Competition, Alignment, and Equilibria in Digital Marketplaces

Meena Jagadeesan, Michael I. Jordan, and Nika Haghtalab (UC Berkeley)

AAAI 2023

Competition in a Digital Marketplace

Google vs bing | Spotify vs pandora | YouTube vs TikTok

This work: a theoretical framework to study how competition between *data-driven* platforms affects users

- *Motivation*: is perfect competition a suitable benchmark for a healthy *digital* marketplace?
- We show that competition need not fully align market outcomes with user utility.
- Misalignment occurs for *separate data* (status quo) and *shared data* (proposed in policy).

Impact on Platform's Learning Process

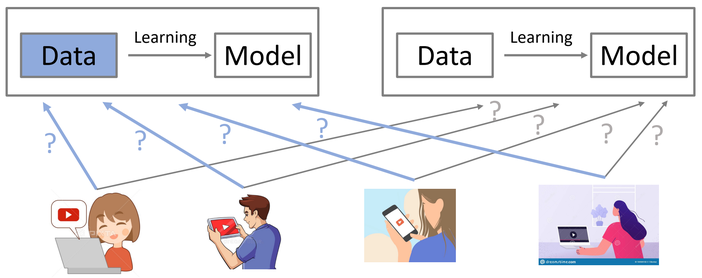

Data comes from users and depends on user choices.

Conventional wisdom about competition

Consider classical marketplaces for products.

- Users "vote with their feet" and choose their favorite product.
- A firm has to fully cater to user choices to retain their user base.

Conventional wisdom: Competition fully aligns market outcomes with user utility in *classical markets*.

How does competition impact users at equilibrium?

Main finding: Competition need not fully align outcomes with user utility in *digital markets*, regardless of data sharing specifics.

Theorem: If platforms have *separate data repositories (status quo)*, there are multiple equilibria whose alignment spans between the two benchmarks.

Weak alignment (optimal user utility in market w/ 1 user)

Full alignment (optimal user utility in market w/ all users)

Theorem: If platforms have a *shared data repository (proposed in policy)*, then there is a unique equilibria with alignment strictly in between the benchmarks.

Intuition for Alignment Results

Separate Data Repositories (Status Quo):

- A platform can make up for a subopt algorithm with more users.
- A platform retains its user base as long as their algorithm achieves at least the optimal utility in a market with 1 user.

Shared Data Repository (Proposed in Policy):

- A platform can't make up for a subopt algorithm with more users!
- But the optimal algorithm for a user when there are other users is *not* the cooperative optimal algorithm.
- Users wish to *free-ride* off the exploration of other users.

Our stylized model of a digital marketplace

Platform's learning task: multi-armed bandits

Content	Arm 1	Arm 2	…	Arm k
Unknown quality level	$q_1 \sim D_1$	$q_2 \sim D_2$	…	$q_k \sim D_k$

Platform action = **choose an algorithm from a predefined class** of general bandit algorithms that may include:

- Greedy and mixtures with uniform exploration
- Thompson sampling and mixtures with exploration
- Optimal algorithm for a given discount factor

Data sharing: platforms may have *separate data repositories for their own observations* or a *shared data repository with observations from both platforms*.

Formalizing how platforms and users interact

Stages of the Stackelberg game:
Stage 1: Each platform commits to an algorithm.
Stage 2: Each user chooses between platforms.

Participant actions and equilibrium concepts

User utility = discounted cumulative quality of recommended arms by chosen platform
Users arrive at a **Nash equilibrium.**

YouTube TikTok

Platform utility = number of users
Platforms arrive at a **Nash equilibrium**.

Measuring alignment with user utility

Definition (User Quality Level): The *user quality level* of a pair of platform algorithms is the *user utility* achieved at a Nash equilibrium for users, given those algorithms.

Our focus: user quality level of *equilibrium* algorithms