Document de travail
Discussion paper

N°2026-05

https://lem.univ-lille.fr

Smells Like Rock Spirit : Attention Concentration in Cultural Markets

Etienne FARVAQUE
Univ. Lille, CNRS, IESEG School of Management, UMR 9221 – LEM, Lille Economie Management, F-59000, France| etienne.farvaque@univ-lille.fr

Arthur HEMDAT
Univ. Lille, CNRS, IESEG School of Management, UMR 9221 – LEM, Lille Economie Management, F-59000, France | arthur.hemdat@univ-lille.fr

Nicolas JEAN
IESEG School of Management, F-59000, France | n.jean@ieseg.fr

Frédéric GANNON
EDEHN, Université du Havre, France |frederic.gannon@univ-lehavre.fr

Smells Like Rock Spirit:
Attention Concentration in Cultural Markets

Etienne Farvaque* Arthur Hemdat†
Nicolas Jean‡ Frédéric Gannon§

March 2026

Abstract

Cultural markets are often characterized by highly unequal distributions of attention, in which a small number of artists capture a disproportionate share of recognition and demand. This paper studies the structure of attention in the music industry by analyzing how popularity is distributed across rock bands. Using a comprehensive dataset of 8,964 bands drawn from the Spirit of Rock Encyclopedia, we measure popularity through the number of declared fans and examine the resulting rank–size distribution. The results reveal strong power-law patterns in the upper tail of the distribution, consistent with highly concentrated attention dynamics. We also document substantial variation in the estimated scaling exponent across musical subgenres and national markets, suggesting that the degree of concentration differs across cultural environments. These findings contribute to the economics of cultural industries by providing new evidence on how recognition is distributed across artists and by offering a simple quantitative measure of concentration in cultural markets.

Keywords: Cultural markets ; Attention distribution ; Superstar effects ; Power laws ; Rock music

JEL: C46, L10, L82, Z11.

*Univ. Lille, CNRS, IESEG School of Management, UMR 9221-LEM-Lille Économie Management, F-59000 Lille, France, and CIRANO (Montréal, Québec, Canada). etienne.farvaque@univ-lille.fr

†Univ. Lille, CNRS, IESEG School of Management, UMR 9221-LEM-Lille Économie Management, F-59000 Lille, France. arthur.hemdat@univ-lille.fr

‡IESEG School of Management, F-59000 Lille, France. n.jean@ieseg.fr

§EDEHN, Université du Havre, France. frederic.gannon@univ-lehavre.fr

The authors would like to thank, for useful remarks and comments, Thomas Baudin, Ekaterina Borisova, Fabrice Defever, Hubert Jayet, Taiju Kitano, Tanguy Le Fur, Simone Moriconi, Tomohito Okabe, as well as other seminar and conferences participants. The usual disclaimer applies.
Part of this research was made while the first author was hosted by Kansai Gakuin University (Nishinomiya, Japan), whose hospitality is gratefully acknowledged.

1 Introduction

Cultural markets are often characterized by extreme inequality in the distribution of attention and success. A small number of artists attract a disproportionate share of public recognition, while a much larger number remain confined to niche audiences or local scenes. Understanding how attention is distributed across cultural producers is therefore central to the economics of cultural industries, where reputational dynamics, network effects, and media exposure can generate highly concentrated outcomes. The literature on “superstars” (Rosen, 1981; Krueger, 2005; Connolly and Krueger, 2006) has long emphasized that small differences in talent, timing, or exposure may translate into very large differences in economic rewards.

In such environments, one would expect popularity to be highly skewed, with a few acts capturing most of the attention and revenue. While cultural economics has extensively studied individual determinants of taste, pricing, and market organization, the overall statistical structure of recognition across artists remains less explored. A few recent contributions have begun to examine how influence or demand is distributed within cultural fields - for instance in contemporary art (Etro and Stepanova, 2018; Prinz, 2022) or music streaming markets (McKenzie et al., 2021), but systematic analyses of popularity distributions across large populations of cultural producers remain relatively scarce.

Recent work in economics suggests that highly unequal outcomes often follow regular statistical patterns, including rank-size relationships commonly approximated by power laws. Power-law distributions, and Zipf’s law in particular, have long attracted attention across the social sciences. Originally formulated in the context of linguistics (Zipf, 1949), they describe situations in which the frequency of an item is inversely proportional to its rank: the second most frequent word appears half as often as the first, the third one-third as often, and so on. These distributions are notable for their “thick tails”, whereby a small number of extreme values account for a disproportionate share of the total.

In economics, such patterns have been documented in a wide range of contexts, including city sizes (Gabaix, 1999), firm sizes (Axtell, 2001), CEO compensation (Gabaix et al., 2014), and trade volumes (del Rosal, 2018). Theoretical models explain the emergence of power laws through mechanisms such as proportionate growth (Gibrat’s law), preferential attachment (Simon, 1955), or optimization under constraints (Mandelbrot, 1953). Surveys of this literature (e.g., Clauset et al., 2009, Gabaix, 2009, 2016, and Kapeller and Steinerberger, 2025) provide comprehensive discussions of these mechanisms and the statistical challenges involved in identifying power-law relationships.

Rock music provides an especially informative empirical setting in which to study the distributional properties of musical success. Its market has relatively clear boundaries and a well-defined internal structure, shaped by decades of stylistic evolution and the formation of distinct fan communities. At the same time, the boundaries of rock are sufficiently porous to encompass commercially mainstream styles, most notably those often classified as Pop, with which it shares historical roots, overlapping audiences, and sometimes permeable artist trajectories. Although pop is not strictly a sub-genre of rock, this degree of overlap makes the rock universe broader and more heterogeneous than many other major genres, where the mainstream segment tends to be narrower or more sharply delimited.

Rock also exhibits an unusual proliferation of sub-genres, ranging from globally recognized styles to highly specialized niches with dense, identifiable fan bases. While other genres, such as Jazz or Hip-hop, also display internal differentiation, the breadth, persistence, and cultural visibility of sub-genre formation in Rock have produced a particularly rich and diverse landscape. This landscape includes both global superstars and thousands of local or niche acts, making rock a natural case for analyzing unequal success. Finally, its long history and prominent digital footprint—spanning artist discographies, its sub-genre classifications, and measures of fan engagement, all facilitate large-scale data collection. Together, these features make Rock an analytically advantageous domain for studying how popularity scales across heterogeneous cultural producers.

Beyond these stylistic and cultural features, the institutional structure of the music industry provides a particularly fertile environment, for the emergence of highly concentrated popularity outcomes. The industry is commonly described as a fringe oligopoly, dominated by a small number of large firms alongside a vast and continuously renewing competitive fringe of independent labels and artists. At the global level, recorded music production and distribution are largely controlled by a handful of multinational "Majors" which together account for a large share of industry revenues. These firms possess unrivalled access to promotion channels, distribution networks, and media exposure, allowing them to scale selected artists to global visibility. By contrast, a large ecosystem of independent labels and self-produced artists operates at the fringe, where entry is relatively easy and creative experimentation is intense.

Innovation in musical styles and talent discovery therefore largely originates from this competitive fringe, while commercial success depends on subsequent selection, amplification, and diffusion by major firms. This two-stage process—characterized by widespread entry, localized experimentation, and highly concentrated scaling capacity—naturally generates extreme dispersion in audience attention. Small initial differences in exposure, timing, or selection can translate into disproportionately large differences in long-run popularity. In this sense, the institutional organization of the music industry provides a structural underpinning for the emergence of heavy-tailed distributions consistent with power-law behavior in artist popularity, complementing existing explanations based on preferential attachment, superstar effects, and cumulative advantage.

This paper examines the concentration of attention in cultural markets. Using a large dataset of rock bands, we analyze how popularity is distributed across artists and show that the resulting rank-size distribution exhibits strong power-law patterns. We define popularity as the number of fans per band in the Spirit of Rock Encyclopedia, a specialized online database that documents thousands of rock bands across countries and subgenres. Unlike platforms such as Wikipedia or Spotify, Spirit of Rock does not filter entries based on prior fame or commercial success, making it a less biased source to tackle the full range of acts in the field. Using this database, we build a dataset of 8,964 bands, each with information on its fan count, country of origin, active status, and subgenre classification.

From a methodological standpoint, we follow recent work on identifying power-law behavior in empirical data (Clauset et al., 2009; Jerez-Lillo et al., 2025; Urzúa, 2020), estimating the optimal exponent α relying on a log-log transformation and minimizing the mean squared error between observed and theoretical frequencies, following recent methodological contributions. We assess the goodness-of-fit of the power-law model, compare it with alternative distributions, and test for heterogeneity across subsamples. In particular, we analyze whether the estimated exponent α differs across subgenres (e.g., Hard Rock, Progressive Rock, Punk/Ska) and across countries (e.g., United States, France, Japan). These comparisons allow us to explore how stylistic fragmentation and national ecosystems shape the distribution of musical recognition, and thus their respective markets.

This paper contributes to the literature in several ways. First, it introduces a new and unusually comprehensive dataset of rock bands with global coverage and detailed metadata, allowing us to analyze the full distribution of popularity rather than focusing only on top acts. Second, it provides new evidence on how the concentration of attention varies across cultural environments, showing that the distribution of recognition differs systematically across musical subgenres and national markets. Third, by estimating and comparing the scaling exponent of popularity distributions, the paper offers a simple quantitative measure of concentration in cultural markets. Finally, the results contribute to broader debates in the economics of superstars and cultural industries by documenting how market structure, genre segmentation, and national ecosystems shape the distribution of fame.

Our findings also have implications beyond the academic literature. Understanding how attention is distributed in cultural markets matters for artists, managers, and platforms alike. If the distribution of recognition is inherently thick-tailed, this may justify business models that rely on blockbuster strategies, catalog monetization, or algorithmic curation.

Conversely, if some markets or genres exhibit flatter distributions, this may point to greater openness, more diverse demand, and broader access to visibility for mid-tier acts. From a policy perspective, such insights could also inform cultural support schemes, especially in markets where a few players dominate exposure and revenue. The distribution of recognition also matters for cultural policy, as the degree of concentration in attention markets may affect the visibility of emerging artists, the diversity of cultural production, and the effectiveness of public support schemes aimed at fostering creative entry.

The rest of the paper is structured as follows. Section 2 presents the theoretical framework of Zipf's law and outlines the estimation procedure. Section 3 describes the dataset and its main features. Section 4 presents the empirical results, both for the full sample and for subgroups by genre and country. The implications of these findings in light of the economics of superstardom and cultural market structure are also discussed. Section 5 concludes.

2 The Number of the Band: Modeling Popularity with Power-Laws

Zipf's law, formulated by the linguist George Kingsley Zipf, posits that, in a large corpus of natural language, the frequency of any word is inversely proportional to its rank in the frequency table. That is, the second-most frequent word appears about half as often as the most frequent one, the third frequent one, third as often, and so on. This power-law distribution reveals a fundamental pattern in language organization, but also appears across various complex systems beyond linguistics. Zipf-like distributions have been observed in city sizes (Gabaix, 1999), firm sizes (Axtell, 2001), income distributions, web traffic, etc. (Brock, 1999, Mitzenmacher, 2004, and Newman, 2005) provide surveys that show that the assumption of power-law distribution has been made in a number of disciplines. These occurrences suggest underlying mechanisms of preferential attachment, self-organization, and efficiency constraints in complex systems. Although Zipf's original work emphasized the cognitive principle of least effort, other models, such as those by Simon (1955) and Mandelbrot (1953), introduced stochastic and information-theoretic explanations. When considering music bands, mechanisms like preferential attachment or stochastic explanations suggest that the popularity[1] of these bands may follow a power-law distribution.

Basically, Zipf's law states that the frequency of an element is inversely proportional to its rank. Formally, if f_r denotes the frequency of the element of rank r, Zipf's law can be expressed as:

[1]In this paper, popularity of a given band is measured by the number of fans of that band in the Spirit of Rock Encyclopedia database, an online repository that catalogs thousands of rock artists. See section 3.

$$f_r^{(\alpha)} = \frac{C}{r^\alpha} \textbf{ for } r = 1, 2, ..., N$$

where r is the rank of the element (1 for the most frequent, 2 for the second most frequent, etc.), $\alpha > 0$ exponent of the law which controls the slope of the decay in a log-log plot, C is a normalization constant that ensures the total probability or frequency sums to 1 (in probability distributions), and N denotes the total number of elements.

When $\alpha = 1$, this corresponds to the strict form of Zipf's law. Otherwise, it is referred to as the generalized Zipf's law.

For the set of frequencies $f_r^{(\alpha)}$ to define a valid probability distribution, it is necessary that $\sum_{r=1}^{N} f_r^{(\alpha)} = 1$. Taking the definition of $f_r^{(\alpha)}$ given by the Zipf's law and introducing the generalized harmonic sum of order α denoted by $H_{N,\alpha} = \sum_{r=1}^{N} \frac{1}{r^\alpha}$, we get:

$$f_r^{(\alpha)} = \frac{1}{r^\alpha . H_{N,\alpha}}$$

This expression represents the normalized form of the generalized Zipf's law. It will serve as our starting point for the logarithmic transformation. Indeed, to better visualize and interpret the empirical data, we need to plot the frequencies using a log-log scale. By applying the standard properties of logarithms to the expression above, we obtain the following equation:

$$\log f_r^{(\alpha)} = \log\left(\frac{1}{r^\alpha . H_{N,\alpha}}\right) = \log\left(\frac{1}{r^\alpha}\right) + \log\left(H_{N,\alpha}\right) = -\alpha \log r - \log\left(H_{N,\alpha}\right)$$

This last equation corresponds to that of a straight line, characterized by a slope equals to $-\alpha$. The final step of our approach consists in estimating the value of α in order to best fit the model to our data[2]. To achieve this, we seek the value of α that minimizes the discrepancy between the observed and the theoretical distribution on a log-log scale. (All the values reported in this section are computed on the basis of the overall sample.) For this purpose, we use the mean squared error (MSE) between the logarithms of the observed and theoretical frequencies:

$$\hat{\alpha} = \arg \min_{\alpha \in [0.1,3]} \left\{ \frac{1}{N} \sum_{r=1}^{N} \left[\log(f_r^{obs}) - \log(f_r^{\alpha})\right]^2 \right\}$$

where:

[2] It is worth noting that the strict form of Zipf's law assumes that the value of α is equal to 1. However, we observed that the optimal value of α corresponding to our data differs from 1, thereby rejecting the strict version of Zipf's law - see below.

$$f_r^{(\alpha)} = \frac{1}{r^\alpha \cdot H_{N,\alpha}}, \quad \text{with} \quad H_{N,\alpha} = \sum_{k=1}^{N} \frac{1}{k^\alpha}$$

Expressed differently, we get:

$$\widehat{\alpha} = \arg\min_{\alpha} \{\text{MSE}_{\log}(\alpha)\}$$

This method allows for estimating the exponent α of the generalized Zipf's law in a way that best fits the observed data. By minimizing the MSE between the logarithms of observed and theoretical frequencies, it adjusts α to most accurately reflect the empirical decay slope of the rank-frequency distribution.

We therefore find the optimal value: ($\widehat{\alpha} \simeq 1.588$), with a mean square error of 0.0469. This low error indicates a very satisfactory fit between the generalized Zipf's law and the empirical distribution of frequencies, confirming the relevance of the model to describe the observed concentration dynamics[3] (see Section 4.1 (Whole Lotta Data: Global Patterns of Popularity) for a more detailed analysis).

The visual fit between the empirical distribution and the generalized Zipf's law can be confirmed by a Kolmogorov–Smirnov test. However, a major issue arises here, as the theoretical law used in the test depends on the parameter α, which is estimated from the data. When a model is fitted directly to the data, the distance between the empirical distribution and the theoretical distribution tends to be underestimated. This artificially increases the apparent quality of the fit. To correct for this bias, we adopt a parametric bootstrap procedure based on Monte Carlo simulations, following the recommendations of Clauset *et al.* (2009).

The Kolmogorov–Smirnov statistic measures the maximum distance between an empirical distribution and a theoretical distribution. In its basic form, it is defined by:

$$D_n = \sup_x |F_n(x) - F(x)|$$

where $F_n(x)$ is the empirical distribution function based on a sample of size n, and $F(x)$ is a fixed theoretical distribution function, independent of the data.

In our case, since we are working directly with ascending ranks, the empirical distribution function above is defined as

$$F^{obs}(r) = \sum_{i=1}^{r} f_i^{obs}$$

[3] This asserts that the mean square deviation between the empirical logarithmic and the theoretical logarithmic frequencies calculated with the optimal alpha is approximately: $\sqrt{0{,}0469} = 0{,}22$

where f_i^{obs} is the relative frequency observed at rank i.

The theoretical distribution function is estimated based on the generalized Zipf's law ($f_r^{(\alpha)}$) for which the parameter α is estimated using MSE (see $\widehat{\alpha}$ above). It is then defined as

$$F^{\widehat{\alpha}}(r) = \sum_{i=1}^{r} f_i^{\widehat{\alpha}}$$

.

Relying on this estimate, we can deduce an adapted version of the Kolmogorov–Smirnov statistic, which becomes:

$$D_{obs} = \max_{1 \leq r \leq n} \left| F^{obs}(r) - F^{\widehat{\alpha}}(r) \right|$$

It is worth noting that since $\widehat{\alpha}$ is estimated from the data, the corresponding theoretical distribution function is also estimated from the data. In this case, we adopt a parametric bootstrap procedure (Clauset *et al.*, 2009) to determine whether the observed difference between empirical and theoretical frequencies can be explained by random fluctuations in the model or whether it is significantly too high. This method consists in comparing D_{obs} not to a fixed theoretical table (classical Kolmogorov–Smirnov test), but to an empirical distribution simulated from Kolmogorov–Smirnov statistics calculated on data generated according to an adjusted model. We first generate B independent synthetic samples containing N draws, where N denotes the number of observed groups in the empirical data, according to the generalized Zipf distribution with the estimated parameter $\widehat{\alpha}$ obtained from the actual data:

$$X^{(b)} = \left\{ x_1^{(b)}, x_2^{(b)}, \ldots, x_N^{(b)} \right\}, \quad b = 1, \ldots, B.$$

Each draw $x_i^{(b)}$ takes values in the fixed set of observed ranks $\{1, \ldots, n\}$, where n denotes the total number of distinct ranks in the empirical data.

For each simulated sample $X^{(b)}$, we compute the corresponding frequency distribution $f_r^{(b)}$ by counting occurrences by rank and normalizing so as to obtain a probability distribution. We then estimate a simulated exponent $\hat{\alpha}^{(b)}$ by fitting the generalized Zipf's law to this distribution, using the same procedure as for the empirical data. The parameter $\hat{\alpha}^{(b)}$ is estimated conditionally on the set of observed ranks, since zero frequencies cannot be included in the log–log fitting procedure. The goodness-of-fit statistic, however, is computed over the full range of ranks, ensuring comparability between the empirical and theoretical distributions.

From this estimate, we calculate two distribution functions: the simulated empirical function

$$F^{(b)}(r) = \sum_{i=1}^{r} f_i^{(b)}$$

and the theoretical function derived from the readjusted Zipf's law

$$F^{\widehat{\alpha}^{(b)}}(r) = \sum_{i=1}^{r} f_i^{\widehat{\alpha}^{(b)}}$$

We then evaluate the simulated Kolmogorov–Smirnov statistic associated with this sample:

$$D_b = \max_{1 \leq r \leq n} \left| F^{(b)}(r) - F^{\widehat{\alpha}^{(b)}}(r) \right|$$

Once the Kolmogorov–Smirnov statistics have been calculated for each simulated sample, we obtain a vector of values:

$$\{D_1, D_2, ..., D_B\}$$

The corrected p-value, denoted $\widehat{p}$, is then estimated as the proportion of these simulated statistics that exceed the observed statistics D_{obs} calculated on the actual data:

$$\widehat{p} = \frac{1}{B} \sum_{b=1}^{B} 1_{\{D_b > D_{obs}\}}$$

The lower the p-value, the more atypical the observed deviation is compared to those generated under Zipf's law. Conversely, a high value indicates that the observed statistic is consistent with the natural fluctuations of the model.

According to Clauset *et al.* (2009), a good rule of thumb to ensure the accuracy of the p-value is to simulate at least $\frac{1}{4\epsilon^2}$ datasets if you want an estimate that is within ϵ of the true value. In our case, to obtain a p-value accurate to two decimal places ($\epsilon = 0.01$), we have to generate at least 2500 simulated samples. Clauset *et al.* (2009) also suggest adopting a relatively conservative rejection threshold, considering that the Zipf's law adjustment is called into question if $\widehat{p} \leq 0.1$, *i.e.*, if there is less than a 10% chance of obtaining such a large deviation randomly under the adjusted model. Applying this to our data, the Kolmogorov–Smirnov statistic observed for the global model is $D_{obs} \simeq 0.0663$, and the corrected p-value obtained by bootstrap is $\widehat{p} = 0$.

At the level of the full sample, the goodness-of-fit test yields a p-value equal to zero, leading to the rejection of a single generalized Zipf model for the aggregated data. However, following Clauset *et al.* (2009), such a rejection does not imply the absence of scaling behavior. Rather, it indicates that a single Zipf's law is statistically too restrictive to capture the structure of a large and heterogeneous dataset.

The global sample aggregates multiple sub-markets and regimes, so that systematic deviations from an idealized Zipf distribution naturally arise and become detectable given the high statistical power of the test. This motivates a disaggregated analysis by country and musical genre, where more homogeneous environments can be examined and where power-law behavior, if present, can be meaningfully assessed.

3 The Data: Counting Fans, Ranking Bands

As stated above, this study uses the Spirit of Rock Encyclopedia, an online repository that catalogs thousands of rock artists across all sub-genres.[4] Unlike Wikipedia, whose editorial guidelines require subjects to attain a predefined level of "notability", Spirit of Rock lists bands regardless of their commercial success or media visibility. Relying on a source that filters entries by prior fame would bias any investigation of how popularity itself emerges.

The resulting dataset contains 8964 distinct acts. For each, we record the year of its first release, the musical sub-genre assigned by the editors, whether the band remains active, and its country of origin. The earliest entry dates back to 1950, and the average band age is 24.8 years. A large majority of the sample (81 percent) is still active, illustrating the platform's willingness to follow ongoing careers as well as legacy acts. Sub-genre labels range from Rock and Progressive Rock to Ska-Punk and Jazz-Rock.[5] Table 1 details the number of observations in the dataset for each genre. The encyclopedic coverage and neutrality of Spirit of Rock have been systematically assessed in Farvaque, 2024, who compares it with Wikipedia and MusicBrainz using random samples. His analysis shows that nearly 30 percent of the artists present in Spirit of Rock have no dedicated Wikipedia page, and that even when bands are listed in MusicBrainz, detailed information is frequently missing. In contrast, Spirit of Rock consistently provides structured and specific entries, even for obscure or defunct bands. Farvaque, 2024 also demonstrates that solo artists are more likely to be included in Wikipedia, confirming a bias in favor of more visible figures. Solo performers appear as separate entries when their artistic output is released under an individual name rather than a collective one. These features hence make Spirit of Rock particularly suitable

[4]https://www.spirit-of-rock.com. This online encyclopedia offers the advantage to be up-to-date when, by definition, even the best books may become outdated the moment they hit the shelves. Moreover, existing books do not cover with the same degree of precision all the genres and sub-genres that are present in the Spirit-of-Rock Encyclopedia.

[5]The genre assigned to each band in the encyclopedia results from community contributions. When a new band entry is created, contributors must select one or several genres from the site's taxonomy. These choices can later be modified or corrected by other users and moderators. Thus, the classification reflects a semi-structured, crowdsourced process rather than a fixed or purely editorial assignment.

for statistical analysis.

Moreover, we have used the Spotify platform API, on which users can also declare if they want to "follow" a band, and are then declared as "fans". Our retrieval reveals that 6800 bands are both on Spotify and in the Spirit of Rock dataset. The correlation between the number of fans declared on Spotify and on Spirit of Rock is equal to 0.6, which is quite high, given that the two popularity measures reflect quite distinct dimensions of musical recognition. The number of "fans" in our database primarily captures recognition within more specialized or genre-specific communities - particularly in subgenres. In contrast, the number of Spotify followers reflects a broader, often more mainstream or international popularity, frequently shaped by algorithmic exposure, platform curation, and maybe even by fraudulent behavior[6].

Moreover, there is a geographical dimension in the popularity of Rock bands: some of them enjoy strong national followings that are sufficient to generate high visibility on global streaming platforms, yet they remain relatively marginal in genre-focused databases. Examples include Duman (Turkey), Radwimps or Greeeen (Japan), or Eraserheads and Parokya Ni Edgar (Philippines). Also, there is a generational effect at play: bands such as Cigarettes After Sex, which appeal strongly to younger audiences through streaming playlists, may be underrepresented in communities more oriented toward traditional or legacy rock culture. Lastly, there is a temporal dimension to consider. Spotify followers reflect a highly current form of popularity — often shaped by recent trends, social media presence, or viral diffusion —, whereas the "fans" variable more likely captures a longer-term or community-based recognition anchored in musical subcultures.

All in all, then, because Spirit of Rock does not privilege already famous artists, it has a larger coverage, and is thus perfectly adequate to study the full rank-size distribution of popularity.

4 Born to Be Ranked: Empirical Analysis

The empirical analysis that follows investigates how attention is allocated across artists within cultural markets. When popularity follows a power-law distribution, the estimated exponent α provides a simple indicator of concentration in the market for recognition. Lower values of α correspond to environments where attention is highly concentrated among a small number of superstar acts, whereas higher values indicate a more fragmented landscape in

[6] https://www.rollingstone.com/music/music-features/lawsuit-spotify-drake-streams-1235457737/. The role played by Artificial Intelligence in producing songs that may also bias the number of "followers" also renders the use of fans-curated bases highly preferable: https://newsroom-deezer.com/2025/09/28-fully-ai-generated-music/

which visibility is distributed across a broader set of performers. From this perspective, differences in α across cohorts, genres, and countries can be interpreted as reflecting variations in market structure, audience segmentation, and the competitive environment faced by artists.

Hence, in this section, we analyze the distribution of popularity among Rock bands, first for the whole dataset, and then by subgenres and countries. For completeness, the Appendix reports the ten most popular bands in the full sample and by major subgroups. The results deliver insights on the structure of each market, on which our discussion focuses.

4.1 Whole Lotta Data: Global Patterns of Popularity

We begin by considering the whole sample of 8964 bands. Before turning to a substantive interpretation of this result, it is important to recall that the goodness-of-fit test rejects the hypothesis of a single generalized Zipf model for the aggregated sample. Consequently, the interpretation that follows should be understood as a descriptive characterization of the observed scaling pattern, rather than as evidence that a single Zipf's law provides an exact statistical fit to the entire dataset. Figure 1 displays the results, revealing a power-law type of distribution for the popularity of all bands, with an exponent equal to 1.588. The value of the exponent reveals a significant degree of concentration of the popularity of Rock bands, which is consistent with the common knowledge of the Rock activity, with superstars (The Rolling Stones, U2, or Coldplay, for example) dominating the stage. Moreover, this degree of skewness of the distribution also reveals the coexistence of a second fringe of bands, characterized by intermediate level of popularity.

A useful way to interpret this concentration pattern is to relate it to the industrial organization of recorded music. The global market operates as a fringe oligopoly, in which a small number of vertically integrated multinational firms - the three "Majors" (Sony Music Entertainment, Universal Music Group, and Warner Music Group), together with more recent entrants such as BMG and Kobalt - control a large share of the production, promotion, and distribution of music. These firms typically account for a large part of global recorded revenues and possess the infrastructure required to transform promising acts into internationally visible superstars.

Alongside this concentrated core lies a vast "independent" fringe made up of hundreds of smaller labels that sign and develop a far larger number of artists. This fringe plays a dual role: it sustains stylistic diversity and constitutes the main source of innovation in the industry. Major labels systematically monitor this pool and sign the most promising acts once they have demonstrated sufficient traction, a process akin to delegated experimentation.

The coexistence of a dominant core and an atomized periphery is thus a structural feature of the market, shaped by high fixed costs, the need for vertical integration, and the mitigation of contractual frictions along the value chain (Nelson, 1970,Krattenmaker and Salop, 1986).

In this framework, the exponent α of the popularity distribution provides a quantitative measure of how skewed outcomes are between this dominant core and the independent fringe. Lower values of α reflect a steeper hierarchy consistent with strong oligopolistic dominance, whereas higher values indicate a more competitive or fragmented environment.

In economic terms, this reveals a contestable market dominated by a few leaders. This first result may however reveal as much as it may hide, given that Rock is a music genre that covers many subgenres, or even subcultures. Moreover, as seen before, when comparing our dataset with the Spotify ranking of bands, some bands that have a high recognition in their own country do not reach the international scene. Hence, in what follows, we consider the potential existence of such sub-markets, by splitting our sample along genres and country lines. Moreover, to assess whether our results are driven by inactive or legacy bands, the Appendix reports the same estimates restricted to bands that are still active at the time of data collection.

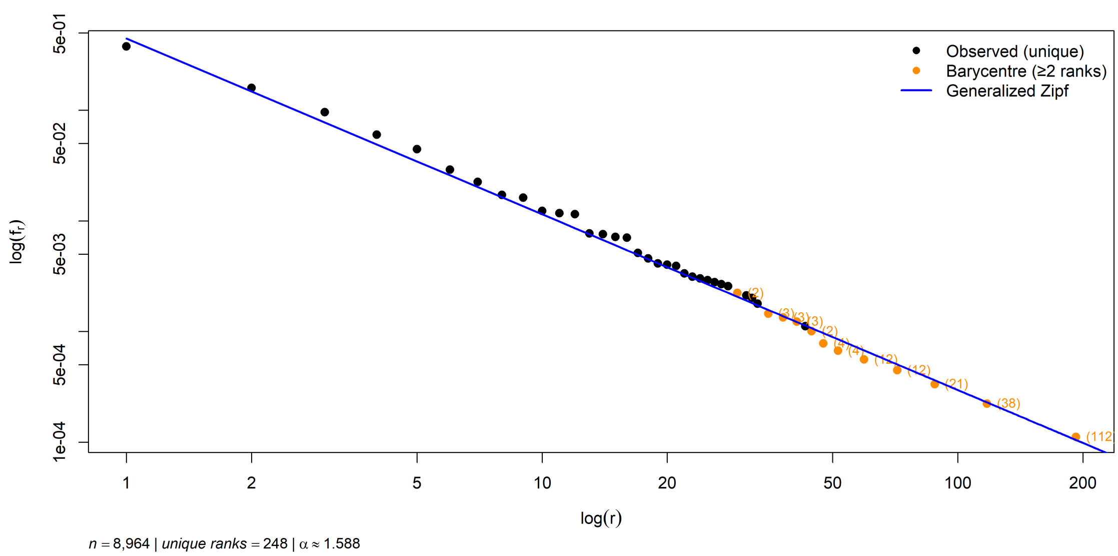

Figure 1: Distribution of the popularity of Rock bands — Full sample

Note: The plotted values correspond to the barycentre of the underlying distribution of ranks (i.e., the mean rank position within the group).

4.2 My Generation: Patterns by Decade of Birth

Figure 2 reports decade-specific Zipf plots for all bands in the database, conditioning the distribution of popularity on the decade of birth of each group. This approach addresses the fact that older bands have simply had more time to produce music, perform, and accumulate followers, which could introduce an age-related bias. By splitting the sample by the decade of birth of each band, we can check if the previous result is related to an "age-bias".[7]

A first striking pattern in Figure 2 is the gradual steepening of the estimated slopes across decades. Bands formed in the 1950s and 1960s display exponents close to one, indicating a highly concentrated distribution in which a small number of groups capture a disproportionate share of declared fans. In contrast, cohorts born in later decades—particularly from the 1980s onward—exhibit increasingly steep slopes, with exponents frequently approaching or exceeding 1.7, revealing a more balanced distribution (see below). This systematic evolution cannot be interpreted without considering a fundamental mechanical feature of the data: older bands have simply had more time to accumulate visibility and followers on the platform.

More precisely, in a cumulative-growth framework (see Gabaix, 2009), the stock of fans is the outcome of repeated exposure opportunities. Groups that entered the market earlier have a longer time horizon over which random growth shocks, performances, recordings, and word-of-mouth interactions can compound. Under standard Gibrat-style dynamics, a longer accumulation window naturally generates heavier upper tails: groups that benefited early from positive shocks have had more time to diverge from the mass of smaller bands, producing a fatter tail and a lower estimated exponent. Conversely, younger bands — those born in the 2000s and later — have had far fewer years in which such cumulative processes can unfold. Even if they were equally promising ex ante, the much shorter time span mechanically compresses the upper tail of their distributions, leading to steeper Zipf slopes.

This interpretation fits squarely within the logic exposed in the first paragraph: the decade split controls for the duration of market presence, allowing us to observe how the distribution behaves when this fundamental source of heterogeneity is held constant. Figure 2 reveals that popularity accumulates over time, and older cohorts naturally exhibit heavier tails. Once this time dimension is neutralized by working within cohorts, later decades appear more fragmented, not because they are inherently more fragmented, but because they have not yet had enough time for the growth process to generate large divergences in

[7] The possibility of dealing with a by-decade analysis has been confirmed by the K-S test, which confirms the validity over the period, except for the very last decades (2000s and 2010s), as the number of ranks (respectively, 23 and 3) is too small.

fan counts.

This cumulative-exposure mechanism provides a parsimonious explanation for the observed monotonic increase in the exponent and aligns with canonical models of proportionate growth. It also clarifies why considering cohorts separately is essential: the raw distribution, mixing bands of vastly different ages, would blur these dynamics and overstate the degree of concentration in earlier decades relative to more recent ones.

Figure 2: Distribution of the popularity of Rock bands - Full sample - Per decade of band birth

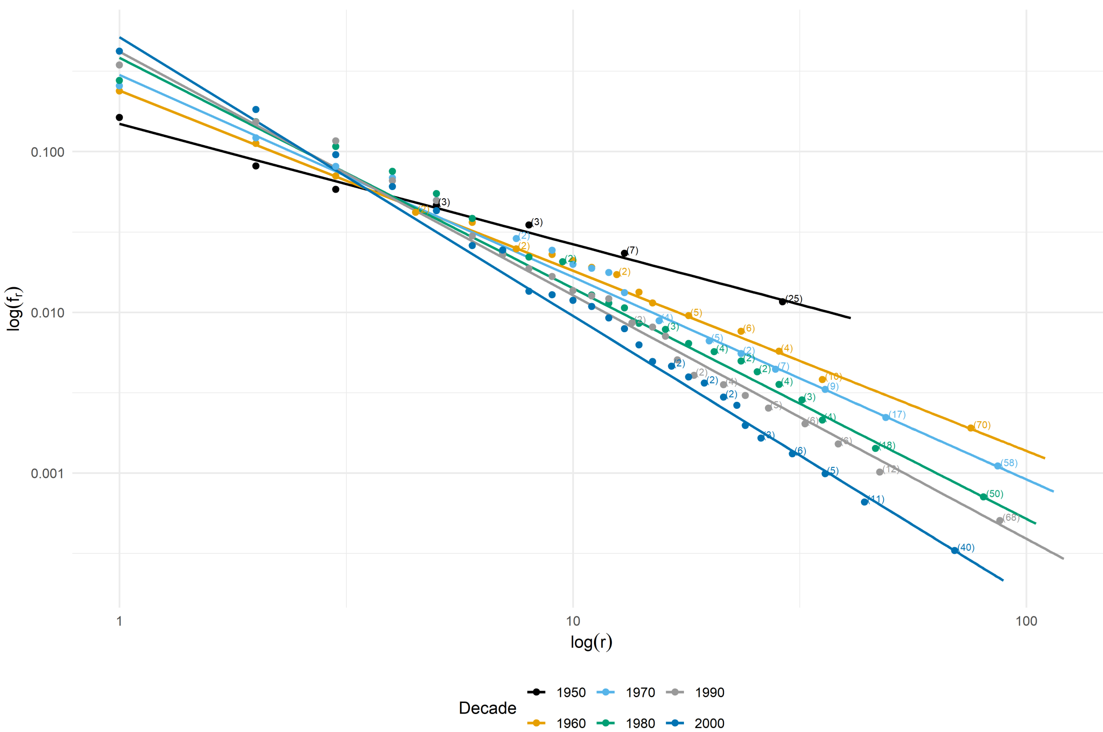

Note: See Figure 1.

A second regularity in Figure 2 is the presence, for all cohorts, of a clear inflection point around 20 to 50 declared fans (a log(rank) between 3 and 4). In log–log space this appears as a discrete change in slope at the lower tail. Because the number of fans recorded on Spirit of Rock reflects only self-declared engagement by users of the platform rather than actual global popularity, this threshold must be interpreted carefully. Even globally famous bands such as AC/DC have fewer than one thousand followers on the site, which underscores that the metric captures platform-specific engagement rather than the true size of an external audience. Under this interpretation, crossing the threshold of 20 to 50 fans signals the emergence of a minimally stable audience: several independent users must have taken the active step of identifying themselves as followers. Bands with lower values are often transient projects, residual profiles, or statistical noise. Once the threshold is reached, however, the band can be said to "exist" on the platform in a meaningful sense.

Notice that this threshold is remarkably stable across decades, despite large differences in the number of bands, the breadth of styles, and the overall density of the musical landscape. The invariance of the inflection point suggests that the minimum level of engaged attention required for a band to "take off" in the data is essentially constant. In the spirit of Gibrat-type growth processes, this can be interpreted as the point beyond which a group escapes the purely stochastic phase of its trajectory. Bands with too small number of fans remain in a region dominated by statistical noise, where observed differences mostly reflect idiosyncratic fluctuations. Reaching the threshold marks the moment when proportionate growth begins to apply meaningfully: only once a band has accumulated this initial nucleus of followers does the multiplicative growth mechanism have room to operate. Under this interpretation, the stability of the inflection point across cohorts reflects a stable "initial size" necessary for Gibrat's law to become a good approximation of subsequent popularity dynamics. Even as the upper tail differs markedly between older and younger decades—largely due to differences in cumulative exposure time—the lower-tail transition into the regime where random growth processes dominate remains fundamentally unchanged.

Altogether, the stable lower-tail threshold combined with the rising exponent across decades is consistent with theoretical models in which technological change increases entry while reducing the extreme dispersion of growth opportunities. The Spirit of Rock data capture this process in a particularly clear form, offering a quantitative window into how the rock landscape has diversified over more than 75 years.

4.3 Enter Subgenres

We split the sample in several parts, using the Spirit of Rock contributors' classification of each band. The genres that we keep are the following: Punk & Ska; Rock; Symphonic Power Rock and Progressive; Pop; Post-Rock; Gothic & Industrial Rock; Jazz, Blues & Fusion; Alternative & Modern Rock; Hard Rock & FM; Folk & Pagan ; Others.

For each subgenre, we analyze the distribution of the bands' popularity among the fans. The results are summarized in Table 1, while Figure 3 shows the equivalent of Figure 1 for the first subgenre (*i.e.*, for Punk and Ska bands). As revealed by Table 1, there is a large discrepancy in the number of bands by subgenre, as well as in the number of the levels of popularity attributed by fans to each band (here converted in ranks). For example, if the Punk/Ska subgenre covers a large number of bands (2058), the number of ranks (85) is relatively similar to the Hard Rock and FM category (79), although the latter contains only 425 bands. The relation between ranks and their distribution thus cannot be attributed to a simple relation with the number of bands present in each category.

A longstanding body of work in economics emphasizes that market structure in differentiated product industries emerges as an adaptation to the distribution of consumer demand. In characteristic-based models (Lancaster, 1971) and in monopolistic competition with differentiated varieties (Dixit and Stiglitz, 1977), firms choose the range and positioning of products in response to heterogeneous preferences, making demand a primary determinant of variety and concentration. Even in frameworks that treat tastes as stable — most notably Stigler and Becker, 1977 — the pattern of preferences across consumers constrains the organization of supply and shapes equilibrium market shares. Moreover, demand-side feed-backs such as network externalities and increasing returns (Arthur, 1989) can amplify small initial differences in visibility or reputation, producing highly skewed outcome distributions. From this perspective, the observed inequality in popularity within a musical genre can be interpreted as an equilibrium response of the supply side to its underlying demand environment: when preferences are globalized, weakly segmented, or highly scale-sensitive, the resulting structure tends to be dominated by a handful of acts, whereas more cohesive or segmented audiences support a more balanced distribution of success.

The analysis of popularity distribution for each subgenre is thus relevant, revealing very different types of distribution, with exponents varying from 1.703 (for Punk/Ska) to 1.221 (for the Hard Rock and FM bands). Figure 4 displays each exponent, allowing for an easy visual inspection of the results. For the Pop Rock and Folk/Pagan categories, the goodness-of-fit test also rejects the generalized Zipf model. While the moderate size of these sub-samples limits the power of the test, the observed rejection nevertheless points to structural heterogeneity within these genres, which encompass distinct sub-scenes and regimes of visibility. As a result, a single Zipf's law appears statistically too restrictive, and the estimated exponents should be interpreted with caution as descriptive indicators of concentration rather than as evidence of fully validated Zipf regimes.

These differences in the estimated exponents across subgenres are not merely descriptive: they significantly impact the industrial structure and competitive dynamics of each segment of the rock music industry. In the industrial organization literature, the degree of inequality in the distribution of popularity - as captured by the exponent α in a power-law relationship - provides a direct measure of concentration and market power. When α is close to 1, the distribution is highly unequal, and the market is dominated by a few "superstar" bands who capture the bulk of attention and demand. This corresponds to a winner-takes-most logic, as theorized by Rosen (1981), where even marginal differences in talent, visibility, or network effects translate into large gaps in success. Conversely, when α is high (e.g., above 1.7), the distribution is more balanced: popularity is spread more evenly across bands, the superstar advantage is mitigated, and the market is more fragmented.

Applying this interpretation to our data, we find evidence of meaningful variation across subgenres. For instance, Punk/Ska shows a relatively high α (1.703), consistent with a more equal distribution and less market concentration. In contrast, Hard Rock & FM displays a much lower exponent (1.090), indicating an extremely skewed distribution of popularity — dominated by a few iconic bands, with limited room for mid-tier competitors, as an inspection of the main festival line-ups across the years would easily confirm. Alternative and Modern Rock, with $\alpha = 1.342$, would fit the superstar model particularly well: a few bands enjoy disproportionate recognition, possibly driven by media exposure, platform algorithms, or - more realistically for these specific subgenres - reputational barriers. However, the obtained exponent is not significant for this genre.

From a strategic and policy standpoint, these patterns thus suggest very differentiated market environments. In low-α segments, competition is intense but asymmetric, entry barriers are high, and innovation tends to be concentrated among dominant players defending their position. In high-α segments, the playing field is flatter, the room for artistic differentiation is broader, and entry is more viable for emerging bands.

Figure 3: Distribution of the popularity of Rock bands – Punk & Ska bands

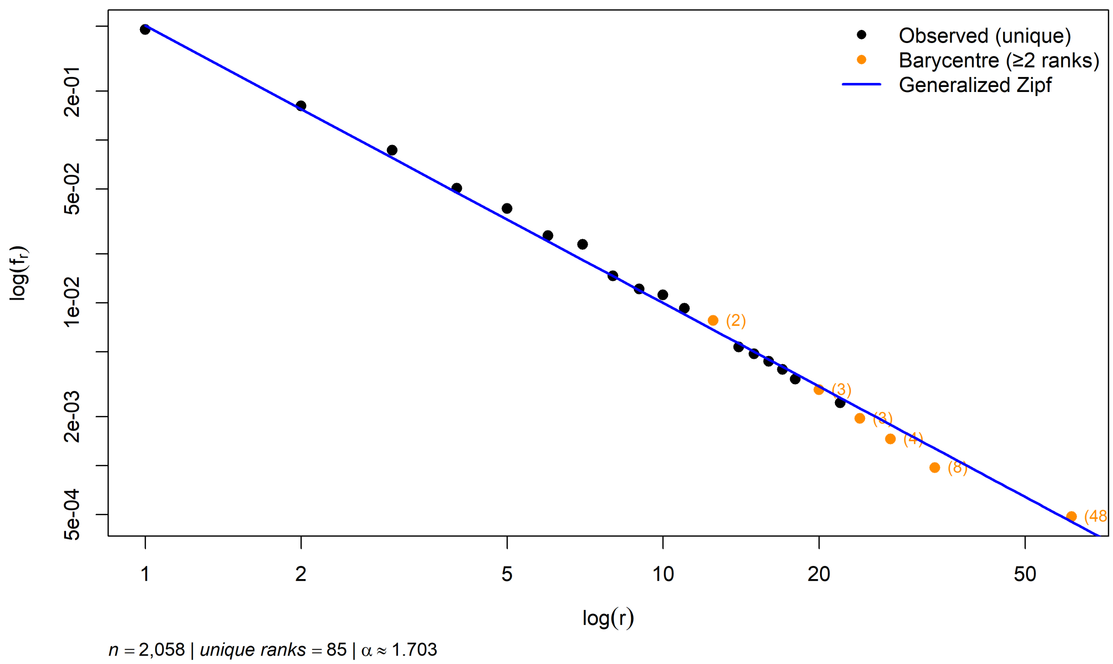

Note: See Figure 1.

4.4 Rock Around the World: Cross-Country Comparisons

It may be considered, in terms of music, that each country is a specific market. This may be either due to language barriers or cultural habits, but many Rock stars in their countries do not benefit from any notoriety in another, even neighboring, one. Many examples can be provided, from France (where Indochine fills stadiums in France since the 1980s, but whose stardom is confined to the French-language market) to Japan, where bands like B'z or Mr. Children have sold millions of records and sold out arenas for decades, yet remain virtually unknown outside the Japanese domestic market due to language barriers and limited international promotion.

Although this might suggest that each national market exhibits its own distinct popularity structure, disconnected from the global pattern, our estimates of the slope (and the associated exponent α) indicate that band rankings nonetheless follow a power-law distribution. Figure 5 displays the exponent for each country for which the number of observations (i.e., ranks) is large enough to allow for such an estimate. As can be seen, although each country has a ranking of popularity that conforms to a power-law, striking differences also exist. Australia,

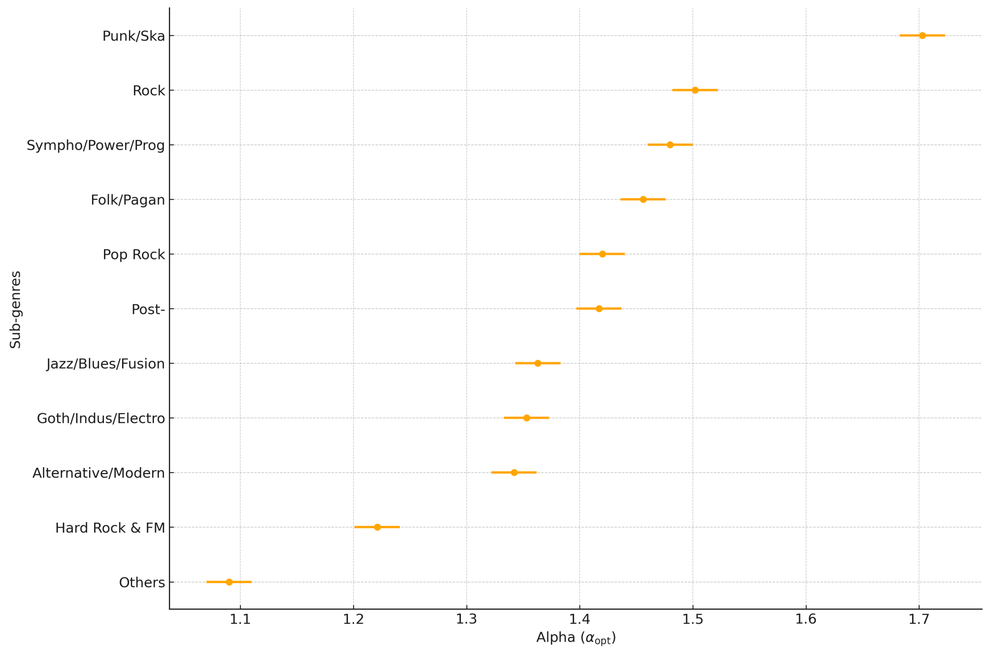

Figure 4: Power-law exponents — Sub-genres
Notes: Each dot represents the estimated optimal exponent α for a given rock sub-genre. Horizontal bars indicate a proxy for uncertainty, proportional to $(1 - |r_s|)$, where r_s is the Spearman correlation measuring the goodness of fit of the power-law. Larger bars reflect weaker fit quality. Sub-genres are sorted by their estimated α.

for example, has an α=1.225, very different from Japan's value (α=1.763) - see Table 2.

Here too, the variation in estimated power-law exponents across countries reflects important underlying differences in how national music markets are structured and, consequently, how recognition is distributed among rock bands. Interpreting again the exponent α of a power-law distribution as a proxy for the concentration of attention or demand, lower values of α imply steeper hierarchies, where a few bands attract the vast majority of interest, while higher values suggest more diffuse systems in which popularity is more evenly shared.

In this perspective, the observed cross-country differences are striking. Countries such as the United Kingdom ($\hat{\alpha} = 1.324$) and the United States ($\hat{\alpha} = 1.444$) - arguably the historical cores of the global rock music industry - display relatively low exponents, compared to the other countries. However, it is important to note that, in the case of the United States, the goodness-of-fit test rejects the generalized Zipf model. Given the large size and pronounced internal heterogeneity of the U.S. market, this rejection is consistent with the high statistical power of the test and indicates that a single Zipf's law is too restrictive to capture the full complexity of the distribution. In this context, the estimated exponent $\hat{\alpha}$ should be interpreted as a descriptive indicator of the overall degree of concentration, rather than as

Table 1: Power Law Estimates by Genre

Genre	n_{groups}	Exponent ($\hat{\alpha}$)	$\text{MSE}_{\log}$	n_{ranks}	p-value
Punk/Ska	2058	1.703	0.104	85	0.663
Rock	2040	1.502	0.053	118	0.188
Symphonic/Power/Prog	518	1.480	0.095	55	0.940
Goth/Indus/Electro	449	1.472	0.058	48	0.298
Folk/Pagan	463	1.456	0.067	48	0.007
Pop Rock	795	1.420	0.101	81	0.017
Post-	523	1.417	0.058	58	0.288
Jazz/Blues/Fusion	454	1.363	0.074	61	0.756
Alternative/Modern	358	1.342	0.123	57	0.598
Hard Rock & FM	425	1.221	0.091	79	0.673
Others	225	1.090	0.057	58	0.961

Notes. The exponent α is estimated by minimizing the mean squared error (MSE) between observed and theoretical frequencies on a log–log scale. The number of "ranks" refers to the distinct popularity levels observed within each subgenre. Approximate standard errors for $\hat{\alpha}$ are computed using the analytic formula $\text{SE}(\hat{\alpha}) = (\hat{\alpha} - 1)/\sqrt{n_{\text{groups}}}$, and the confidence intervals reported in the figures correspond to $\hat{\alpha} \pm 1.64\,\text{SE}(\hat{\alpha})$, i.e. an approximate 90% interval in the sense of Clauset *et al.* (2009). The reported p-values are obtained from a bootstrap–corrected Kolmogorov–Smirnov test (Clauset *et al.*, 2009), which assesses the goodness of fit of the power-law specification. High p-values (e.g., above 0.1) indicate that the empirical distribution is statistically consistent with a generalized Zipf's law, while low values point to a poor fit.

evidence of a fully validated Zipf regime.

This is consistent with highly concentrated popularity landscapes, where a handful of canonical bands dominate the national (and even international, in their case) markets. Australia shows an even more pronounced pattern $\hat{\alpha} = 1.225$, suggesting an even narrower funnel through which rock stardom is channeled, perhaps due to a smaller domestic industry or stronger winner-takes-all dynamics in an overseas market.[8]

At the other end of the spectrum, Japan stands out with the highest exponent in the sample ($\hat{\alpha} = 1.763$), pointing to a more balanced distribution of visibility and success. This is consistent with the fragmentation often noted in the Japanese music market, where genre-specific fan bases and strong domestic loyalty sustain a wider array of bands, even in the absence of international recognition. Intermediate cases include France ($\hat{\alpha} = 1.575$), Germany ($\hat{\alpha} = 1.592$), and Canada ($\hat{\alpha} = 1.404$), where the distribution is less skewed than in Anglo-American markets but still reflects a hierarchical structure of popularity.

These patterns likely reflect not only differences in market size or industrial concentration

[8]The presence of AC/DC in the database, which is the band with the highest number of fans, does not impact much the results. Dropping the band delivers an = 1.174, hence does not substantially modify the results, which is logical as our model relies on the number of ranks, not directly on the number of fans.

but also broader institutional and cultural factors such as the strength of national media ecosystems, the segmentation of audiences by language, the structure of live performance circuits, and the extent of public support for the arts. While our data do not allow us to disentangle these mechanisms empirically, some systematic contrasts are suggestive. Countries with strong linguistic identities and relatively autonomous domestic music markets (Japan, France, Germany, and Sweden) tend to display higher exponents, consistent with more balanced distributions in which attention is spread across a wider set of national acts. By contrast, English-speaking countries (the United States, the United Kingdom, Canada, and Australia) exhibit lower exponents, a pattern compatible with the pressures of competing directly in a large, globally integrated English-language market where a small number of internationally dominant acts capture substantial attention.

These results can also be interpreted in terms of entry conditions for emerging artists. In markets where the popularity distribution is extremely concentrated, the path to visibility is likely to be mediated by powerful gatekeepers such as major labels, large promotion networks, or dominant streaming platforms. Under these conditions, small initial differences in exposure may be amplified through cumulative advantage mechanisms, making it difficult for new entrants to reach a critical threshold of audience attention. By contrast, environments characterized by flatter popularity distributions may offer more room for gradual recognition, allowing a larger number of artists to build sustainable audiences without necessarily reaching superstar status. In this sense, the estimated scaling exponent can be interpreted not only as a statistical descriptor of popularity but also as an indirect indicator of how accessible different cultural markets are to new creative entrants.

Table 2: Power-law fit diagnostics by country

#	Country	n_{groups}	Exponent ($\hat{\alpha}$)	$\text{MSE}_{\log}$	n_{ranks}	p-value
1	Japan	402	1.763	0.074	28	0.216
2	Germany	469	1.592	0.051	39	0.7
3	Sweden	251	1.576	0.080	28	0.570
4	France	1161	1.575	0.042	70	0.620
5	United States	2843	1.444	0.043	167	0
6	Canada	343	1.404	0.056	45	0.886
7	United Kingdom	1411	1.324	0.084	150	0.512
8	Australia	167	1.225	0.061	37	0.934

Notes. α and the mean squared error (MSE) are defined as in Table 1. The Spearman coefficient measures the monotonic association between empirical ranks and theoretical expectations. The number of "ranks" corresponds to the distinct fan-count strata observed in each country. As in Table 1, the p-values are based on a bootstrap–corrected Kolmogorov–Smirnov statistic, providing a diagnostic of the adequacy of the power-law model for each national distribution.

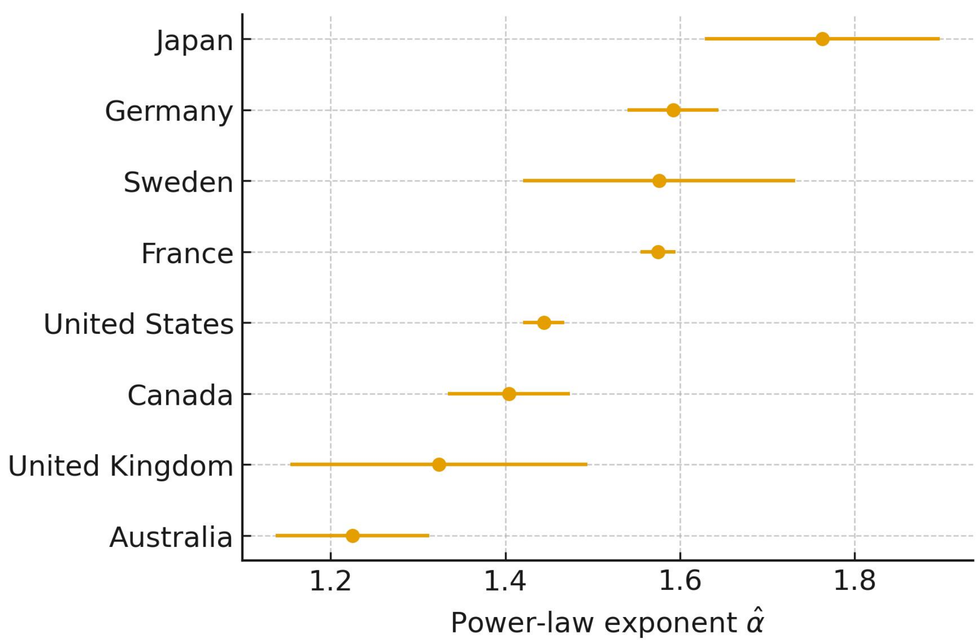

Figure 5: Power-law exponents — Countries

Notes: Each dot represents the estimated optimal exponent α for a given country. Horizontal bars indicate a proxy for uncertainty, proportional to $(1 - |r_s|)$, where r_s is the Spearman correlation measuring the goodness of fit of the power-law. Larger bars reflect weaker fit quality. Countries are sorted by their estimated α.

5 Conclusion

In this paper, we investigate the popularity dynamics of rock bands based on the number of fans. We have shown that the general model follows a power-law distribution, typical of a generalized Zipf's law. Generally speaking, the rock music market is dominated by a few major bands that attract large numbers of fans, followed by many bands with much less popularity. In terms of industrial economics, the popularity of rock bands can be described as a fringe oligopoly: a small number of companies capture a large share of demand, while a long tail of smaller firms share the crumbs.

Although this market structure is true for all subgenres and all countries, there are nuances in the length of the tail and/or the power of the major bands. As for subgenres, some follow the same distribution as in the general case (*e.g.*, Alternative and Modern Rock and Hard Rock/FM): specific genres in which major bands probably benefit from greater media exposure, thereby reinforcing the reputational barriers - playing as a ratchet effect - that protect them. Other subgenres such as symphonic/power/prog rock or, even more notably, ska/punk are characterized by a higher value of the exponent α, and therefore a less concentrated market. These less mainstream musical styles reflect audiences that are probably more knowledgeable and specialized. Within these 'niche' markets, competition is fiercer and it is more difficult for individual bands to stand out on the international stage.

Taken together, these results suggest that the exponent of the popularity distribution provides a simple quantitative indicator of how concentrated cultural markets are and how accessible they may be to new entrants.

Finally, the model also highlighted idiosyncratic dynamics when analyzing cross-country data. For some countries, there is clear evidence of a Zipf distribution (US, UK, Australia). In this set of countries, only a few rock giants dominate the national market, enjoying significant popularity returns and imposing stronger barriers to smaller emerging outliers. Other countries exhibit higher fragmentation of their rock music market (*e.g.* Japan, France, Germany)

Appendix

A Robustness Check: Active Bands Only

We replicate the main analysis restricting the sample to bands that are still active at the time of data collection. This restriction addresses potential survivorship concerns and allows us to assess whether the observed power-law patterns are driven by historical accumulation of popularity among inactive or defunct acts. The results are robust, with correlation levels between each set of results being never inferior to 0.98.

A.1 Cross-Country Results

Table 3: Power Law Estimates by Country: Full Sample vs Active Bands

Country	**Exponent ($\hat{\alpha}$), full sample**	**Exponent ($\hat{\alpha}$), active bands**
Japan	1.76	1.68
Germany	1.59	1.59
Sweden	1.58	1.61
France	1.57	1.56
United States	1.44	1.44
Canada	1.40	1.39
United Kingdom	1.32	1.30
Australia	1.23	1.21

A.2 Sub-Genre Results

Table 4: Power Law Estimates by Genre: Full Sample vs Active Bands

Genre	**Exponent ($\hat{\alpha}$), full sample**	**Exponent ($\hat{\alpha}$), active bands**
Punk/Ska	1.70	1.70
Rock	1.50	1.50
Sympho/Power/Prog	1.48	1.42
Goth/Indus/Electro	1.47	1.42
Folk/Pagan	1.46	1.47
Pop Rock	1.42	1.41
Post-	1.42	1.41
Jazz/Blues/Fusion	1.36	1.40
Alternative/Modern	1.34	1.33
Hard Rock & FM	1.22	1.20
Others	1.09	1.06

A.3 Cohort Results

Table 5: Power Law Estimates by Decade of Band Formation: Full Sample vs Active Bands

Decade	**Exponent ($\hat{\alpha}$), full sample**	**Exponent ($\hat{\alpha}$), active bands**
2020	2.86	2.86
2010	2.20	2.20
2000	1.73	1.73
1990	1.52	1.47
1980	1.43	1.35
1970	1.26	1.17
1960	1.12	0.991

B Top-Ranked Bands by Popularity

For completeness, we here report the ten most popular bands according to the number of fans recorded in the Spirit of Rock Encyclopedia. Popularity is measured consistently with the main analysis and corresponds to the total number of registered fans per band.

B.1 Top 10 Bands in the Full Sample

Table 6: Top 10 Bands in the Full Sample

Name of the band	**Country**	**Genre**
Pink Floyd	United Kingdom	Sympho/Power/Prog
Queen	United Kingdom	Rock
Green Day	United States	Punk/Ska
Nirvana	United States	Others
Muse	United Kingdom	Rock
The Offspring	United States	Punk/Ska
Red Hot Chili Peppers	United States	Jazz/Blues/Fusion
Sum 41	Canada	Punk/Ska
The Rolling Stones	United Kingdom	Rock
AC–DC	Australia	Hard Rock & FM

B.2 Top 10 Bands by Selected Subgroups

Tables reporting the top-ranked bands by major subgenres and countries are available upon request or can be found in the online supplementary material.

Table 7: Top 10 Bands by Genre

Genre	**Band**	**Country**
Alternative/Modern	Foo Fighters	United States
Alternative/Modern	My Chemical Romance	United States
Alternative/Modern	30 Seconds To Mars	United States
Alternative/Modern	The Smashing Pumpkins	United States
Alternative/Modern	REM	United States
Alternative/Modern	Three Days Grace	Canada
Alternative/Modern	Nickelback	Canada

Continued on next page

Genre	Band	Country
Alternative/Modern	Rise Against	United States
Alternative/Modern	The Pixies	United States
Alternative/Modern	Kings of Leon	United States
Others	Nirvana	United States
Others	Alice In Chains	United States
Others	Bad Religion	United States
Others	Pearl Jam	United States
Others	The Exploited	United Kingdom
Others	Soundgarden	United States
Others	Evanescence	United States
Others	Dead Kennedys	United States
Others	Anti-Flag	United States
Others	Creed	United States
Folk/Pagan	Bob Dylan	United States
Folk/Pagan	Jethro Tull	United Kingdom
Folk/Pagan	Dropkick Murphys	United States
Folk/Pagan	Louise Attaque	France
Folk/Pagan	Blackmore's Night	United Kingdom
Folk/Pagan	Flogging Molly	United States
Folk/Pagan	America	United States
Folk/Pagan	Les Ramoneurs de Menhirs	France
Folk/Pagan	The Pogues	United Kingdom
Folk/Pagan	Simon and Garfunkel	United States
Goth/Industrial/Electronic	The Cure	United Kingdom
Goth/Industrial/Electronic	The Sisters of Mercy	United Kingdom
Goth/Industrial/Electronic	The Prodigy	United Kingdom
Goth/Industrial/Electronic	Massive Attack	United Kingdom
Goth/Industrial/Electronic	Dead Can Dance	Australia
Goth/Industrial/Electronic	BlutEngel	Germany
Goth/Industrial/Electronic	The 69 Eyes	Finland
Goth/Industrial/Electronic	Wednesday 13	United States
Goth/Industrial/Electronic	Fields of the Nephilim	United Kingdom

Continued on next page

Genre	Band	Country
Goth/Industrial/Electronic	Alien Sex Fiend	United Kingdom
Hard Rock & FM	AC–DC	Australia
Hard Rock & FM	Led Zeppelin	United Kingdom
Hard Rock & FM	Deep Purple	United Kingdom
Hard Rock & FM	Guns N' Roses	United States
Hard Rock & FM	Scorpions	Germany
Hard Rock & FM	Aerosmith	United States
Hard Rock & FM	Kiss	United States
Hard Rock & FM	Alice Cooper	United States
Hard Rock & FM	Motörhead	United Kingdom
Hard Rock & FM	Van Halen	United States
Jazz/Blues/Fusion	Red Hot Chili Peppers	United States
Jazz/Blues/Fusion	Eric Clapton	United Kingdom
Jazz/Blues/Fusion	Janis Joplin	United States
Jazz/Blues/Fusion	Santana	United States
Jazz/Blues/Fusion	Rory Gallagher	Ireland
Jazz/Blues/Fusion	Stevie Ray Vaughan	United States
Jazz/Blues/Fusion	Shaka Ponk	France
Jazz/Blues/Fusion	B.B. King	United States
Jazz/Blues/Fusion	Joe Cocker	United Kingdom
Jazz/Blues/Fusion	Ten Years After	United Kingdom
Pop Rock	The Beatles	United Kingdom
Pop Rock	Coldplay	United Kingdom
Pop Rock	Simple Plan	Canada
Pop Rock	Indochine	France
Pop Rock	Oasis	United Kingdom
Pop Rock	Fall Out Boy	United States
Pop Rock	Avril Lavigne	Canada
Pop Rock	John Lennon	United Kingdom
Pop Rock	Toto	United States
Pop Rock	Blondie	United States
Post-Rock	Depeche Mode	United Kingdom

Continued on next page

Genre	Band	Country
Post-Rock	Joy Division	United Kingdom
Post-Rock	Bauhaus	United Kingdom
Post-Rock	The Red Jumpsuit Apparatus	United States
Post-Rock	Siouxsie and the Banshees	United Kingdom
Post-Rock	Eurythmics	United Kingdom
Post-Rock	Simple Minds	United Kingdom
Post-Rock	Duran Duran	United Kingdom
Post-Rock	Archive	United Kingdom
Post-Rock	Tears for Fears	United Kingdom
Punk/Ska	Green Day	United States
Punk/Ska	The Offspring	United States
Punk/Ska	Sum 41	Canada
Punk/Ska	Blink 182	United States
Punk/Ska	Sex Pistols	United Kingdom
Punk/Ska	The Ramones	United States
Punk/Ska	The Clash	United Kingdom
Punk/Ska	Good Charlotte	United States
Punk/Ska	Iggy Pop	United States
Punk/Ska	NOFX	United States
Rock	Queen	United Kingdom
Rock	Muse	United Kingdom
Rock	The Rolling Stones	United Kingdom
Rock	U2	Ireland
Rock	Jimi Hendrix	United States
Rock	The Doors	United States
Rock	The Who	United Kingdom
Rock	Dire Straits	United Kingdom
Rock	The Police	United Kingdom
Rock	Noir Désir	France
Symphonic/Power/Progressive	Pink Floyd	United Kingdom
Symphonic/Power/Progressive	Genesis	United Kingdom
Symphonic/Power/Progressive	Yes	United Kingdom

Continued on next page

Genre	Band	Country
Symphonic/Power/Progressive	Supertramp	United Kingdom
Symphonic/Power/Progressive	King Crimson	United Kingdom
Symphonic/Power/Progressive	Marillion	United Kingdom
Symphonic/Power/Progressive	Mike Oldfield	United Kingdom
Symphonic/Power/Progressive	Porcupine Tree	United Kingdom
Symphonic/Power/Progressive	Asia	United Kingdom
Symphonic/Power/Progressive	Camel	United Kingdom

Table 8: Top 10 Bands by Country

Country	**Band**	**Genre**
Germany	Scorpions	Hard Rock & FM
Germany	Kraftwerk	Post-Rock
Germany	Tokio Hotel	Pop Rock
Germany	BlutEngel	Goth/Industrial/Electronic
Germany	Tangerine Dream	Rock
Germany	Nina Hagen	Punk/Ska
Germany	Cinema Bizarre	Rock
Germany	Eloy	Symphonic/Power/Progressive
Germany	Die Toten Hosen	Punk/Ska
Germany	Die Ärzte	Punk/Ska
Australia	AC–DC	Hard Rock & FM
Australia	Airbourne	Hard Rock & FM
Australia	Dead Can Dance	Goth/Industrial/Electronic
Australia	Silverchair	Others
Australia	Nick Cave and the Bad Seeds	Rock
Australia	Jet	Rock
Australia	Wolfmother	Doom/Sludge/Stoner
Australia	Texas	Hard Rock & FM
Australia	The Vines	Rock
Australia	Pendulum	Goth/Industrial/Electronic
Canada	Sum 41	Punk/Ska
Canada	Simple Plan	Pop Rock
Canada	Avril Lavigne	Pop Rock
Canada	Three Days Grace	Alternative/Modern
Canada	Neil Young	Rock
Canada	Billy Talent	Punk/Ska
Canada	Nickelback	Alternative/Modern
Canada	Rush	Hard Rock & FM
Canada	Bryan Adams	Rock

Continued on next page

Country	Band	Genre
Canada	Silverstein	Emo/Screamo
United States	Green Day	Punk/Ska
United States	Nirvana	Others
United States	The Offspring	Punk/Ska
United States	Red Hot Chili Peppers	Jazz/Blues/Fusion
United States	Jimi Hendrix	Rock
United States	The Doors	Rock
United States	Blink 182	Punk/Ska
United States	The Ramones	Punk/Ska
United States	Foo Fighters	Alternative/Modern
United States	My Chemical Romance	Alternative/Modern
France	Noir Désir	Rock
France	Indochine	Pop Rock
France	Téléphone	Rock
France	Bérurier Noir	Punk/Ska
France	Shaka Ponk	Jazz/Blues/Fusion
France	Manu Chao	Rock
France	Mano Negra	Alternative/Modern
France	Superbus	Pop Rock
France	Louise Attaque	Folk/Pagan
France	Damien Saez	Rock
Japan	Maximum the Hormone	Thrash
Japan	Asian Kung-Fu Generation	Pop Rock
Japan	Arc˜en˜Ciel	Rock
Japan	Yousei Teikoku	Goth/Industrial/Electronic
Japan	Mono	Post-Rock
Japan	One Ok Rock	Rock
Japan	M	Others
Japan	Gossip	Others
Japan	Anna Tsuchiya	Punk/Ska
Japan	Mika Nakashima	Pop Rock
United Kingdom	Pink Floyd	Symphonic/Power/Progressive
United Kingdom	Queen	Rock
United Kingdom	Muse	Rock
United Kingdom	The Rolling Stones	Rock
United Kingdom	The Beatles	Pop Rock
United Kingdom	Sex Pistols	Punk/Ska
United Kingdom	Led Zeppelin	Hard Rock & FM
United Kingdom	The Who	Rock
United Kingdom	Dire Straits	Rock
United Kingdom	The Police	Rock
Sweden	The Hives	Punk/Ska
Sweden	Europe	Heavy/Speed

Continued on next page

Country	Band	Genre
Sweden	Blues Pills	Rock
Sweden	Roxette	Pop Rock
Sweden	Millencolin	Punk/Ska
Sweden	The Flower Kings	Symphonic/Power/Progressive
Sweden	Hardcore Superstar	Rock
Sweden	Spiritual Beggars	Doom/Sludge/Stoner
Sweden	The Cardigans	Rock
Sweden	Änglagård	Symphonic/Power/Progressive

References

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, *99*(394), 116–131.

Axtell, R. L. (2001). Zipf distribution of us firm sizes. *Science*, *293*(5536), 1818–1820.

Brock, W. (1999). Scaling in economics: A reader's guide. *Industrial and Corporate Change*, *8*(3), 409–446. https://doi.org/10.1093/icc/8.3.409

Clauset, A., Shalizi, C. R., & Newman, M. E. (2009). Power-law distributions in empirical data. *SIAM Review*, *51*(4), 661–703.

Connolly, M., & Krueger, A. B. (2006). Rockonomics: The economics of popular music. In V. Ginsburgh & D. Throsby (Eds.), *Handbook of the economics of art and culture* (pp. 667–719, Vol. 1). Elsevier.

del Rosal, I. (2018). Power laws in eu country exports. *Empirica*, *45*, 311–337.

Dixit, A. K., & Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *American Economic Review*, *67*(3), 297–308.

Etro, F., & Stepanova, E. (2018). Power-laws in art. *Physica A: Statistical Mechanics and its Applications*, *506*, 217–220.

Farvaque, E. (2024). For those about to rock… is stability a determinant of rock bands' success? *Journal of Cultural Economics*, *48*(1), 145–166.

Gabaix, X. (1999). Zipf's law for cities: An explanation. *The Quarterly Journal of Economics*, *114*(3), 739–767.

Gabaix, X. (2009). Power laws in economics and finance. *Annual Review of Economics*, *1*(1), 255–294.

Gabaix, X. (2016). Power laws in economics: An introduction. *Journal of Economic Perspectives*, *30*(1), 185–206.

Gabaix, X., Landier, A., & Sauvagnat, J. (2014). Ceo pay and firm size: An update after the crisis. *The Economic Journal*, *124*(574), F40–F59.

Jerez-Lillo, N., Rodrigues, F., Ferreira, P. H., & Ramos, P. L. (2025). *Beyond the power law: Estimation, goodness-of-fit, and a semiparametric extension in complex networks* [Manuscript under review].

Kapeller, J., & Steinerberger, S. (2025). *Why are there so many power laws in economics?* [ICAE Working Paper Series No. 160, Johannes Kepler University Linz].

Krattenmaker, T. G., & Salop, S. C. (1986). Anticompetitive exclusion: Raising rivals' costs to achieve power over price. *Yale Law Journal*, *96*(2), 209–293. https://doi.org/10.2307/796483

Krueger, A. B. (2005). The economics of real superstars: The market for rock concerts in the material world. *Journal of Labor Economics*, *23*(1), 1–30.

Lancaster, K. J. (1971). *Consumer demand: A new approach*. Columbia University Press.

Mandelbrot, B. B. (1953). An informational theory of the statistical structure of language. *Communication theory*, *84*, 486–502.

McKenzie, J., Crosby, P., & Lenten, L. J. A. (2021). It takes two, baby! feature artist collaborations and streaming demand for music. *Journal of Cultural Economics*, *45*, 385–408.

Mitzenmacher, M. (2004). A Brief History of Generative Models for Power Law and Lognormal Distributions. *Internet Mathematics*, *1*(2). https://doi.org/10.1080/15427951.2004.10129088

Nelson, P. (1970). Information and consumer behavior. *Journal of Political Economy*, *78*(2), 311–329. https://doi.org/10.1086/259630

Newman, M. (2005). Power laws, pareto distributions and zipf's law. *Contemporary Physics*, *46*(5), 323–351. https://doi.org/10.1080/00107510500052444

Prinz, A. L. (2022). The concentration of power in the market for contemporary art: An empirical analysis of artreview's "power 100". *SN Business & Economics*, *2*, 11.

Rosen, S. (1981). The economics of superstars. *The American Economic Review*, *71*(5), 845–858.

Simon, H. A. (1955). On a class of skew distribution functions. *Biometrika*, *42*(3-4), 425–440.

Stigler, G. J., & Becker, G. S. (1977). De gustibus non est disputandum. *American Economic Review*, *67*(2), 76–90.

Urzúa, C. M. (2020). A simple test for power-law behavior. *The Stata Journal*, *20*(3), 604–612.

Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Addison-Wesley.