANTITATIVE MARKETING

Prof. Andrey Fradkin

Lecture 12: AI in the Wild

TODAY'S AGENDA

- AI in the wild: Pricing
- AI in the wild: Advertising Algorithms
- Discussions: Agarwal et al., Karlinsky-Shichor
- Time to talk about projects if you would like.

ONLINE COMPETITION AND ALGORITHMS

- Online markets were expected to lead to near-perfect competition.
- In practice: large price dispersion and dominance by few firms.
- Key question: What role do pricing algorithms play?
- Even simple algorithms can fundamentally change market outcomes.

THREE STYLIZED FACTS (BROWN AND MACKAY (2023))


- 1: Firms have heterogeneous pricing technology (update frequency).
- 2: Faster firms respond more quickly to rivals' price changes.
- 3: Faster firms tend to have persistently lower prices.

The image shows the homepage of repricer.com. At the top, there is a navigation bar with the company logo, links for 'Features', 'Pricing', 'Contact', and 'Blog', and two buttons: 'Book a demo' and 'Sign Up'. The main headline reads 'Beat your competitors with the fastest Amazon Repricer on the market'. Below this, a sub-headline states: 'Generate more sales, maximize profit and manage your pricing with super-fast repricing. Reach millions of new customers with our new Google Shopping integration.' There are three buttons: 'Get Free 14 Day Trial', 'Amazon repricer.com', and 'Join Partner'. A woman in a red dress is shown on the right side of the page, with large orange arrows pointing to the right behind her. Below the woman is a colorful wave graphic. Text below the graphic says 'Join the 5,000+ retailers who trust Repricer.com to automate pricing across 15 Amazon marketplaces'. A row of 15 national flags is displayed below that. At the bottom, it says 'also supporting' followed by logos for 'ebay', 'Walmart', and 'Google'. The footer contains the text '© Mackay Competition in Pricing Algorithms'.


HETEROGENEITY IN PRICING TECHNOLOGY

Figure 2: Heterogeneity in Pricing Technology


Retailer A and B



Retailer C



Retailer E and D



MODEL OVERVIEW

- Firms choose price p_j and possibly a pricing algorithm $\sigma_j(p_{-j})$.
- Algorithms update at intervals γ_j .
- Model nests static Bertrand, sequential games, and algorithmic pricing.
- Firms face differentiated demand.

HOW ALGORITHMS RAISE PRICES

- Fast firms commit to undercutting rivals.
- This discourages rivals from competing on price.
- Leads to prices between Bertrand and Stackelberg levels.
- Algorithmic commitment softens competition.

ENDOGENOUS TECHNOLOGY ADOPTION

- Firms can adopt faster pricing at a cost.
- Equilibrium may involve asymmetric adoption.
- Result: higher prices and higher profits for both firms.
- Simultaneous pricing not an equilibrium.

COUNTERFACTUAL SIMULATION

- Simulate market under symmetric (Bertrand) pricing.
- Baseline: firms have asymmetric pricing technology.
- Algorithmic pricing raises prices by 5.2%, profits by 10%.
- Estimated \$300 million/year reduction in consumer surplus.

CALVANO, CALZOLARI, AND PASTORELLO (2020)

- Note, Brown and MacKay did not have algorithms that learn.
- Key question for CCP2020: Can reinforcement learning algorithms autonomously learn to collude?
- Approach: Simulate Q-learning agents in repeated oligopoly pricing games.
- Finding: Algorithms often learn to sustain supracompetitive prices.

WHAT IS Q-LEARNING?

- Reinforcement learning algorithm that learns value of actions from rewards.
- Agents choose actions based on a state, update value estimates via rewards.
- Q-function: expected long-term reward for each state-action pair.
- Exploration vs. exploitation trade-off via ϵ -greedy strategy.

COMPUTATIONAL EXPERIMENT SETUP

- Stage game: logit demand, differentiated products, repeated Bertrand pricing.
- Agents are symmetric; action space discretized between Bertrand and Monopoly.
- Agents observe rival prices with bounded memory.
- Parameters tuned for slow, persistent learning.

KEY RESULT: SUPRACOMPETITIVE PRICES

- Convergence takes a long time, often > 100,000 iterations.
- Algorithms converge to pricing strategies well above Nash.
- Strategies learned through trial-and-error.
- Prices below monopoly, but significantly above competitive levels.
- Emergent collusion without communication.




Figure 1. The distribution of prices charged by reinforcement-learning price algorithms in the virtual market created in [Calvano et al. \(2020\)](#). The price that would maximize the firms' joint profit is just above 1.93. The algorithms routinely learned to collude.

MECHANISM: REWARD-PUNISHMENT STRATEGY

- Algorithms punish defection with temporary price drops.
- Punishment followed by gradual return to cooperation.
- Similar to stick-and-carrot strategies, not grim-trigger.
- Deviations typically reduce long-run profits.




FIGURE 4

Notes: Prices charged by the two algorithms in period τ after an exogenous price cut by one of them in period $\tau = 1$. The forced cheater deviates to the static best response, and the deviation lasts for one period only. The figure plots the average prices across the 1,000 sessions. For sessions leading to a price cycle, we consider deviations starting from every point of the cycle and take the average of all of them. This counts as one observation in the calculation of the overall average.

ROBUSTNESS OF FINDINGS

- Results hold with:
 - Asymmetric costs or demand
 - More competitors ($n = 3$ or 4)
 - Stochastic environments
- Collusion persists even with limited exploration and high noise.

ASSAD ET AL. (2024)

- The first empirical study of algorithmic pricing (AP) adoption and competition.
- Focus: German retail gasoline market, 2016–2018.
- Key idea, adoption of AP took place at different times in different stations. Identify when this happened by looking at changes in pricing patterns.
- But adoption is endogenous, so instrument: fraction of brand's stations classified as adopters (HQ-led rollouts).

MAIN RESULTS: STATION-LEVEL EFFECTS

- Adoption increases margins by 1.3 cents/litre (15
- No effect for monopolists.
- In duopoly/triopoly markets, effect only if **all** stations adopt.
 - No adoption: baseline margins
 - Some adopt: no change
 - All adopt: 3.2 cents/litre higher margins (approx. 38%)

POLICY IMPLICATIONS

- Tacit collusion can emerge algorithmically.
- Traditional antitrust frameworks focus on explicit agreements.
- AI-driven pricing raises new enforcement challenges.
- Potential need to regulate pricing behavior, not just communication.

DISCUSSION QUESTIONS

1. None of the above studies considered LLMs as algorithms that set prices. How might LLMs be different?
2. What are some promising settings for studying algorithmic pricing that have not been considered above?
3. Are there research questions that are left unanswered by these papers?

ALGORITHMIC BIAS? AN EMPIRICAL STUDY OF APPARENT
GENDER-BASED DISCRIMINATION IN THE DISPLAY OF STEM
CAREER ADS
ANJA LAMBRECHT AND CATHERINE TUCKER

INTRODUCTION: THE PROBLEM

- Algorithms increasingly automate decisions (e.g., ad delivery).
- Concern: Automated choices might produce discriminatory outcomes.
- Previous research documented discriminatory patterns in ad display (Sweeney 2013, Datta et al. 2015) but often didn't identify the *why*.
- This study focuses on ads for STEM (Science, Technology, Engineering, Math) careers.
- Policy goal: Encourage more people, especially women, into STEM.
- Question: Why might algorithms display STEM ads differently based on gender, even when intended to be neutral?

IMAGE AND AD SETTINGS

Figure 1. (Color online) Sample Ad



in the study. These countries tend to be wealthier ones,

Figure 2. Ad Targeting Settings: Ad Intended to Be Shown to Both Men and Women Aged 18+ Years

Location	People who live in this location ✓
	United States ✓
Age	18 + ✓
Gender	All Men Women ✓

FIELD TEST METHODOLOGY

- Ran ad campaigns on Facebook directing users to the site.
- Tested in 191 countries.
- **Key Setting:** Ad explicitly targeted to be gender-neutral (Men and Women, 18+).
- Facebook uses an auction mechanism considering advertiser bids and a "relevance score" (quality score based on expected interaction).
- Initial bid: \$0.20/click, raised up to \$0.60 in some (wealthier) countries if needed to reach >5,000 views.

DATA OVERVIEW

- Data aggregated by Facebook at the demographic group (Age x Gender) - Country level.
- Key metrics: Impressions (times ad shown), Reach (unique people shown), Clicks, Cost per Click (CPC).
- Age groups: 18-24, 25-34, 35-44, 45-54, 55-64, 65+.

SUMMARY STATISTICS

Table 1. Summary Statistics Reported at the Demographic Group–Country Level

	Mean	Standard deviation	Minimum	Maximum
Impressions	1,911.8	2,321.4	0	24,980
Clicks	3.00	4.52	0	42
Unique clicks	2.78	4.15	0	40
Cost per click	0.085	0.090	0	0.66
Reach	615.6	850.7	0	13,436
Frequency	4.38	4.32	1	53

CORE FINDING: FEWER WOMEN SEE THE AD

- **Observation:** The ad was shown significantly more often to men than to women (over 20% difference overall).
- The difference was particularly pronounced for younger individuals (prime career years).
- Regression analysis confirms: Women, especially younger women (25-44), had significantly lower ad reach than men, controlling for country effects.
- *Reach* (unique individuals) shows a similar pattern to *Impressions*.

EXPLANATION 1: LEARNED FROM USER BEHAVIOR?

- **Hypothesis 1a:** Algorithm learned women click less, so shows ad less to optimize clicks.
- **Finding:** *Rejected.* If shown the ad, women were *more* likely to click than men (0.167% vs 0.131% CTR).
- **Hypothesis 1b:** Fewer women available on the platform.
- **Finding:** *Rejected.* Data suggests women use Facebook more actively than men globally and in the US. Sufficient pool of women interested in STEM exists.

EXPLANATION 2: LEARNED FROM CULTURAL BIAS?

- **Hypothesis:** Algorithm learned country-specific gender biases (e.g., women less likely to pursue STEM in certain cultures).
- Tested using World Bank data on:
 - Female labor market participation
 - Female education levels (primary, secondary)
 - Gender equality index (CPIA)
 - Country GDP (rich vs. poor)
- **Finding:** *Rejected.* These country-level factors did *not* significantly explain the gender difference in ad reach. The pattern held across different cultural/economic contexts.

EXPLANATION 3: ECONOMICS OF AD DELIVERY (COMPETITIVE SPILLOVERS)

- **Hypothesis:** Showing ads to women is more expensive due to competition from other advertisers targeting women.
- Collected secondary data on Facebook's *suggested bids* for different demographics.
- **Mechanism:** A cost-optimizing algorithm, faced with higher prices for female eyeballs, will naturally show the ad less often to women to stay within budget or maximize reach/clicks per dollar, even if the ad itself is gender-neutral.

WOMEN COST MORE TO REACH

Table 7. In General, Women Are More Expensive to Advertise to on Social Media: Competitive Spillovers from Other Advertisers' Decisions May Explain Our Finding

	(1) <i>Average suggested bid</i>	(2) <i>Average suggested bid</i>	(3) <i>Average suggested bid</i>
<i>Female</i>	0.053* (0.02)	0.05* (0.02)	-0.05 (0.04)
<i>Female × Age 18–24 years</i>			0.06* (0.04)
<i>Female × Age 25–34 years</i>			0.17* (0.09)
<i>Female × Age 35–44 years</i>			0.15*** (0.04)
<i>Female × Age 45–54 years</i>			0.08 (0.05)
<i>Female × Age 55–64 years</i>			0.13** (0.04)
<i>Age 18–24 years</i>	-0.01 (0.03)	-0.01 (0.03)	-0.04 (0.04)
<i>Age 25–34 years</i>	0.08 (0.05)	0.08 (0.05)	-0.01 (0.04)
<i>Age 35–44 years</i>	0.07* (0.03)	0.07* (0.04)	-0.00 (0.04)
<i>Age 45–54 years</i>	0.06 (0.04)	0.06 (0.04)	0.02 (0.06)
<i>Age 55+ years</i>	0.02 (0.03)	0.02 (0.03)	-0.04 (0.04)
<i>Country controls</i>	No	Yes	Yes
<i>Observations</i>	2,096.00	2,096.00	2,096.00
<i>Log-likelihood</i>	-2,096.47	-1,219.82	-1,214.99
<i>R²</i>	0.00	0.57	0.57

Notes. Ordinary least squares estimates. Dependent variable is average suggested bid. Omitted demographic groups are those aged between 13 and 17 years and those of the male gender. Robust standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$.

WHY ARE FEMALE EYEBALLS MORE EXPENSIVE?

- Marketing literature/business press suggests women (esp. younger) are a prized demographic.
- Often control household spending.
- **Spillover Effect:** Competition in sectors like retail (where targeting women might be profitable) increases ad costs for women, impacting ad delivery in unrelated sectors like STEM careers.

GENERALIZABILITY ACROSS PLATFORMS

- Tested similar STEM ad campaigns on other major platforms (US market):
 - **Google Display Network:** Men received significantly more impressions (51% vs 36%). Women had higher CTR but slightly higher CPC.
 - **Instagram:** Men received vastly more impressions (85% vs 15%). Women had lower CTR and much higher CPC here.
 - **Twitter:** Men received more impressions (56% vs 44%).
- **Conclusion:** The pattern of men seeing more STEM ads appears characteristic of the broader online advertising ecosystem, not just Facebook.

IMPLICATIONS AND CONCLUSION

- Gender-neutral ads can be delivered in an apparently discriminatory way due to economic forces (higher ad costs for women driven by competition), not necessarily biased algorithms or user behavior.
- **Policy Challenge 1: Regulation Difficulty**
 - Algorithmic transparency might not reveal the issue (code shows cost optimization, not intent).
 - Observing a biased outcome doesn't automatically mean discriminatory intent.
- **Policy Challenge 2: Anti-Discrimination Laws**
 - Trying to *correct* the imbalance by running separate campaigns for employment ads is often prohibited by platforms.
 - Creates a tension: Can't easily fix algorithmic imbalances using targeting tools.

DISCUSSION QUESTIONS

1. What are the limitations of using suggested bid data as a proxy for actual auction dynamics and competitive pressure? How else could this mechanism be tested?
2. The study focuses on gender. How might similar economic spillover effects manifest for other protected characteristics (race, age, etc.) in different ad contexts (housing, credit)?
3. Critically evaluate the proposed solution of platforms offering "equal distribution" options. What are the potential economic costs, implementation challenges, and unintended consequences? Could it be gamed?

OTHER ALGORITHMIC CONSEQUENCES PAPERS:

- Uber's surge pricing algorithm (Castillo's paper), some papers about Airbnb's algorithms.
- CS literature on ad delivery: See the work of Alan Mislove or Arvind Narayanan.
- Lots of work on algorithms for pre-trial detention and other outcomes in the judicial system. Old-school AI.

LLMS IN THE WILD

- Brynjolfsson, Li, and Raymond (2025, QJE) - customer service agents, but pre-ChatGPT.
- Noy and Zhang (2023) - Giving people access to chatGPT makes them more productive.
- Dell'Acqua et al. (2023) - Consultants can use LLMs.
- van Inwegen (Wiles) et al. - Writing assistance when applying to jobs.
- Literature is particularly interested about inequality of benefits.

NEXT TIME: AI AS A RESEARCH TOOL

- Discuss assignment.
- Read Horton (2025), Park et al. (2023).
- Intros to Compiani et al. (2025) and Ludwig and Mullainathan (2024).