ARTIFICIAL INTELLIGENCE ALGORITHMIC RECOMMENDATION AND COMPETITION

Giacomo Calzolari
(European University Institute)

with E. Calvano, V. Denicolò, S. Pastorello

4th Cambridge Symposium on Competition Policy
March 10, 2025

The world we live

- Choice set for consumers is nowadays immense, mostly unknown, e.g. products, news, movies, songs, assets, posts to read, papers ...
  - products in Amazon Marketplace: 353m
  - songs on Spotify: 90m
  - movies on (US) Netflix catalog: 6000
  - videos on YouTube: 26bn
  - millions of news on Facebook
  - ...

Ocean of products

- We will never be able to fully explore this ocean of products,
  - **too many alternatives** and even if we knew they exist...
  - we would **not know our own tastes** for these products

- We need some help!

  What is the new sextant to navigate this ocean?

Recommender Systems

- Def. ***Recommender Systems (RS)*** *are software programs providing personalized recommendations to users/consumers about specific items/products.*

- They predict users' preferences by collecting information on users' valuations of items tried in past, also by other users: a **collaborative tool**
- With predicted match-value, RS make different items prominent to different users (*personalized prominence*)

Recommender Systems

RecSys Team at X

"*Our high-level goal is to make content discovery effortless and to free the user from the need for manual curation.
Personalized recommendations are essential to a wide range of technology-enabled products, and X is a prime example of this.
The primary aim is to facilitate effortless content discovery for users, thereby eliminating the requirement for manual curation.*"

Recommender Systems: applications

- products (Amazon and Google)
- music (Spotify)
- movies (Netflix, Amazon)
- videos (YouTube)
- socials (Facebook)
- apps (Google Apps)
- financial products (Betterment, robo-advisor)
- academic articles (Elsevier)
- referees (Elsevier)
- ...

Why do we care?

Recommender Systems (RS) are already shaping users' choices

- Recommended: Netflix movies 75%, Amazon views 35%, Spotify songs 40%, YouTube videos 60%

Worries and claims about algorithmic recommendations

- Heated policy debate, risks for competition and democracy
- DSA "*large online platforms [...] may need to mitigate the negative effects of personalised recommendations.*"
- Rich-get-richer: RS exasperate popularity
- Endogeneity/feedback-loop issue with AI in markets: (re)trained on data which contributed to generate
  $\rightarrow$ "bias in the data" vs. bias in algo

Recommender systems and markets

Buyers

- Uninformed buyers search for products and RS give personalized recommendations, inducing an "algorithmic-mediated demand":
  1. informing of unknown products
  2. giving prominence to certain products

Sellers

- Adapt their pricing strategies to algorithmic demand

Platforms

- Design their RS, e.g. best vs. manipulated recommendations

Recommender systems and economic modelling

- In typical search model, idiosyncratic preferences, for simplicity ...
- ... but then typical RS would be useless
- In reality, preferences and product characteristics show systematic similarities and differences
- and platforms observe feedback for only a tiny fraction of all consumer-product pairs, estimates subject to small-sample biases

- Integrating these key elements results in a highly complex framework that challenges the application of analytical methods.

Methodology to study AI and markets

Experimental (simulation) approach

- We operate *realistic* AI systems in synthetic and controlled environments
  - Synthetic = we generate key economic dimensions
  - Controlled = we control training data of the algo
- Challenges: (i) Algos must be similar to those used in markets, (ii) Economic environments must be realistic

Intended contribution

1. Methodological: use sound economic analysis with realistic AI algorithms
2. Specific: studying the implications of AI in applications (e.g. pricing algorithms, recommender systems)

Literature

- IO/theory literature of information to imperfectly informed consumers
  - Closest are Zhong (2022) and Zhou (2022): like us with differentiated products, heterogeneous consumers, and consumer search
  - They show providing pre-search information reduces prices
  - For tractability assume idiosyncratic preferences (RS useless here)
- Trade-off with more information, Armstrong and Zhou (2022), Anderson and Renault (2002): better match but also higher prices
  - But they abstract from both consumer heterogeneity and search
  - Adding these ingredients (but with idiosyncratic tastes) Zhong (2022) and Zhou (2022) find price reduction and no trade-off
- Manipulation of recommendations
  - Hagiu Jullien (2011), de Corniere Taylor (2019), Teh Wright (2020), Bourreau Gaudin (2022), Peitz Sobolev (2022), Bar-Isaac Shelegia (2023)
  - We confirm price-reduction of manipulation, but in a different model
- Methodology: realistically including all elements needed makes theory (so far) intractable
  - Similar methodology, Lee and Wright (2021) and Castellini, Fletcher, Ormosi, and Savani (2023)
  - Do not consider individual search benchmark to contrast with the RS
  - But do not study product market competition and effects on prices

Conceptual framework

Basic ingredients

- Set of items
- Set of users
- Dataset (Rating Matrix): contains observed *ratings* of items effectively consumed by some users, but just for **few** user/item combinations

Goals of a RS is to match items to users:

1. predict preferences and thus ratings for *unobserved* user-item combinations, i.e. those items that a given user did not consume yet
2. rank items for any user
3. provide *personalized* recommendations: show highest ranked items

Conceptual framework

The Rating Matrix R

- Rating matrix R (IxJ)
- I users & J items
- Some observed ratings r_{ij}
- R very large and very **sparse** (typically 1-5% non-blanks)

Items

Users		A	B	C	D
	1		4.5	2.0	
	2	4.0		3.5	
	3		5.0		2.0
	4		3.5	4.1	1.0

The task is:
predict missing ratings (matrix completion)
$\rightarrow$ make personalized recommendations

Environment & Recommendation Algorithm

Economy:

- We use a standard discrete choice model of differentiated products with systematic correlations across users/items: RS *possibly useful*
- We create J synthetic items with specific characteristics, controlling for horiz./vertical differentiation + central/niche products
- We create I users, each knowing own personal characteristics but not products' characteristics (incomplete information): thus searching (at cost) for individual best match-value item

Algorithm:

- From the Computer Science toolbox, we use Model-based Collaborative-Filtering RS: popular, handling huge rating matrices (matrix decomposition with users & products embeddings)
- We build Rating Matrix R with realistic properties: low density $\approx 1.2\%$
- Analytical properties of such algorithmic estimators in small samples inherently difficult/impossible to characterize

Model: Products and Preferences

A discrete choice model of differentiated-substitute products with systematic correlations across users/items: RS *possibly useful*

We create I synthetic users and J synthetic items:

- item-j characteristics $\beta_j = (\beta_{j1}, ...\beta_{jk})$
- user-i preferences $\theta_i = (\theta_{i1}, ..., \theta_{ik})$
- a random utility model $\tilde{u}_{ij} = \theta_i \cdot \beta_j + \tilde{\varepsilon}_{ij}$ with normal iid errors

We generate (θ, β) controlling for horizontal-vertical differentiation, and central/niche items, e.g. distributed uniformly over arc

Experimental protocol, data and information

- We build dataset R replicating key features in Netflix Challenge: I/J ($24000/800 \approx 30$), $\#obs/\#param (\approx 5)$, density ($\approx 1.2\%$)
- For each environment, we run 100 simulations

Two types of data:

- *Exogenous data*: R built with ratings of randomly chosen items per-user (rating reported into R with noise)
- *Endogenous data* to assess feedback-loop: initialize R with few random ratings per-user, make recommendations, that users follow reporting ratings (with noise), thus populating R till target density

Two worlds to compare, identifying effects of RS:

- *Individual search benchmark* with no RS (Anderson Renault Wolinsky)
- *RS environment*: each user receive a *personalized* recommendation (prominence), decides to follow if good, or start searching

Baseline environment

Part I: We initially assume:

- Subscription-based platforms (active items' sellers Part II)
- Platform always recommends best match (manipulation in Part III)
- Reported ratings are the realized utilities (robust.)

Results

Part I: Concentration and Superstars

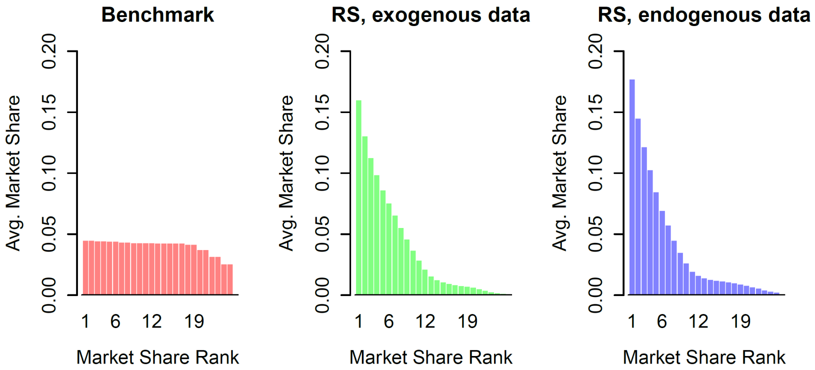

Figure: Items' market shares, ordered (≈horizontal differentiation, 25 items)

Part I: Product market concentration

% change RS vs. Benchmark	horizz.diff. ($\alpha = 0$)	interm.diff. ($\alpha = \frac{1}{2}$)	vertic.diff. ($\alpha = 1$)
Market-share of central products	20.61% (11.380%)	140.67% (2.196%)	72.30% (0.339%)
HHI	100.15% (1.822%)	214.90% (4.669%)	153.08% (0.989%)
Market-share peripheral products	−20.66% (2.806%)	−16.39% (3.449%)	−38.11% (0.561%)

Result (Concentration Bias) *The RS induces:*

1. *strong super-star effect of central products*
2. *strong increase of HHI*
3. *decrease of tail-peripheral products (no-long tail effect)*
4. *minor role of feedback loop*

Preferences and Products Estimations

Giving the RS the right model, it estimates $r_{ij} = \hat{\theta}_i \cdot \hat{\beta}_j$

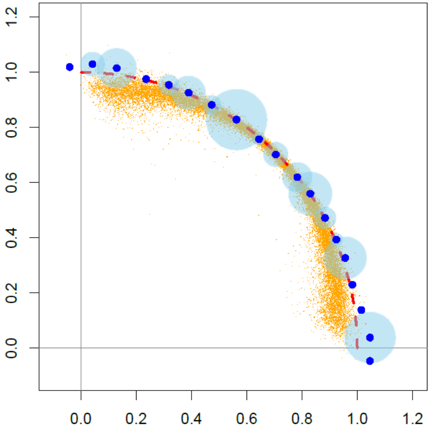

Is Concentration Bias an issue?

% change RS vs. Benchmark	horizz.diff. ($\alpha = 0$)	interm.diff. ($\alpha = \frac{1}{2}$)	vertic.diff. ($\alpha = 1$)
Users' surplus	$+1,21\%$ (0.01%)	$+2.72\%$ (0.01%)	$+5.82\%$ (0.01%)

- **Result (Consumer Surplus)**: *RS consistently increases users' surplus thanks to: (i) better user-item match and (ii) saving search costs.*

- Per-se the bias on concentration does not tell much: it could be weak/strong competition with opposite welfare implications
- How RS and its biases affect competition?

Part II (Amazon-like platform with active items' sellers)
We measure intensity of competition with the RS

- calculate Nash equilibrium prices with and without RS

Algorithmic demand

Now users also imperfectly informed about $(p_1, ..., p_J)$

Individual search benchmark

- Sellers anticipate users' search decisions and calculate the expected demand $D_j(p_j, p_{-j})$ (passive beliefs and no directed search)

With Recommender system

- for comparability, recommendations are independent of prices
- sellers anticipate recommendations and how they affect buyers' search and decisions: i.e. they rely on the **algorithmic demand**

Compare equilibrium prices when sellers use $D_j(p_j, p_{-j})$ or $D_j^{algo}(p_j, p_{-j})$

Equilibrium (weighted-)price %-change

	→ increasing vertical differentiation				
%. Diff. mean	+15.10	+13.01	+9.06	+7.39	+7.13
% Diff. sd	0.41	0.28	0.15	0.08	0.03

Results (Price Effect)

- *RS substantially increases prices*
- *Feedback loop (unreported) slightly reduces prices*

Consumer surplus %-change

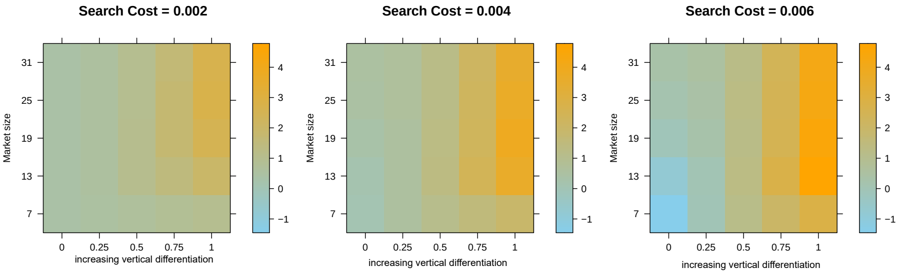

Results (Net Consumer Surplus)

- *Price reaction to RS eats most of match-value increase, and can even reduce consumer surplus (horiz. differentiation).*
- *Feedback loop (unreported) very similar.*

Anatomy of algorithmic demand

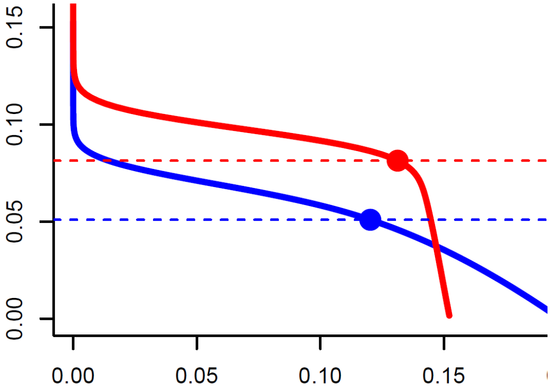

Demands: Red RS – Blue Benchmark

1. higher WTP users $\rightarrow$ **upward-vertical shift** ($p \uparrow$)
2. more similar users $\rightarrow$ **counter-clockwise rotation** (p?)
3. prominence to other items $\rightarrow$ **rectangularization** ($p \uparrow$)
4. bias in favor(against) superstars(niche) $\rightarrow$ **heterog. horiz. shift** (p?)

Part III: Manipulated recommendations

- If platform earns more from transactions on certain item(s)
- It may over-recommend it, constrained by users disregarding very poor recommendations
- Manipulation rate: favored item is recommended to x% of close but less ideal users

Results:

- **Result** *Manipulation reduces prices.*
  - favored seller receives more heterogeneous buyers and reduces price to retain mismatched users, other sellers respond (complements)
- Price reduction limits negative mismatch of manipulations for consumers and also the incentive to manipulate:
  profit maximizing manipulation rate $<< 100\%$
- Since disfavored competitors forced to reduce prices: manipulation is more an exclusionary abuse than exploitative

Part III: Manipulated recommendations

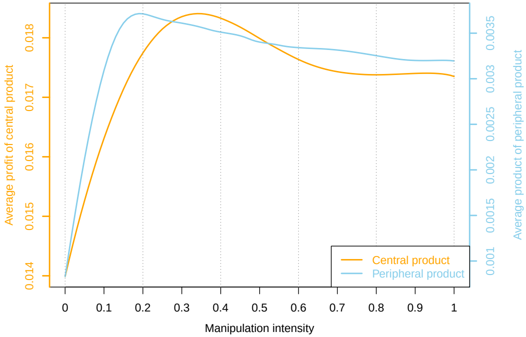

- Platform's optimal manipulation rate: % of users prompted with the favored rather than best-estimated product $\leq 32\%$
- Hence, the favored seller first gains, then loses: too much manipulation induces too intense price competition

Quality of recommendations

Quality of reccommendations depend on the RS' information:

- Less informative ratings: discretized Likert ratings, varying error to reported ratings
- Varying density and dimensions of rating matrix R
- Mis-specified model and algo deciding on dimensions
- Algo cross-validating its hyper-parameters
- Endogenous data

- Results qualitatively confirmed with these robustness
- But what is the role of different levels of RS' information on effects?

Quality of recommendations and welfare

Altering information available to RS and thus its quality?

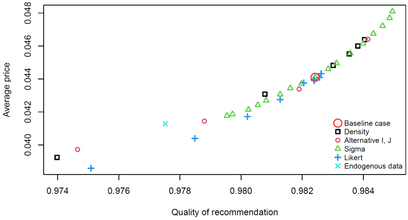

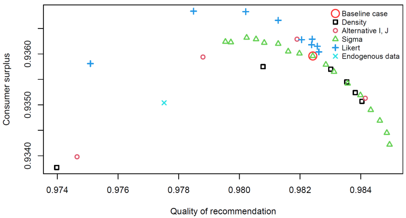

Each point is a different environment (varying density, I, J, shocks, Likert ratings, number of factors,...)

- More/better data to the RS → higher prices
- Inverted-U effects on consumers surplus: scope for some privacy

EXTENSIONS

- Multiple recommendations
- Price-directed search and price-mediated RS
- RS modifies market structure: entry/exit
- Competing Recommender systems
- Complementing RecSys with LLM, e.g. for cold start probl.

Key Findings and Policy

1. RS greatly increase market concentration (super star-effect)
2. 'feedback loop' second order: bias in the algorithm
3. RS anti-competitive for pricing, generally increasing consumer surplus
4. If platform manipulates, prices decrease, but CS decreases too
5. A notion of too-much-information for consumer surplus (privacy)

Although it does not look good, the picture is nuanced, and given the general positive impact on CS:

- no presumption of negative impact
- we need close monitoring

New stuff with AI and market analysis

Multi-product firms

- With recommender systems (ongoing)
- Selling with pricing algorithms: multi-markets contacts facilitate algorithmic price-collusion

Mergers

- AI and profitability of horizontal mergers
- AI searching for best merger deals

AI agentification: Strategic algorithmic buyers

- Durable good, heterogeneous buyers, strategically delayed purchases
- Monopolist' curse or Pacman surplus extraction?
- Colluding buyers protect against colluding sellers: but inefficiencies

Thank you

"AI, Algorithmic Recommendation and Competition"

Giacomo Calzolari