Playing to the Algorithm: How Spotify's Recommendations Affect Music Production

Max E. Schnidman

October 2024

Department of Economics
University of Virginia

Question and Motivation

- **How have recommender systems affected the characteristics of songs that rightsholders choose to release?**

- **Motivation**:
  - Recommender systems have transformed digital markets (Amazon, TikTok, Spotify)
  - Streaming platforms that use them have taken over the music industry: 84% of music revenue comes from streaming platforms (RIAA 2023); growing antitrust concern about algorithms
  - Cultural commentary about the effects of algorithms on music production (shorter, more homogeneous songs), and industry insiders assert that they have changed their music to fit them (Singer and Rosenblatt 2023)

This Paper

- Structural model with three sets of agents:
    - Rightsholders (forward-looking): choose whether to release a song
    - Spotify (single platform): determines the probability of recommending a song
    - Consumers: choose whether to listen or skip songs in a streaming session
- Entry condition: rightsholders release up to the breakeven point, so expected revenue from the worst-performing (marginal) song equals the fixed cost
- Estimate dynamic model using data from Spotify
- Run counterfactuals on the effects of recommender systems on song characteristics

Preview of Results

- Structural Model Estimates:
    - Difference between the consumer and recommender system preferences estimates, and producers target a combination of the two
    - Estimate an average fixed cost of $80,000, in line with industry estimates and academic research (Aguiar and Waldfogel 2018)
- Counterfactual Analyses:
    - Randomized recommendations: fewer songs get released (have negative expected profit), songs are on average 8 seconds longer and more heterogeneous, and consumer surplus is 4% lower
    - Popularity-based recommendations: significantly reduces the number of songs released, consumer surplus is 15.7% lower, but the released songs are more homogeneous and profitable

Related Literature and Contribution

- **Economics of Music**: Mortimer et al. (2012), Berry, Eizenberg, and Waldfogel (2016), Aguiar and Waldfogel (2018), Aguiar, Waldfogel, and Waldfogel (2021),
  - **Contribution**: A structural model of the music streaming industry with recommender systems
- **Recommender Systems**: Calvano et al. (2020), Melchiorre et al. (2021), Aridor and Gonçalves (2022), Bourreau and Gaudin (2022)
  - **Contribution**: Embedding product recommendations in an economic model with endogenous entry
- **Platforms and Intermediation**: Goeree (2008), Lee (2013), Aguiar and Waldfogel (2021), Anderson and Bedre-Defolie (2022), Reimers and Waldfogel (2023)
  - **Contribution**: Empirical application of product-based recommender systems on a digital platform

Roadmap

- Industry and Data
- Model
    - Consumers
    - Recommender System
    - Producers (Rightsholders)
- Estimation
- Results
- Counterfactual Analyses
- Conclusion

Industry and Data

Sales in the Recorded Music Industry

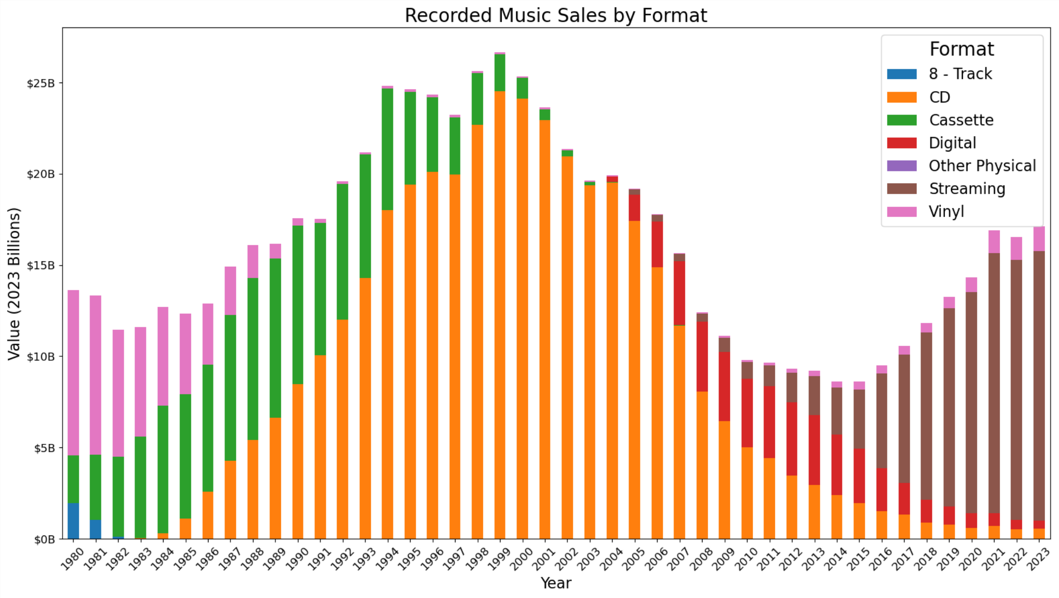

Agents in the Music Industry

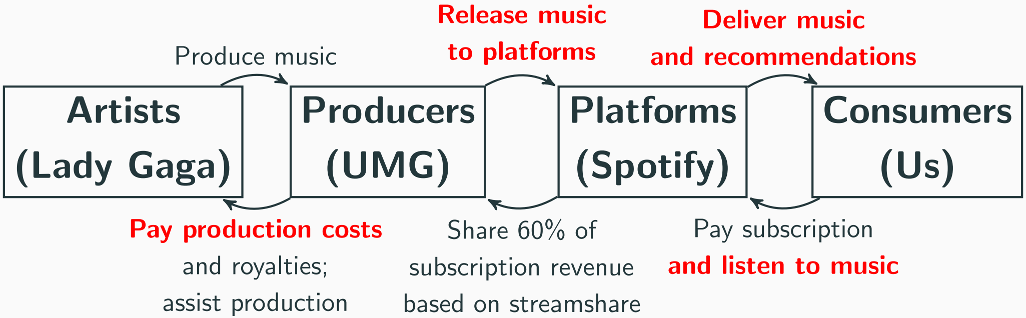

▸ Concentration Data

Spotify Data

1. Music Streaming Session Dataset (MSSD) (Brost, Mehrotra, and Jehan 2018)
    - 160mn consumer streaming sessions, 10-20 songs each, from July-September 2018
    - Song and consumer characteristics, and listen/skip point for each song (binned)
    - How consumers accessed the song (e.g., library, search, algorithmic playlist, etc.)
    - Use approximately 10 percent of the data, stratified by song, to estimate demand and recommender system parameters
2. Charts
    - Daily top 200 songs on Spotify from 2017-2021 with streaming counts
    - Use to estimate supply and entry, combined with demand and recommender system parameters

▸ Song Characteristics ▸ Sampling Strategy

Variation over Time

▸ OLS Motivation ▸ Time Use ▸ Demand Robustness Check ▸ Additional Characteristics

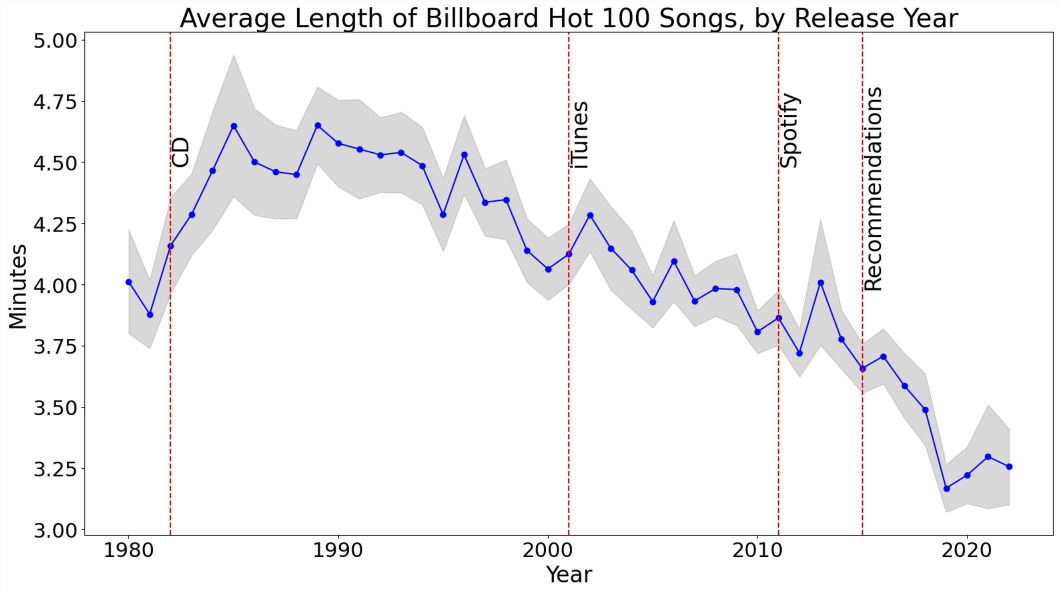

Song Descriptive Statistics - MSSD ($N = 3.7mn$ songs)

	Mean	Median	Standard Deviation	Min	Max
Duration (seconds)	233.19	217.91	108.40	30.00	1800.00
Loudness (dB LUFS)	-9.60	-8.08	5.73	-60.00	6.28
Tempo (BPM)	120.07	119.95	30.43	0.00	249.99
Release Year	2009	2013	11	1950	2019
Acousticness	0.35	0.22	0.34	0.00	1.00
Danceability	0.56	0.57	0.19	0.00	1.00
Energy	0.59	0.63	0.26	0.00	1.00
Instrumentalness	0.21	0.00	0.34	0.00	1.00
Liveness	0.21	0.13	0.19	0.00	1.00
Speechiness	0.10	0.05	0.14	0.00	0.97
Valence	0.48	0.47	0.27	0.00	1.00

▸ Spotify Charts Data ▸ Song Lifecycle ▸ Characteristic Correlations

Consumer Descriptive Statistics - MSSD ($N = 10mn$ consumers)

	Mean	Standard Deviation
Session Length	18.07	6.91
% Premium Subscribers	0.84	0.31
% RBS	0.58	0.31
% Completion	0.34	0.31
% Morning Listen	0.24	0.31
% Afternoon Listen	0.39	0.31
% Evening Listen	0.29	0.31
% Night Listen	0.08	0.27
% Algorithmic Playlist Listen	0.03	0.16
% Algorithmic Radio Listen	0.15	0.35
% Catalog Listen	0.24	0.43
% Editorial Playlist Listen	0.15	0.35
% User Collection Listen	0.42	0.49

Model and Estimation

Structural Model Overview

Producers make release decisions paying fixed cost F_j

↓

Spotify's recommender system determines the probability of recommending song j

↓

Consumers choose to listen or skip each song j they receive

- Consumer demand: random utility model
- Recommender system: logistic regression
- Supply: entry and fixed costs
    - Firms release a song based if expected profit is positive
    - Expected revenue of the marginal song equals the fixed cost
- Solution Concept: Oblivious Equilibrium

Consumer Microfoundation and Choice Structure

- Consumers on Spotify listen to a recommended playlist, their own playlist, or a song they have in mind
- **Maintained Assumption**: consumers do not consider how their listens affect the recommender system
- To count as a listen, a consumer must listen to at least **30 seconds** of a song
- Choice structure:

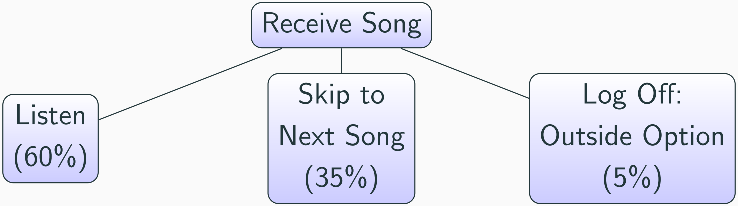

Utilities

- Consumers i have the following utility over listening to a song j:

$$U_{L,ijs} = \beta X_j + \gamma Y_i + \zeta_s + \epsilon_{ijs} \tag{1}$$

  - X_j: vector of linear and quadratic song characteristics
  - Y_i: vector of consumer characteristics
  - ζ_s: session position fixed effect
  - ϵ_{ijs}: Gumbel-distributed (Type 1 Extreme Value) error term

- Consumers have adaptive expectations of skipping, based on previous songs:

$$U_{S,ijs} = \beta E_{is}[X_j | X_{j,s-1}] + \gamma Y_i + \zeta_s + \epsilon_{ijs} \tag{2}$$

- Normalize the utility of the outside option to zero:

$$U_{i0s} = \epsilon_{i0s} \tag{3}$$

▸ Expected Skip Characteristics

Demand Estimation

- Choice probability (estimating equation), conditional on receiving song j:

$$
\begin{aligned}
& P(i \text{ listens to } j | i \text{ receives } j) \\
&= \frac{\exp(\beta X_j + \gamma Y_i + \zeta_s)}{1 + (\exp(\beta X_j + \gamma Y_i + \zeta_s) + \exp(\beta E_{is}[X_j | X_{j,s-1}] + \gamma Y_i + \zeta_s))}
\end{aligned} \tag{4}
$$

- Maximum Likelihood Estimation over the song-consumer interactions in the MSSD

- Identification comes from the variation in song characteristics and consumer characteristics for each song-consumer interaction

Recommender System - Model and Estimation

- Following Spotify research, I estimate the recommender system using logistic regression:

$$P(i \text{ receives } j) = \frac{\exp(\eta_1 X_{1j} + \eta_2 \prod_{n=2}^{N} X_{nj} + \eta_3 Y_i)}{1 + \exp(\eta_1 X_{1j} + \eta_2 \prod_{n=2}^{N} X_{nj} + \eta_3 Y_i)} \tag{5}$$

- I estimate the probability that consumer i receives song j as the probability consumer i completes song j
- The outcome variable is whether a consumer **completes** a song
- Estimation:
  - Maximum Likelihood Estimation over song-consumer interactions in the MSSD
  - Estimated with and without interaction terms

Recommender Systems Generally

Combining Song Demand and the Recommender System

- The probability a consumer i listens to song j at time t is:

$$
\begin{aligned}
P(i \text{ listens to } j) &= P(i \text{ receives } j) \times P(i \text{ listens to } j | i \text{ receives } j) \\
&= \frac{\exp(\eta_1 X_{1j} + \eta_2 \prod_{n=2}^{N} X_{nj} + \eta_3 Y_i)}{1 + \exp(\eta_1 X_{1j} + \eta_2 \prod_{n=2}^{N} X_{nj} + \eta_3 Y_i)} \\
&\times \frac{\exp(\beta X_j + \gamma Y_i + \zeta_s)}{1 + (\exp(\beta X_j + \gamma Y_i + \zeta_s) + \exp(\beta E_{is}[X_j | X_{j,s-1}] + \gamma Y_i + \zeta_s))}
\end{aligned}
\tag{6}
$$

- Producers face this unconditional demand

Microfoundation - Producers

- Each period (day), a producer has a potential song in its catalog (e.g., through a contract with an artist)
- Producers consider both the characteristics of rival songs, and the recommender system's intermediation of music demand in their profit maximization problem
- **Maintained Assumptions**: release timing is exogenous, and each song has an independent producer (i.e., no multiproduct firms)
- Once a song has been released, it is always available on the platform

Modelling Forward-Looking Behavior

- Producers do not know what songs will be available next period, so they form rational expectations over two variables:
  1. The characteristics of songs in the market $\mathcal{X}_t$
  2. The probability of being recommended ϕ_{jt}
- Define $\mathcal{X}$ as the vector of mean characteristics of all songs in the market
- Propose first-order Markov processes for the evolution of mean song characteristics and the recommender system

Markov Process Estimation

- Song Characteristics:

$$\mathcal{X}_{t+1} = \nu_0 + \nu_1 \mathcal{X}_t + \epsilon_t^{\mathcal{X}} \tag{7}$$

  - Estimate VAR on daily mean song characteristics in market-level Spotify Charts data
  - Also estimated independent AR(1) models for each characteristic

- Recommender System:

$$\phi_{j,t+1} = \psi_0 + \psi_1 \phi_{jt} + \epsilon_{jt}^{\phi} \tag{8}$$

  - Compute daily mean recommendation probability in market-level Spotify Charts data, using recommender system estimates
  - Estimate an AR(1) regression on these probabilities

Solution Concept: Oblivious Equilibrium (Weintraub, Benkard, and Van Roy 2005)

- Each player chooses a strategy that maximizes their expected payoff, based only on its own state and the long-run industry state
- Asymptotically approximates a Markov Perfect Nash Equilibrium.
- In recorded music:
    - Labels release thousands of songs daily, making it unlikely they consider every rival's release strategy
    - Labels also cannot know exactly how Spotify's recommender system will evolve

Profit Function

- Producers decide whether to release a song based on its expected profit:

$$E[\pi_j(X_j)] = 0.6 \left(\sum_{t=0}^{3} \delta^t R_t \left(\frac{P(i \text{ listens to } j \text{ with characteristics } X)}{\sum_K P(i \text{ listens to } k \text{ with characteristics } \mathcal{X})} \right) \right) - F_j \tag{9}$$

  - Fixed cost of producing F_j
  - R_t: (exogenous) total revenue generated by Spotify
  - δ: discount factor

- The fraction in the profit function is the estimated streamshare of song j, which is how Spotify allocates revenue to producers

Estimating Expected Revenue

- Expected revenue is a function of the demand model, the recommender system model, the Markov processes, and data

$$\text{Expected Revenue} = 0.6\left(\sum_{t=0}^{T}\delta^t R_t\left(\frac{P(i \text{ listens to } j \text{ with characteristics } X)}{\sum_K P(i \text{ listens to } k \text{ with characteristics } \mathcal{X})}\right)\right)$$

- I compute the expected revenue for songs that released in the first three quarters of 2018 and entered the Top 200

- Calibrate $\delta = 0.9956$ and $T = 1,095$ (three years)

▸ Streamshare

Entry Condition

- Producers will release if the following condition holds:

$$0.6\left(\sum_{t=0}^{3}\delta^{t}R_{t}\left(\frac{P(i\text{ listens to }j\text{ with characteristics }X)}{\sum_{K}P(i\text{ listens to }k\text{ with characteristics }\mathcal{X})}\right)\right)\geqslant F_{j} \tag{10}$$

- Producers release up to the point of indifference, so the fixed cost is equal to the expected revenue from the worst-performing song
- I apply this entry condition to each day in the first three quarters of 2018 to compute fixed costs for each day

▸ Song Network (External Validity)

Results

Estimates - Selected Song Characteristics

	Royalty-Bearing Stream (Demand)		Song Completion (Recommender System)	
	Estimate	Odds Ratio	Estimate	Odds Ratio
Age	-0.012***	0.988	-0.044***	0.957
	(0.0005)		(0.002)	
Duration	-0.009***	0.763	-0.222***	0.801
	(0.0008)		(0.0003)	
Duration^2	-0.0006***	–		
	(0.00009)			
Mode	0.027***	1.028	0.016***	1.016
	(0.001)		(0.0003)	
Tempo	0.0007	1.012	-0.011***	0.990
	(0.0009)		(0.0003)	
Tempo^2	-0.018***	–		
	(0.0005)			
Time Signature	0.045***	1.046	0.035***	1.036
	(0.003)		(0.001)	
Observations	148,822,923		180,061,351	
$\bar{\rho}^2$ / Pseudo R^2	0.0002		0.006	

Estimates - Selected Song Characteristics (Continued)

	Royalty-Bearing Stream (Demand)		Song Completion (Recommender System)	
	Estimate	Odds Ratio	Estimate	Odds Ratio
Acousticness	0.060***	0.994	0.043***	1.044
	(0.000639)		(0.0002)	
Acousticness^2	-0.046***	–		
	(0.00011)			
Energy	-0.008***	1.100	0.020***	1.021
	(0.001)		(0.0004)	
Energy^2	0.046***	–		
	(0.000398)			
Instrumentalness	0.0002	0.990	0.056***	1.057
	(0.0005)		(0.0002)	
Instrumentalness^2	-0.010***	–		
	(0.00005)			
Valence	0.025***	0.889	0.036***	1.036
	(0.0007)		(0.0002)	
Valence^2	-0.074***	–		
	(0.0002)			
Observations	148,822,923		180,061,351	
$\bar{\rho}^2$ / Pseudo R^2	0.0002		0.006	

Estimates - Selected Consumer Characteristics

	Royalty-Bearing Stream (Demand)		Song Completion (Recommender System)	
	Estimate	Odds Ratio	Estimate	Odds Ratio
Morning	0.133***	1.142	0.154***	1.167
	(0.0005)		(0.0004)	
Afternoon	0.073***	1.076	0.097***	1.102
	(0.0004)		(0.0004)	
Night	0.124***	1.132	0.171***	1.187
	(0.0007)		(0.001)	
Premium	0.016***	1.016	-0.065***	0.937
	(0.0004)		(0.0004)	
Observations	148,822,923		180,061,351	
$\bar{\rho}^2$ / Pseudo R^2	0.0002		0.006	

Discussion of Estimates

- Song Characteristics:
  - Gap between the demand and recommender system estimates, especially for Acousticness, Instrumentalness, and Valence
  - Recommender system is likelier to surface happier, more acoustic, and more instrumental songs
- Consumer Characteristics:
  - General alignment between the demand and recommender system estimates
  - Exception: subscribers more likely to provide RBS, but less likely to complete songs
- Markov processes for song characteristics and the recommender system exhibit statistical significance and stationarity

▸ Markov Process Estimates

Expected Revenues

Alternate Assumption

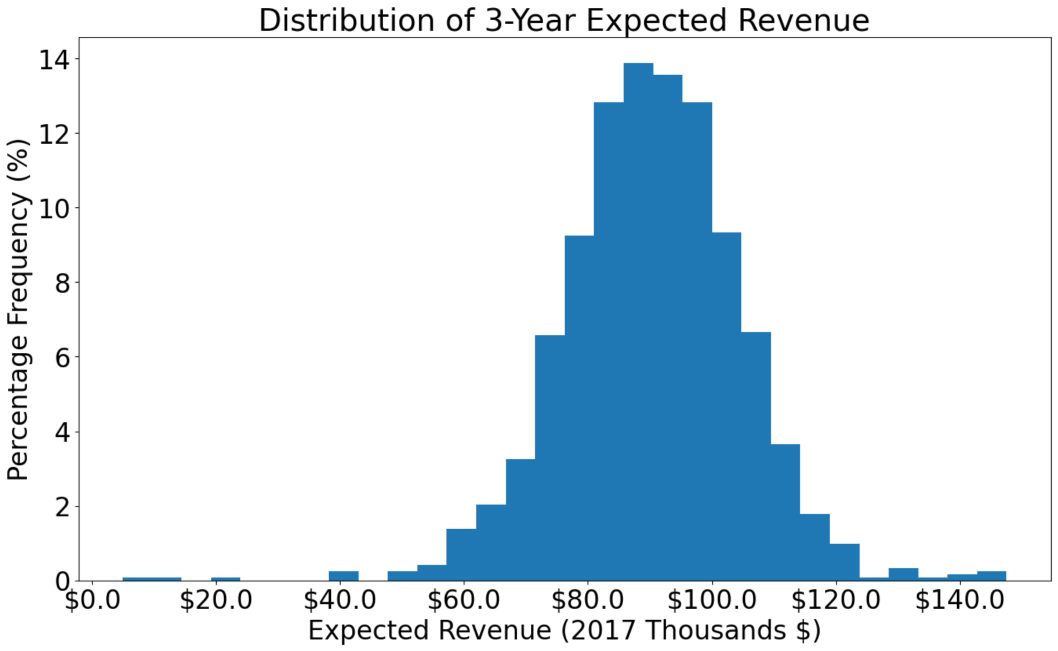

Fixed Costs ▸ Alternate Assumption

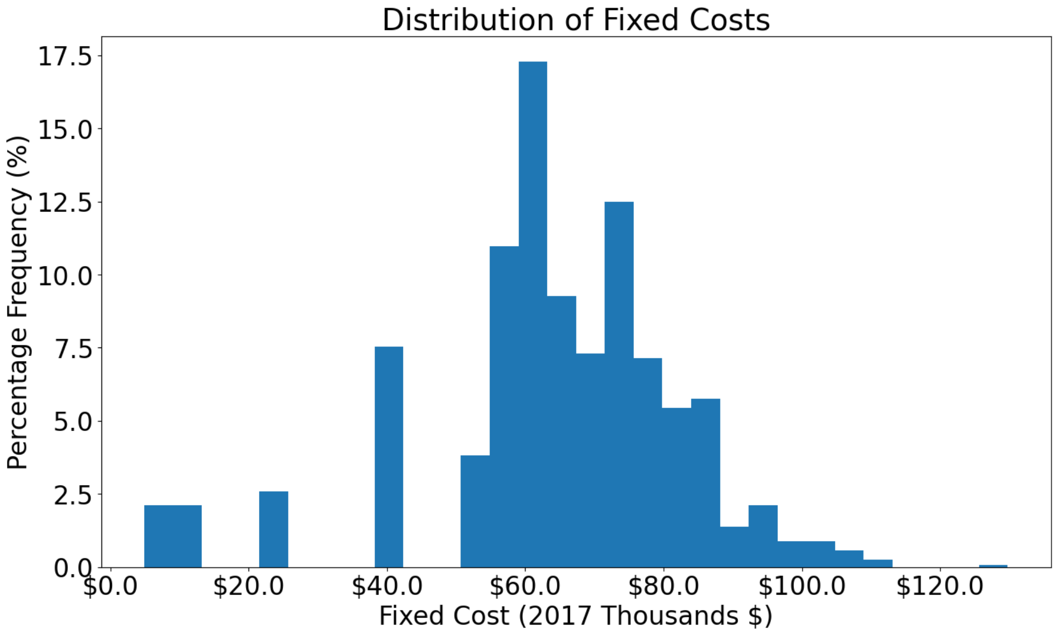

Discussion of Fixed Costs

- Average fixed cost of approximately \$80 thousand for songs in the top 200
- Similar to a Chace 2011 report, which places the non-marketing cost of Rihanna's "Man Down" at approximately \$88 thousand
- Compared to all songs: I estimate an average fixed cost of approximately \$38 for all songs in the data
  - Aguiar and Waldfogel 2018 estimate fixed cost at approximately \$21
  - Their model has focuses on one year of revenue, digital sales, and assumes fixed cost is the same over a year
- Average price-cost margin is approximately 25.6%

Counterfactual Analyses

Counterfactuals

- Recall that the choice probability a firm faces is:

$$P(i \text{ listens to } j) = P(\text{RS surfaces } j \text{ to } i) \times P(i \text{ listens to } j | \text{RS surfaces } j \text{ to } i)$$

- Counterfactuals will change the first term
- Random Recommendations:
  - Each song has a uniform random chance of being recommended
  - Similar to a very naive consumer search process
- Popular Recommendations:
  - Each song is recommended proportional to its listening shares
  - Similar to a ban on personalized recommendations, or the digital sales era

Counterfactual Methodology

1. Draw 500 consumers and have them listen to 15 vintage songs at random, to generate a streaming session
2. Give the consumer a new song, and compute the listening probability
3. Repeat for all 1232 new releases
4. Compute the expected revenue for each song, assuming recommendations are uniform random or proportional to listens from the simulation process

Counterfactual Objects of Interest

- Expected profit: how do profit levels change, and how does this affect entry? What are the new price-cost margins?

- Song characteristics: how do the characteristics of profitable and unprofitable (in expectation) songs differ?

- Consumer welfare: define consumer surplus as the log-sum of the utility of all profitable songs in the choice set:

$$CS = \log \left(\sum_{j \in \mathcal{J}} \exp(\beta X_j + \gamma Y_i + \zeta_s) \right) \tag{11}$$

- How does consumer surplus change under the counterfactuals?

Random Recommendations - Profit

Alternate Assumption

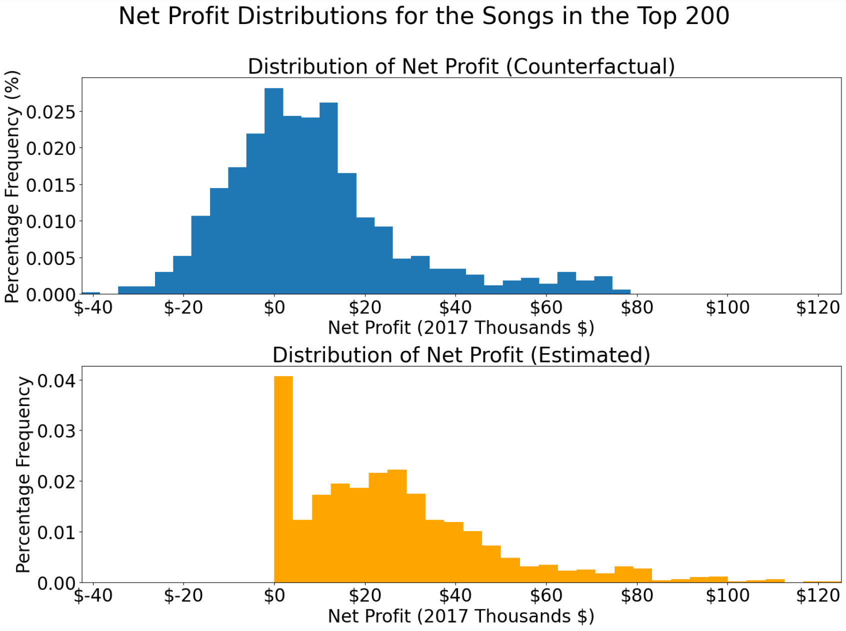

Random Recommendations - Characteristics

Feature	Profitable	Unprofitable	Diff	KS-Test P-Value
Duration (s)	210.774	202.661	-8.113**	0.1194
	(1.990)	(2.087)		
Tempo	124.670	123.687	-0.983	0.0376
	(1.049)	(1.535)		
Energy	0.622	0.651	0.029**	0.0009
	(0.006)	(0.007)		
Valence	0.448	0.424	-0.023	0.0000
	(0.007)	(0.011)		
Acousticness	0.214	0.180	-0.034*	0.0000
	(0.008)	(0.011)		
Instrumentalness	0.005	0.013	0.008*	0.8837
	(0.001)	(0.004)		
Loudness	-6.593	-6.009	0.584***	0.0000
	(0.084)	(0.106)		

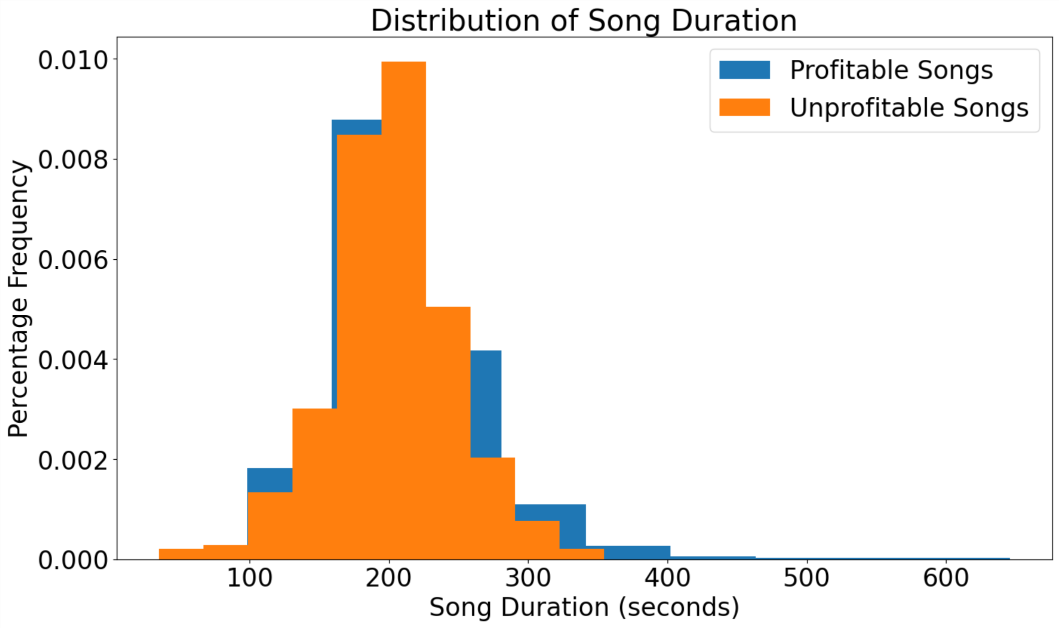
Distribution of Song Duration
Profitable Songs
Unprofitable Songs
Percentage Frequency
0.000
0.002
0.004
0.006
0.008
0.010
100
200
300
400
500
600
Song Duration (seconds)

Random Recommendations - Energy

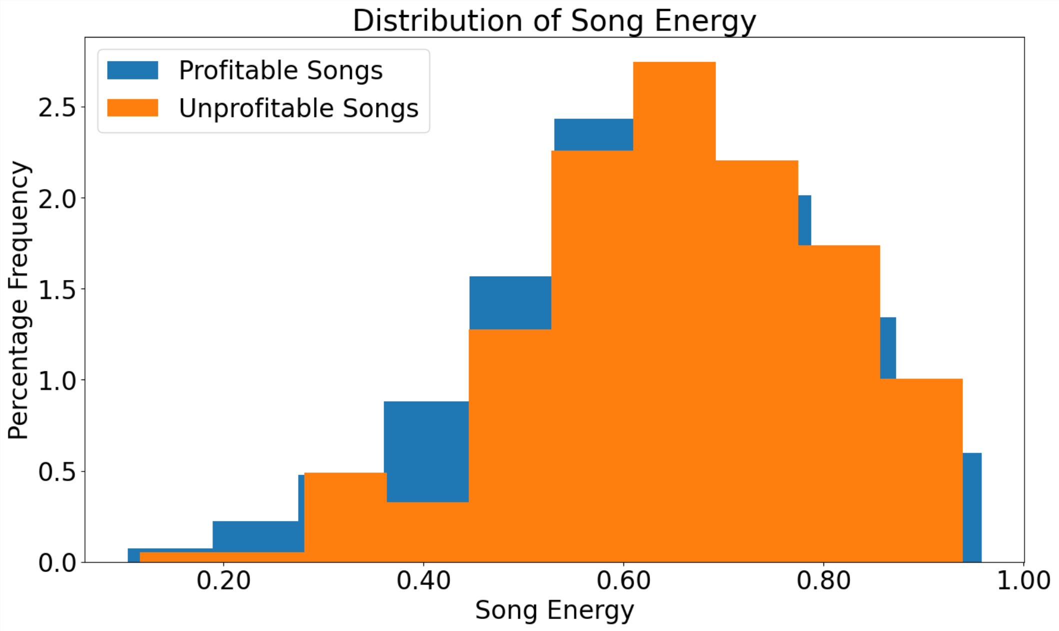

Random Recommendations - Welfare and Discussion

- 447 of the 1232 songs (36 percent) released in the first three quarters of 2018 would not have been released
- The recommender system allows shorter, more energetic, and more "clubby" songs to be released
- These songs also end up being more homogeneous in their characteristic distributions
- Price-cost margins are lower at 17.5%
- Consumer surplus decreases by 4%

Popular Recommendations - Profit

Alternate Assumption

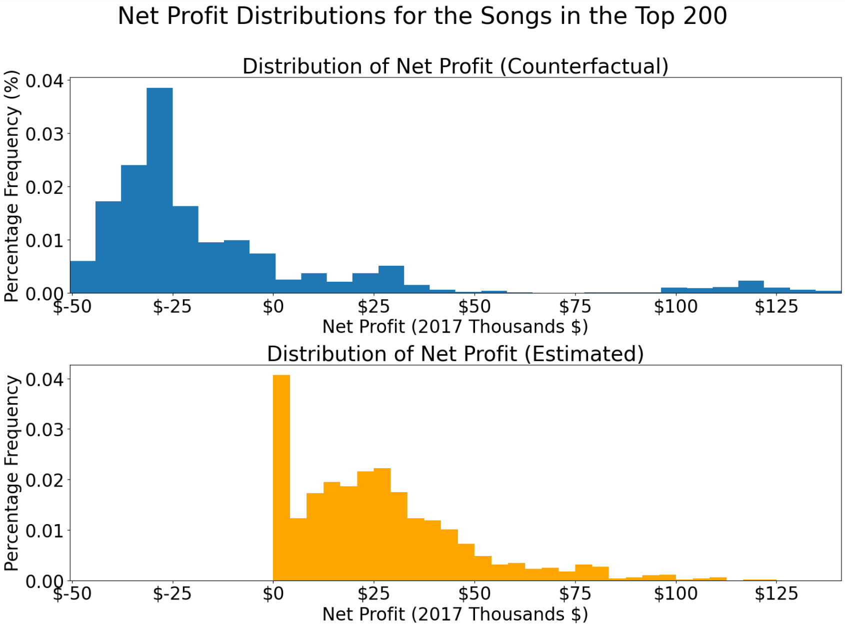

Popular Recommendations - Characteristics

Feature	Profitable	Unprofitable	Diff	KS-Test P-Value
Duration (s)	204.570	208.540	3.970	0.3674
	(3.403)	(1.644)		
Tempo	122.776	124.647	1.872	0.0595
	(1.859)	(0.979)		
Valence	0.452	0.436	-0.016	0.0270
	(0.013)	(0.007)		

Popular Recommendations - Valence

▸ Alternate Assumption

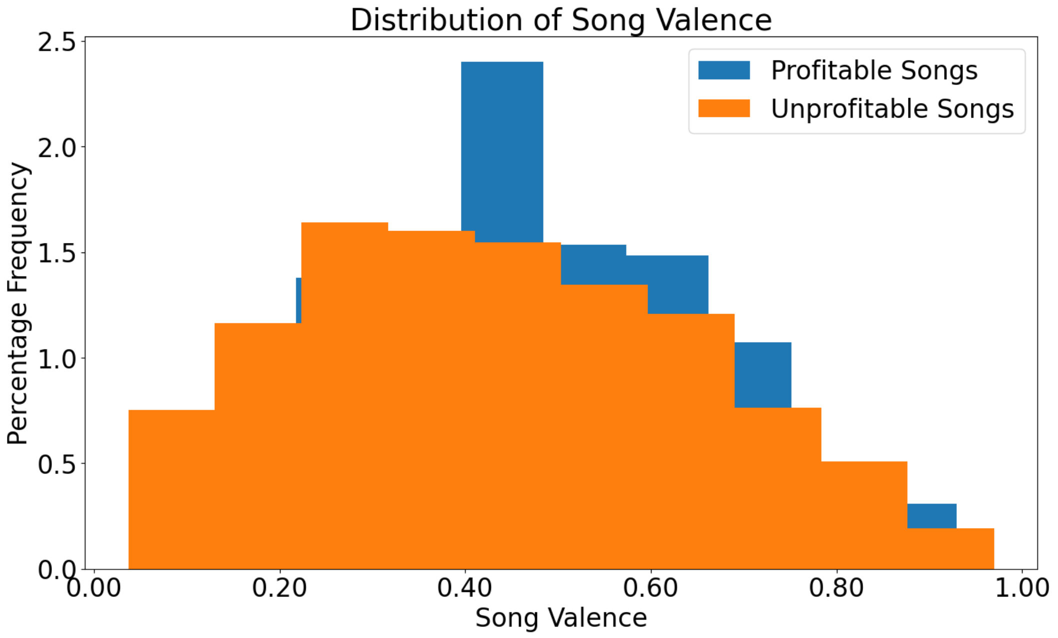

Popular Recommendations - Tempo

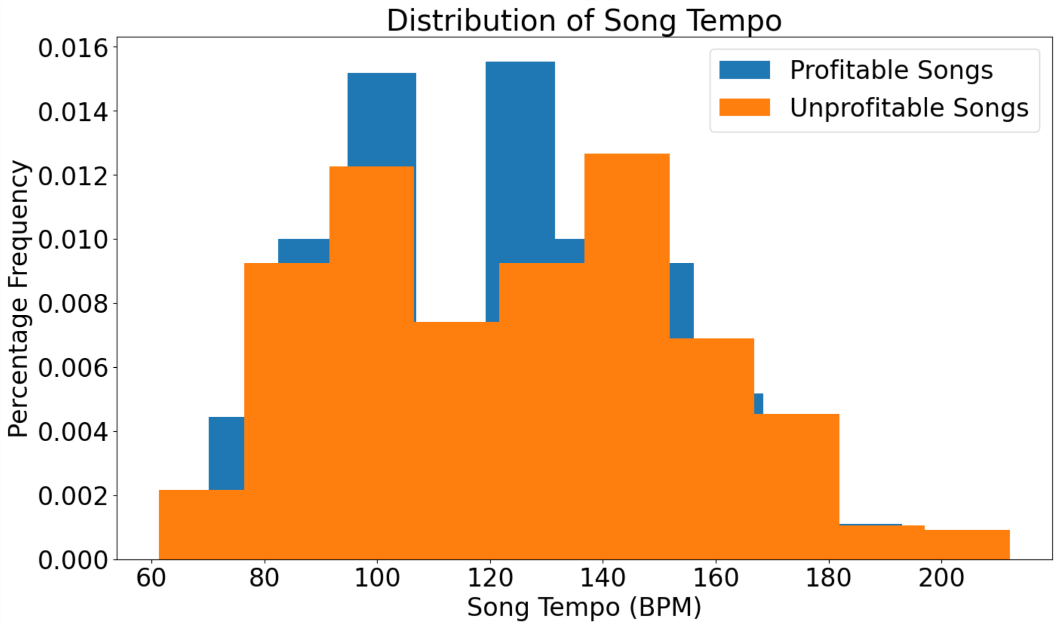

Popular Recommendations - Discussion

- 1012 of the 1232 songs (82 percent) released in the first three quarters of 2018 would not have been released
- Those songs didn't get enough interest in the simulated consumer sessions, so they wouldn't get recommendations more widely
- The songs that do release, however, capture more of the market and earn greater profits.
- Price-cost margins increase to 48.6%
- Those songs also are more homogeneous in tempo and valence, sounding happier and slower
- Consumer surplus decreases by 15.7%

Conclusion

Conclusion

- I have estimated a model of the music industry, focusing on the effects of recommender systems on music production
- My model demonstrates a gap between consumer preferences and the recommender system, and producers are likely to respond to a mixture of these preferences
- I have estimated a fixed cost of production of $80,000 on average, in line with Chace 2011
- My counterfactual estimates suggest that recommender systems allow more songs to enter the market, and that these songs are more homogeneous in their characteristics

Current and Future Research Directions

- **Current focus:**
  - Analyzing the impact of consumer heterogeneity on our findings
  - Exploring the characteristics of a consumer-optimal recommender system
- **Future research:**
  - Investigating Spotify's strategic considerations in recommender system design
  - Examining the consumer's extensive margin decision (i.e., choosing to subscribe)

Appendix

Industry Structure — Concentration

Back

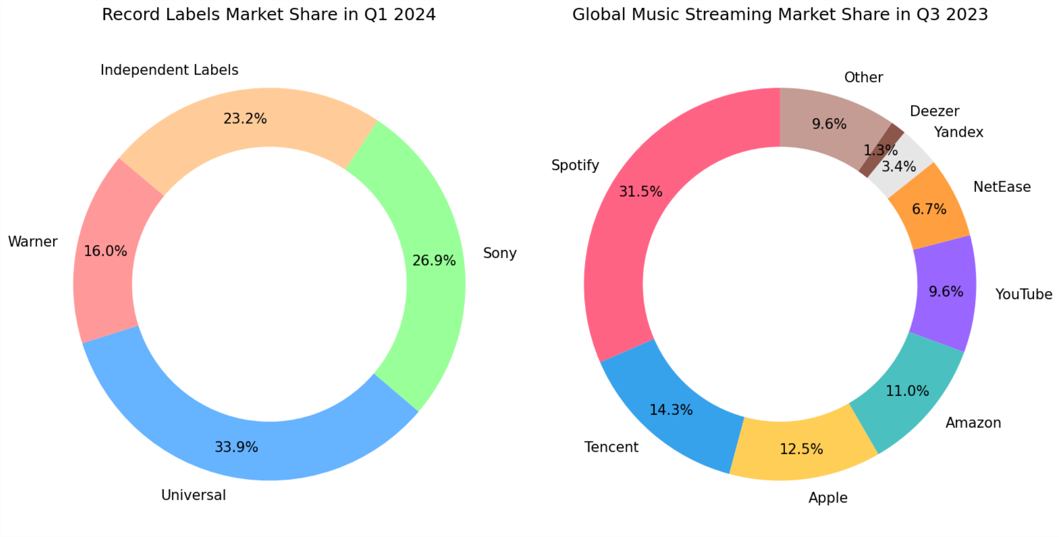

Music Characteristics

Name	Artist	Duration (min)	Tempo (BPM)	Key
Sympathy for the Devil	Rolling Stones	6.3	116	A
Bohemian Rhapsody	Queen	5.9	71	C
Sweet Dreams	Eurythmics	3.6	125	C
Bad Romance	Lady Gaga	4.9	119	C
My Universe	BTS, Coldplay	3.8	105	A

Name	Danceability	Energy	Speechiness	Valence
Sympathy for the Devil	0.7	0.67	0.21	0.56
Bohemian Rhapsody	0.41	0.40	0.05	0.22
Sweet Dreams	0.69	0.71	0.03	0.88
Bad Romance	0.7	0.92	0.04	0.71
My Universe	0.59	0.7	0.04	0.44

Characteristic Definitions (from Spotify's API Reference)

- Danceability: "Danceability describes how suitable a track is for dancing based on a combination of musical elements including tempo, rhythm stability, beat strength, and overall regularity. A value of 0.0 is least danceable and 1.0 is most danceable."
- Energy: "Energy is a measure from 0.0 to 1.0 and represents a perceptual measure of intensity and activity. Typically, energetic tracks feel fast, loud, and noisy. Perceptual features contributing to this attribute include dynamic range, perceived loudness, timbre, onset rate, and general entropy."
- Key: The key of the track, using pitch class notation
- Speechiness: "Speechiness detects the presence of spoken words in a track. The more exclusively speech-like the recording (e.g., talk show, audiobook, poetry), the closer to 1.0 the attribute value."
- Tempo: "The overall estimated tempo of a track in beats per minute (BPM). Tempo is the speed or pace of a given piece and derives directly from the average beat duration."
- Valence: "A measure from 0.0 to 1.0 describing the musical positiveness conveyed by a track. Tracks with high valence sound more positive (e.g., happy, cheerful, euphoric), while tracks with low valence sound more negative (e.g., sad, depressed, angry)."

Back

Sampling Strategy

- For each song in the MSSD, I sample the greater of the following:
    1. 0.5 percent of the total number of consumers who have this song in their streaming session
    2. 1 consumer who has this song in their streaming session
- When I sample a consumer, I sample all the songs in their streaming session, not just the song in question
- Remove duplicate streaming sessions after sampling

Back

Time Use Data

Back

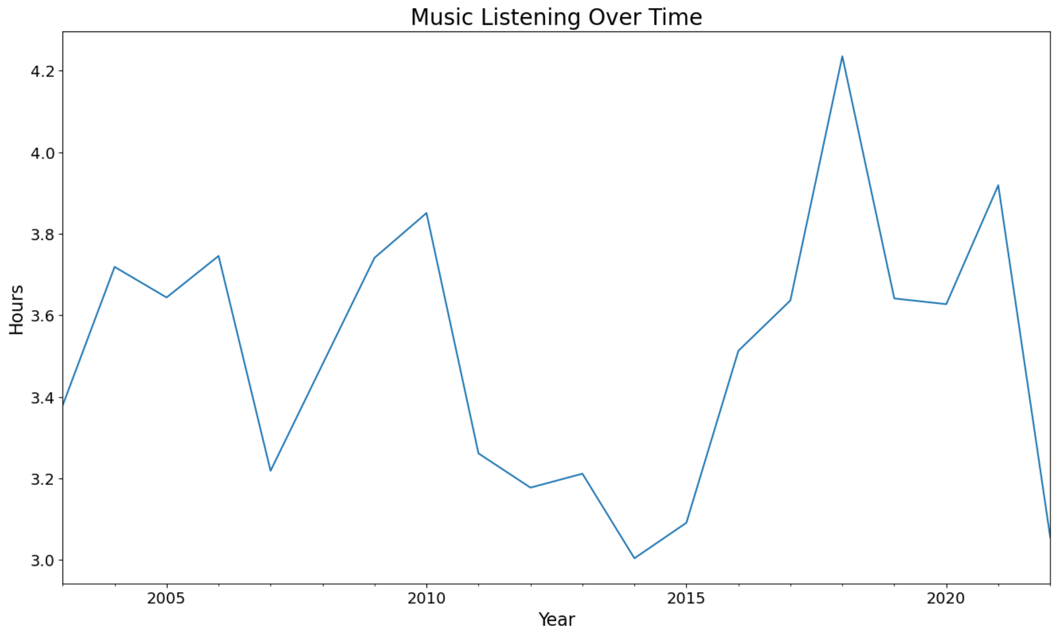

Firms and Platform Payoffs

- Spotify pays rightsholders based on royalty-bearing streams (RBS), defined as a listen of 30 seconds or more
- Rightsholders earn revenue based on their streamshare:

$$\text{Streamshare}_j = \frac{RBS_j}{\sum_{k \in \mathcal{K}} RBS_k}$$

- Example of two six-minute streaming sessions:

Person A	Cruel Summer (T. Swift): 2:58	Cruel Summer (T. Swift): 2:58	6 Minutes
Person B	Like a Prayer (Madonna): 5:40		6 Minutes

Payoff Calculations:

Taylor Swift: $\frac{2 \text{ RBS}}{3 \text{ RBS}} \times \$10$ per subscriber $\times$ 2 subscribers $= \$13.33$

Madonna: $\frac{1 \text{ RBS}}{3 \text{ RBS}} \times \$10$ per subscriber $\times$ 2 subscribers $= \$6.67$

Earnings Flowchart

How a Song Earns Royalties ($\tau = 0.6$, $w_s = \$6$, $w_a = \$0.0225$)

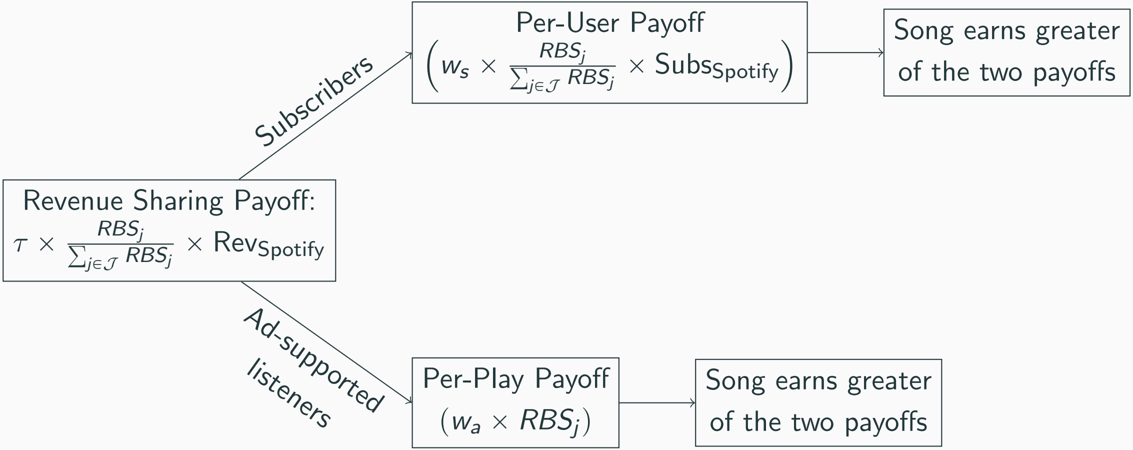

Back

Motivating Panel Regression

- I estimate a panel regression to assess the correlation between the introduction of new music formats and song duration

	Duration (minutes)
Spotify	−0.208***
	(0.047)
Recommendations	-0.394***
	(0.081)
Observations	6,879
Adjusted R^2	0.237
Note:	*p<0.1; **p<0.05; ***p<0.01

Back

Variation in Other Characteristics

Back

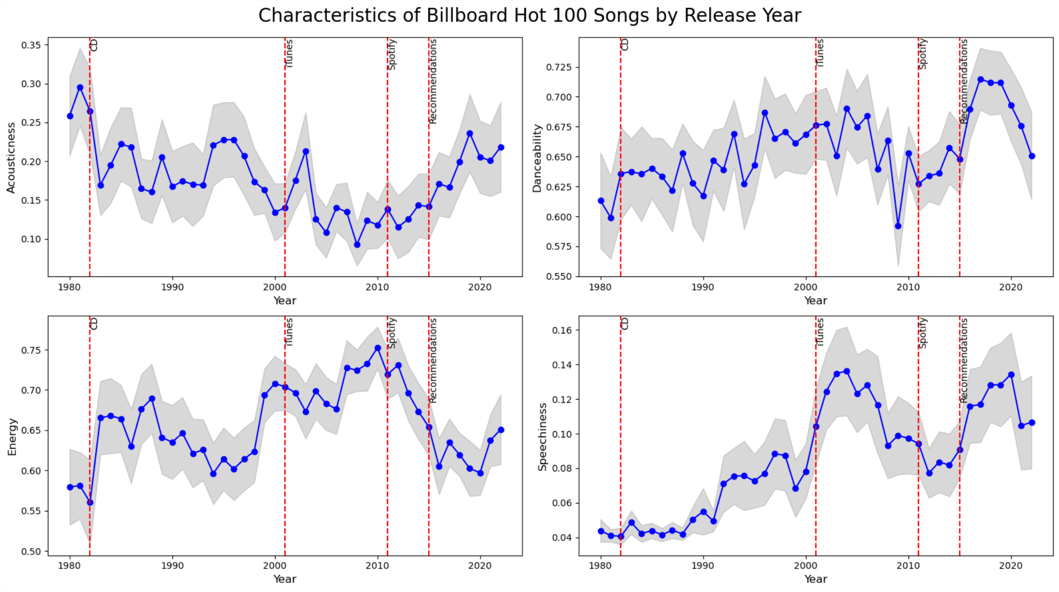

Song Lifecycle

Back

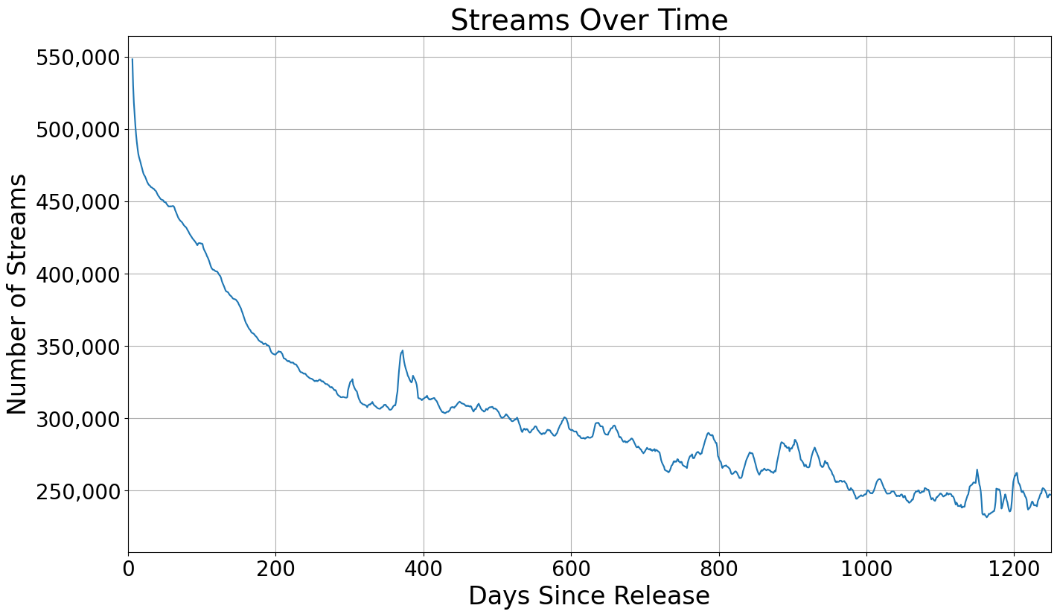

Characteristic Correlation

Back

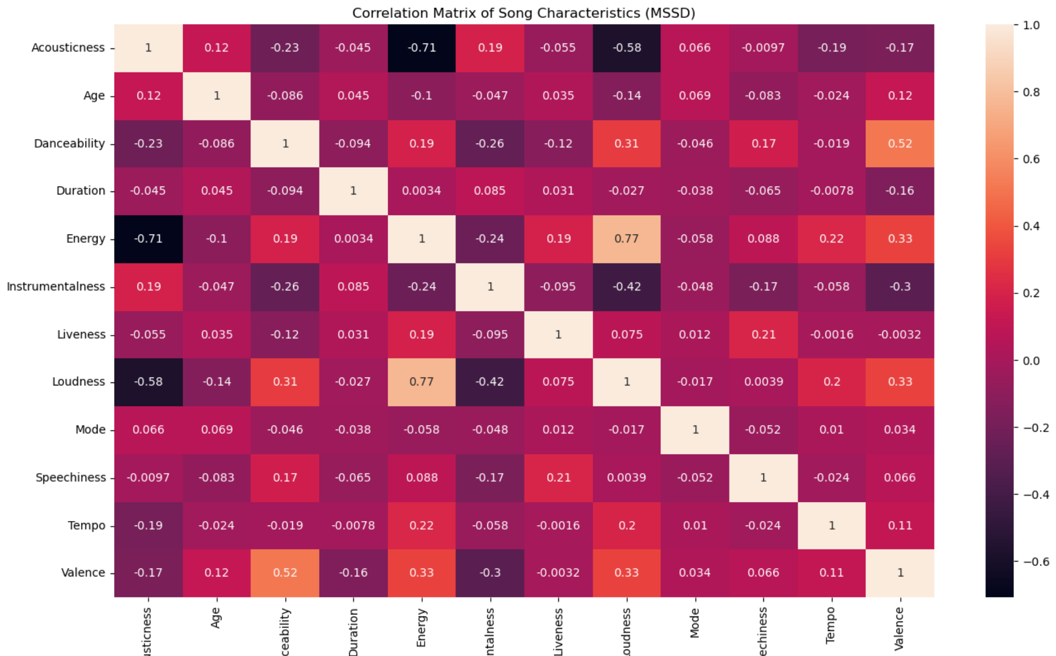

Song Descriptive Statistics - Spotify Charts ($N = 9,244$ songs) Back

	Mean	Median	Standard Deviation	Min	Max
Duration (seconds)	203.27	199.32	54.28	30.13	943.53
Release Year	2019	2019	1.39	2017	2021
Acousticness	0.23	0.13	0.25	0.00	0.99
Danceability	0.67	0.68	0.15	0.06	0.98
Energy	0.62	0.63	0.17	0.01	1.00
Instrumentalness	0.01	0.00	0.09	0.00	0.96
Liveness	0.18	0.13	0.14	0.02	0.97
Loudness (dB LUFS)	-6.83	-6.38	2.71	-38.86	0.35
Mode	0.61	1.00	0.49	0.00	1.00
Speechiness	0.15	0.09	0.13	0.02	0.97
Tempo (BPM)	122.41	122.08	30.04	40.32	212.06
Time Signature	0.97	1.00	0.16	0.00	1.00
Valence	0.46	0.46	0.22	0.03	0.98

Expected Skip Characteristics

- If consumers are listening to an algorithmic playlist or radio station, then their expected utility of skipping comes from the average characteristics of the songs they have received in their streaming session so far.
- If consumers are listening to their own catalog or playlist, or a song they searched for, then their expected utility of skipping comes from the average characteristics of the songs in their entire streaming session.
- If consumers are listening to editorial playlists or top 200 playlists, then their expected utility of skipping depends on whether they shuffle the playlist: if they do, expected utility comes from the characteristics of songs received so far; if not, then the expected utility comes from the average characteristics of the songs in streaming session.

Back

Recommender Systems — Overview

- Recommender systems attempt to predict whether consumers will enjoy/use a product, and suggest it to them
- Recommender systems are ubiquitous in the digital economy: Netflix, Amazon, Spotify, etc.
- There are typically three types of recommender systems: content-based, collaborative, and hybrid:
    - Content-based systems recommend products based on their underlying characteristics
    - Collaborative systems recommend products based on the characteristics of products that similar consumers have enjoyed (e.g., "people who liked this also liked")
    - Hybrid systems combine both content-based and collaborative systems.

Spotify's Recommender System

- Spotify's recommender system recommends songs based on their own characteristics, and the characteristics of songs that similar consumers have enjoyed
- Spotify formulates this model in the following equation:

$$r(j, x) = \sigma(\theta + \theta_j 1_j + \theta_x X) \tag{12}$$

- $r(j, x)$ is the probability that a consumer will complete a song j with characteristics x
- σ is the logistic function, θ are parameters, and 1_j is a song-specific dummy variable and X is a vector of song and consumer characteristics

Spotify's Recommender System — Serving Recommendations

- To determine which songs to recommend to a consumer, Spotify solves a multi-armed bandit problem
- Spotify seeks to maximize engagement, and each song they recommend is a "bandit arm"
- Spotify balances exploration (recommending new songs) and exploitation (recommending songs that have been enjoyed by similar consumers) to maximize engagement while sharpening the recommender system

Back

Robustness Check: Demand over Time

- Estimate demand for the period July - September 2018, but assume preferences are consistent over time
- As a robustness check, I estimate demand across multiple years using Spotify chart data from 2017-2021
- Two approaches:
  - DiD regression of streamns on song characteristics, interacted with time dummies
  - Discrete-choice model of song selection, with characteristics interacted with time dummies

Back

OLS Demand over Time

- Estimating equation:

$$\begin{aligned} log(\text{Streams}_{jt}) &= \alpha + \beta_1 \text{Duration}_j + \beta_2 \text{Time Trend}_t \\ &+ \delta(\text{Duration} \times \text{Time Trend})_{jt} + \gamma X_j + \rho_t + \epsilon_{jt} \end{aligned} \tag{13}$$

- Controls include song characteristics and week fixed effects

OLS Robustness Results

	Log(Streams)
Duration	−0.027***
	(0.002)
Time Trend	0.00003***
	(0.000)
Duration × Time Trend	0.00002***
	(0.000)
Observations	364,081
Adjusted R^2	0.019
F Statistic	98.576***
Note:	*p<0.1; **p<0.05; ***p<0.01

◂ Back

Discrete-Choice Demand over Time

- Random utility model over songs in Spotify's top 50, with the remaining songs in the top 200 as the outside option
- Utility function:

$$\begin{aligned} U_{ijt\,i} &= \alpha + \beta \text{Duration}_j \\ &\quad + \delta(\text{Duration} \times \text{Time Trend}_{jt}) + \gamma X_j + \rho_t + \epsilon_{ijt} \end{aligned} \tag{14}$$

- Instrument duration using other characteristics (BLP instruments)
- Controls include song characteristics and month fixed effects

DCM Robustness Results

	Dependent variable: $log(Market\ Share) - log(Outside\ Share)$
Duration	-0.487^{***}
	(0.059)
Duration^2	0.006
	(0.005)
Duration $\times$ Time Trend	0.003^{***}
	(0.0003)
Observations	10,350
Note:	*p<0.1; **p<0.05; ***p<0.01

Back

Computing Streamshare

- Rival songs in an oblivious equilibrium have the same set of long-run average characteristics
- Two assumptions for number of rivals: top 200 songs, or all songs on Spotify (40mn in 2018)
- For the top 200, I estimate the percent of total streams on Spotify are in the top 200, and downscale the revenue to match:
  - Assume that the average listener listens to 125 minutes per day
  - Divide 125 minutes by the average length of songs in the top 200 to get the number of streams per listener per day
  - Multiply by the number of listeners each quarter to get the total number of streams

Back

Song Network Structure

Back

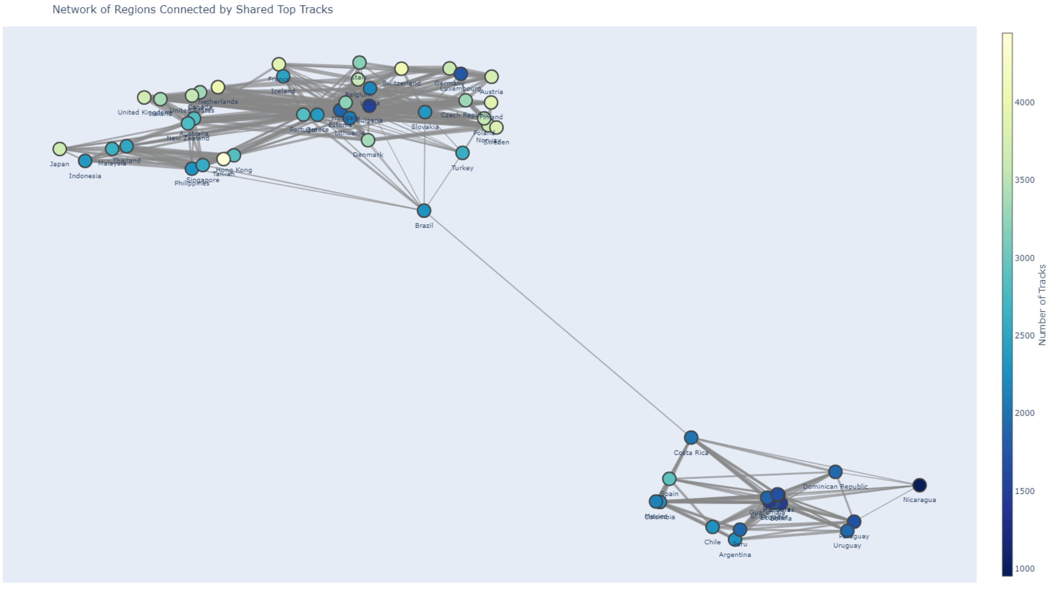

Song Characteristic Markov Process Estimates

	Acousticness	Age	Danceability	Duration	Energy	Instrumentalness
Constant	0.447**	-0.325	-0.320*	0.028	-0.320**	-0.131***
	(0.144)	(0.265)	(0.127)	(0.075)	(0.111)	(0.040)
Drift	0.000**	0.000	-0.000***	-0.000***	-0.000	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Own-characteristic$_{t-1}$	**0.927*** **	**0.832*** **	**0.784*** **	**0.907*** **	**0.858*** **	**0.815*** **
	(0.027)	**(0.028)**	**(0.023)**	**(0.011)**	**(0.028)**	**(0.015)**
Observations	1820	1820	1820	1820	1820	1820
	Liveness	Loudness	Mode	Speechiness	Tempo	Valence
Constant	-0.104	-0.015	0.061	-0.044	-0.122	-0.190*
	(0.064)	(0.122)	(0.040)	(0.095)	(0.069)	(0.086)
Drift	0.000***	-0.000	0.000*	-0.000***	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Own-characteristic$_{t-1}$	**0.870*** **	**0.900*** **	**0.918*** **	**0.887*** **	**0.860*** **	**0.924*** **
	(0.014)	**(0.032)**	**(0.011)**	**(0.017)**	**(0.012)**	**(0.012)**
Observations	1820	1820	1820	1820	1820	1820

Note: Standard errors in parentheses. *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$. Own-characteristic lag terms are in bold.

Recommender System Markov Process Estimates

	Dependent variable: Predicted Probability (ϕ_t)
ϕ_{t-1}	0.695***
	(0.020)
Drift	0.000***
	(0.000)
Constant	0.103***
	(0.007)
Observations	1826

Expected Revenues

Back

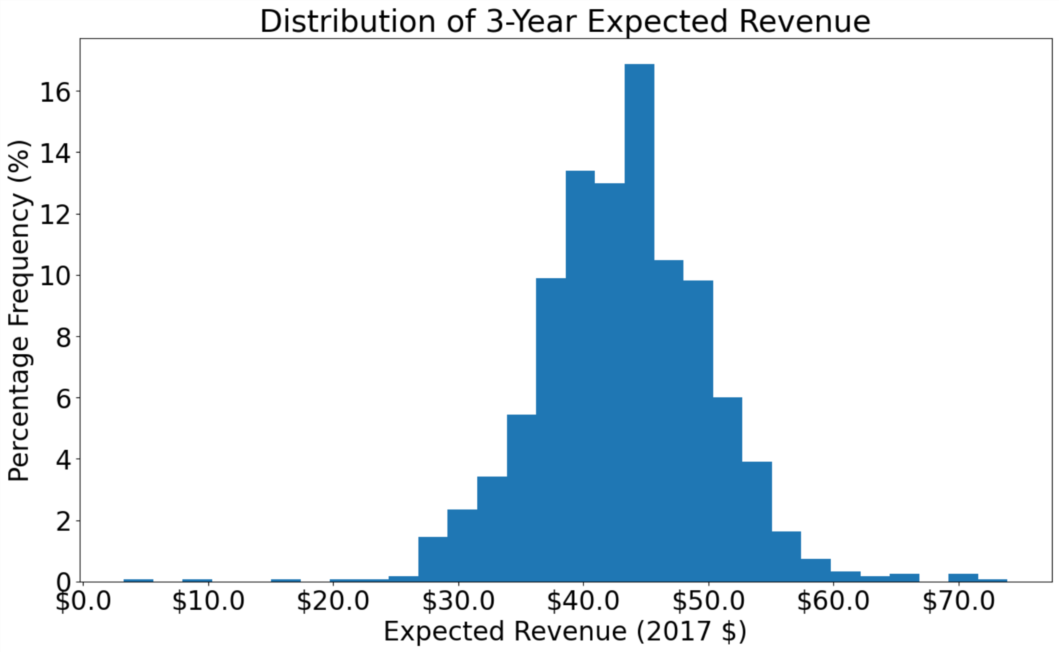

Fixed Costs

Back

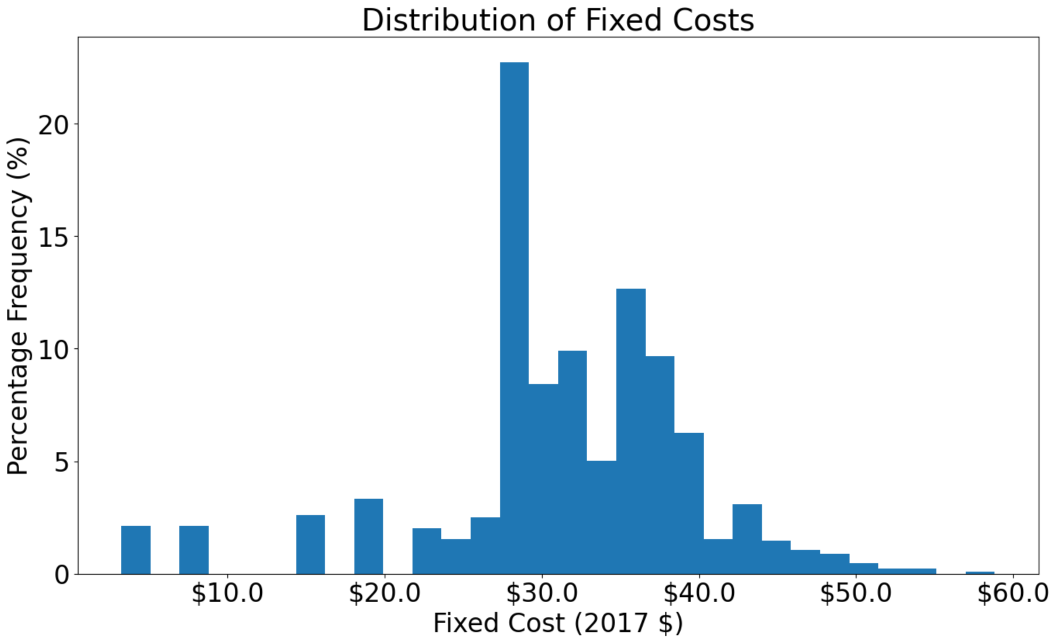

Random Recommendations - Profit

Back

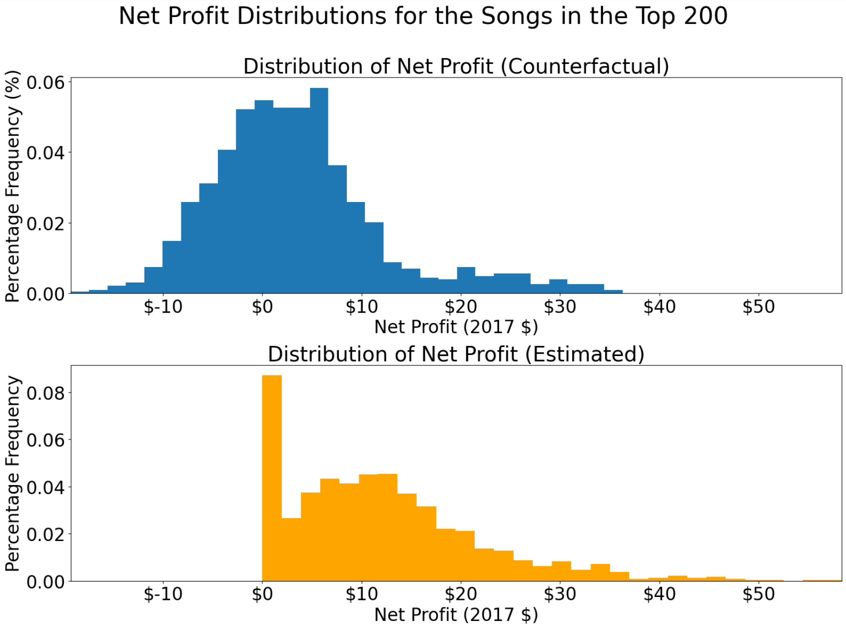

Random Recommendations - Duration

Back

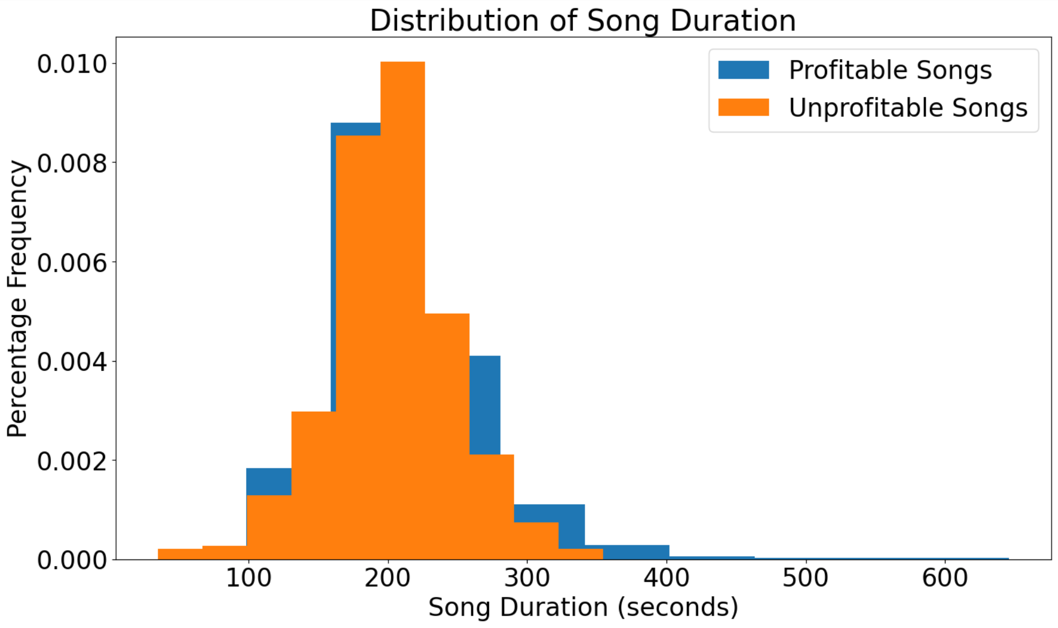

Popular Recommendations - Profit

Back

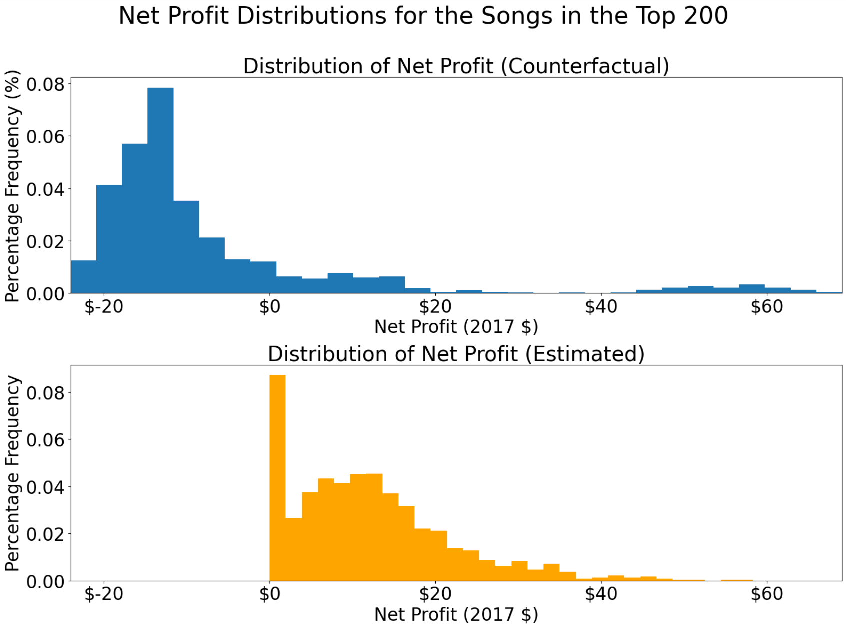

Popular Recommendations - Valence

Back