Recommendation Systems and Product Diversity in Online Platforms

Chinmay Lohani and Jaume Vives-i-Bastida

U. Penn and MIT

Goal: study empirically and theoretically how recommendation systems affect content creation in platforms.

Today:

1. Motivate the problem.
2. Suggestive evidence that content variety is decreasing.
3. Highlight a channel with a toy model.

Motivation

- **A lot at stake**: Online content platforms are big markets.
  - **YouTube**: 2 billion monthly active users, 30 million paid subscribers, 37 million channels, $28 billion yearly revenue.
  - **Spotify**: 260 million active users, 160 million paid subscribers, 3 million artists, $10 billion yearly revenue.
- **Recommendation systems** are at the core of these platforms:
  - 70% of YouTube views and 75% of Netflix views come from recommendations.
  - Tiktok's main feature does not even allow consumers to choose.
  - One 2019 vendor survey: 31% of the revenues in the global e-commerce industry.
- **Concerns**:
  - **Consumer side**: content diversity has been decreasing over the years.
  - **Supply side**: artists protest about unfair compensation in streaming platforms.

Motivation: in the news I

The New York Times

ON TECH

Streaming Saved Music. Artists Hate It.

Many musicians aren't sharing in streaming riches. Can digital music economics change to benefit everyone?

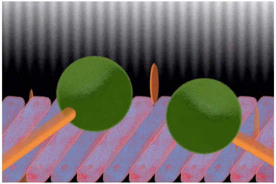

By Alexis Jamet

Pop music these days: it all sounds the same, survey reveals

Pop music is too loud and melodies have become more similar, according to a study of songs from the past 50 years conducted by Spanish scientists

▲ A sea of homogeneity? ... revellers at the park stage at Glastonbury 2011. Photograph: Adrian Dennis/AFP/Getty Images

Motivation: in the news II

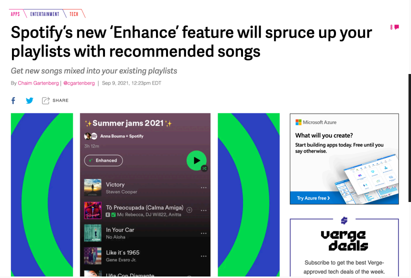

Questions

Questions we want to explore:

1. How do recommendation systems affect content diversity and consumer diversity in platforms? [Today]
2. What happens when recommendation systems become more accurate? [Today]
    - Who loses?
    - Who wins?
    - Exploration vs. Exploitation.
3. Markups/Market power: do recommendation systems create super stars?

Related literature

- **Recommendation systems:** Hosanagar et al. 2014, Aguiar and Waldfogel 2018, Dinerstein et al. 2018, Chen et al 2019, Yeomens et al. 2019.
- **Multisided markets:** Rochet and Tirole 2003, Farrell and Klemperer 2007, Weyl 2010, Tan and Zhou 2020.
- **Platforms and Data:** Reisinger et al. 2009, Nosko and Tadelis 2015, Acemoglu et al. 2020, Ichihashi 2020, Bergemann et al. 2021, Cao et al. 2021, Johnson et al. 2021.
- **Consumer search:** Varian 1980, Diamond 1981, Ellison and Ellison 2004, Athey and Ellison 2011, Ellison and Wolitzky 2012, Blake et al. 2016.

Suggestive Evidence

Anecdotal Evidence

- **More of the same**: as recommendation systems get better platform content becomes less diverse. [Today]
    - Randomised trials by Spotify: personalised recommendations lower consumption diversity.
- **More revenue/usage**: recommendation systems that do more exploitation than exploration lead to more revenue/usage.
    - Randomised trials by Spotify: personalised recommendations increase sales.
    - YT music executive: exploitation increases revenue.
- **More markups**: as content becomes more concentrated the top content producers bargain for higher fees per view.

Spotify API: limited access to user and song level data:

- **Personalization API**: get user - date recommendations based on *affinity* metric (i.e. recommendation system).
- **Playlists API**: get user - date playlists that users make.
- **Tracks API**: get songs data with technical information.

Data: 0.5 M songs sampled from the Spotify library in 2021.

energy	liveness	acousticness	loudness	valence	tempo	time_signature	duration_ms	year	popularity	name
0.592	0.3830	0.0467	-6.738	0.4420	140.038	4	166286	2021	57	EXTENDO
0.933	0.7650	0.1140	-6.476	0.4420	137.915	4	148447	2021	44	Marek Hamšík
0.719	0.0938	0.0170	-5.972	0.3580	169.939	4	170667	2021	57	Rich
0.145	0.1640	0.9460	-23.367	0.0395	113.445	3	212851	2021	1	At Sunset
0.615	0.3050	0.2060	-6.212	0.4380	90.029	4	142003	2021	58	A Day At A Time

Figure 3: Sample datum.

Spotify data: similarity over time

More of the same: average cosine similarity increases by release year.

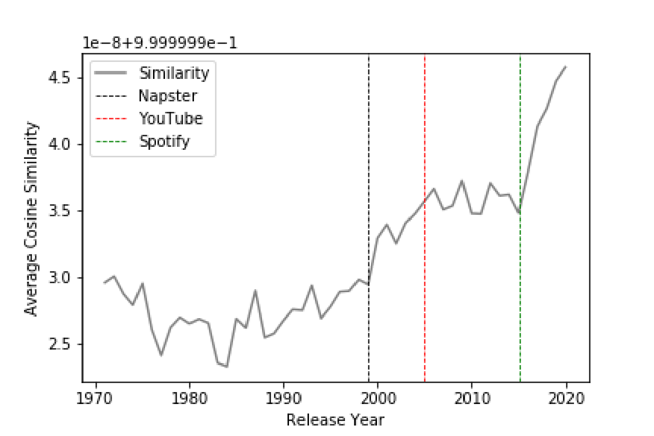

Figure 4: Cosine similarity trend.

- Variance of the music features is decreasing.
- Songs nowadays are **louder**, more **energetic** and have higher **tempo** and **time signature**.

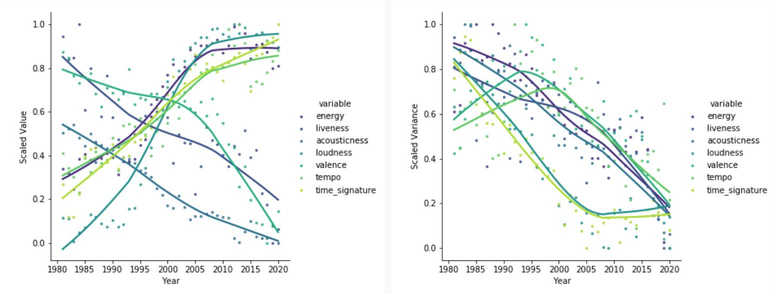

Figure 5: Music features trends and variance.

- Positive relationship between similarity and popularity.
- Very popular songs are close to the mediod.

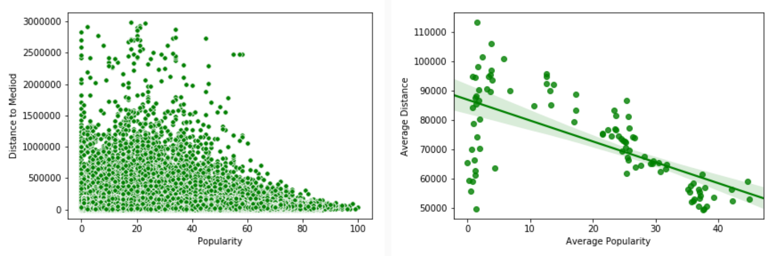

Figure 6: Distance from mediod and popularity.

Desired features of a model

Today:

- **Two sided platform** where consumers and content creators are matched.
- **Content space** where the match utility depends on the distances between a producers and consumer $\Longrightarrow$ **order**.
- A **recommendation system**'s goal is to serve consumers according to their preferences.
- **Channel**: **Screening through prices**, the platform sets fees for consumers and pays producers optimally to maximise profit taken the recommendation system as given.

Not Today:

- **Data externalities**: More consumers help make better recommendations.
- **Market power**: more popular producers can bargain for higher prices.
- **Dynamic**: exploration as a way to learn preferences.

Simple set up:

- **Content Space**: **2 consumer masses** and **N producers** are located in a content space χ.
- **Platform**: brings together the consumer and a set $\mathcal{J}$ of producers. Charges the consumer p^B and each producer p^S.

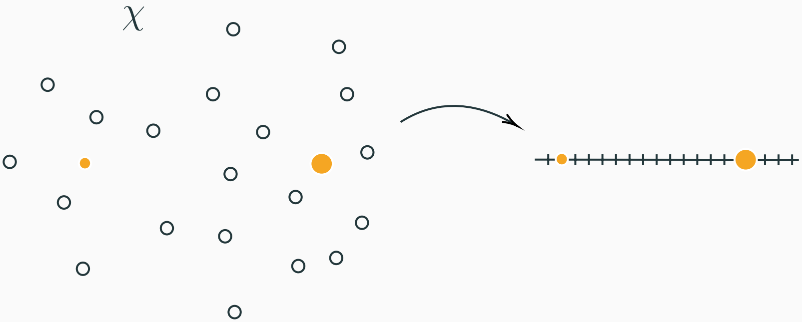

Toy Model II

- **Ordering**: consumer x has a value distribution $G(x, y)$, over producers that is induced by the distance metric in χ.
- **Consumers**: each consumer has unit demand and is offered a bundle over producers according to f during a free period. Then decides whether to pay the p^B. His value of joining is given by:

$$V^B(x) = \sum_{j \in \mathcal{J}} g(x, y_j) f(x, y_j) - s.$$

- $\alpha-$**recommendation system**: The probability that the consumer is offered a producer from a set of $\mathcal{J}$ producers is:

$$f(x, y, \mathcal{J}) = \alpha \frac{g(x, y)}{\sum_{z \in \mathcal{J}} g(x, z)} + (1 - \alpha) \frac{1}{|\mathcal{J}|}.$$

- **Producers**: outside option of joining c. They decide to join *before* the free period with knowledge of f and the (per unit) price p^S they will receive. For now we abstract away from beliefs on others and equilibrium concepts.

Toy Model III

The **platform** problem is:

$$\max_{p^B, p^S} \sum_{i \in \mathcal{P}} W_i p^B - \sum_{j \in \mathcal{J}} \sum_{i \in \mathcal{P}} W_i f(x_i, y_j) p^S \quad \text{s.t.}$$

$$\text{(Producer PC)} \quad \mathcal{J} = \{j \in [N] | \sum_{i \in \mathcal{P}} W_i f(x_i, y_j) p^S \geq c\},$$

$$\text{(Consumer PC)} \quad \mathcal{P} = \{i \in \{1, 2\} | \sum_{j \in \mathcal{J}} g(x_i, y_j) f(x, y_j) - s \geq p^B\}.$$

- The platform knows g, but it can't control f.
- W_i is the weight on each consumer type mass ($W_1 = 1, W_2 = W \geq 1$).
- The idea is that α is given by technological constraints rather than platform optimization.

Results Summary I

Case: **no recommendation system** ($\alpha = 0$)

- If the platform is profitable, then all producers will choose to participate and $\mathcal{J} = [N]$.
- Producer prices are high: $p^S = cN/(1+W)$.
- Consumer prices are $p^B = \overline{G} - s$, where $\overline{G} = \frac{\sum_{j \in [N]} g(x_i, y_j)}{N}$.
- Profits are low: $\pi = (1+W)(\overline{G} - s) - cN$, if $c \leq \frac{1+W}{N}(\overline{G} - s)$.

Case: **perfect recommendation system** ($\alpha = 1, N = 2$)

- Each of the consumer types x_1, x_2 has a most preferred producer type y_1 and y_2 respectively.
- There exists a cutoff weight W^*. Below W^*, the platform will serve both consumer types, above W^* it will only serve type 2.
- When $N \geq 2$ cutoff exists, but we need additional conditions to determine who will be in the market below the cutoff.

Results Summary II

Case: **perfect recommendation system** ($\alpha = 1$, $N = 2$)

- For $W \geq W^*$:
  - Only serve consumer 2 and producer y_2.
  - High profits: $\pi_{exclusive} = W(g(x_2, y_2) - s) - c$.
- For $W < W^*$:
  - Serve both consumers and producers y_1 and y_2.
  - $p^S = \frac{c}{\min_{j \in \{1,2\}} \sum_{i \in \{1,2\}} W_i f(x_i, y_j)}$.
  - $p^B = \min_{i \in \{1,2\}} \sum_{j \in \{1,2\}} f(x_i, y_j) g(x_i, y_j) - s$.
  - Lower profits: $\pi_2 = (W+1)p^B - c(1+\gamma)$, $\gamma > 1$, where $\gamma - 1$ is the positive profit made by producer y_2.

Case: **imperfect recommendation system** ($0 < \alpha < 1$, $N = 2$)

- Uniform distribution does no change the order: same cutoff structure as in perfect case.
- When $W \geq W^*$ same as before.
- Exploitation vs. Exploration: When $W \leq W^*$ then p^B and p^S are both lower than in the perfect case.

- **Main takeaway**: recommendation system strength leads the platform to include less content producers.

$\alpha = 0$

All producers included

Low profits

Low p^B, high p^S

$\alpha = 1$

Only close producers included

High profits

High p^B, low p^S

Trade off lower p^S and p^B

Future Work

- **Empirical**:
  - Get **user** level data and study the effect of a change in recommendation system (structural model).
  - Get **producer** level data and study the effect on entry and markups (YT data on price per view by content category).
- **Theoretical**:
  - General framework for the problem.
  - Add data externalities, different producer prices/outside options, consumer search.

Case: No recommendation system

- When $\alpha = 0$ the consumer is equally likely to consume from any producer: $f(x, y_j) = 1/N$.
- Participation constraint of producers can be satisfied with equality: $p^S = cN/(1+W)$.
- If the platform is profitable, then all producers will choose to participate and $\mathcal{J} = [N]$.
- The platform maximizes profits by setting $p^B = \overline{G} - s$, where $\overline{G} = \frac{\sum_{j \in [N]} g(x_i, y_j)}{N}$: utility of the average bundle.
- $\pi = (1+W)(\overline{G} - s) - cN$,

Profit is positive if $c \leq \frac{1+W}{N}(\overline{G} - s)$.

Consider the case with N=2.

Each of the consumer types x_1, x_2 has a most preferred producer type y_1 and y_2 respectively.

There exists a cutoff weight W^*. Below W^*, the platform will serve both consumer types, above W^* it will only serve type 2.

Case: Perfect recommendation system II

- Case: $1 \leq W \leq W^*$- Platform serves both consumer types-
  - $p^S : min_{j\in\{1,2\}} \sum_{i\in\{1,2\}} W_i f(x_i, y_j)p^S = c$. The platform sets p^S high enough so that the least profitable producer is indifferent. When distances are symmetric, this binds for type 1.
  - $p^B : min_{i\in\{1,2\}} \sum_{j\in\{1,2\}} f(x_i, y_j)g(x_i, y_j) - s = p^B$. The platform sets price p^B low enough so that the least utility consumer type is indifferent. When distances are symmetric, both consumer types have the same utility.
  - Platform profits are $\pi_{all} = (W+1)p^B - c(1+\gamma)$, $\gamma > 1$, where $\gamma - 1$ is the positive profit made by the more profitable producer. For the symmetric case, $\gamma = 1 + \frac{c(W-1)(D-1)}{(W+D)}, \quad D := \frac{g(x_2,y_2)}{g(x_2,y_1)}$.

Case: Perfect recommendation system III

- Case: $W > W^*$- Platform serves consumer type 2 (one with higher mass)-
  - $p^{S*} : Wf(x_2, y_2)p^S = c$ Platform sets a price p^{S*} low enough so that only producer y_2 will be able to stay on the platform.
  - $p^{B*} : g(x_2, y_2) - s = p^B$ Platform sets price p^{B*} to extract all the surplus from consumer 2.
  - Platform profits are $\pi_{exclusive} = Wp^{B*} - c$
- When N>2: There still exists a cutoff W^*, above which the platform screens out all producers other than $max_{j\in[N]} g(x_2, y_j)$. We need regularity conditions on the utility function $g(x, y)$ to determine the producers and consumers who will be present below the cutoff.

When $\alpha \in (0, 1)$- Similar equilibria can be maintained

- $W > W^*$
  - If $\alpha > 0$ the producer y_2 still has a strictly lower reservation price than other producers, i.e., they are willing to be present a lower price than other producers
  - Given this α, however small, the platform can screen out other producers by setting a very low price $p^S = p^{*S}$. The platform will choose to do this when $W > W^*$.
  - Similarly, given that only y_1 enters, the consumer will face the same price as before p^{B*}.
  - This is the case because the platform knows g so it can screen out the sellers if the recommender system is a little bit informative.

Case: In between II

- $1 \leq W \leq W^*$- Platform serves both consumer types-
  - $p^S : min_{j \in \{1,2\}} \sum_{i \in \{1,2\}} W_i(\alpha \frac{g(x,y)}{\sum_{z \in \mathcal{J}} g(x,z)}) + (1-\alpha)\frac{1}{\mathcal{J}})p^S = c$. The platform sets p^S high enough so that the least profitable producer is indifferent.
  - $p^B : min_{i \in \{1,2\}} \sum_{j \in \{1,2\}} (\alpha \frac{g(x,y)}{\sum_{z \in \mathcal{J}} g(x,z)} + (1-\alpha)\frac{1}{\mathcal{J}})g(x_i, y_j) - s = p^B$. The platform sets price p^B low enough so that the least utility consumer type is indifferent.
  - An equilibrium with the same set of producers who were present when $\alpha = 1$ can be maintained here.