

ASSESSING THE IMPACT OF SAMPLING, REMIXES, AND COVERS ON ORIGINAL SONG POPULARITY

Guilherme Soares S. dos Santos Flavio Figueiredo

Universidade Federal de Minas Gerais

{guilhermesoares, flaviovd} @dcc.ufmg.br

ABSTRACT

Music digitalization has introduced new forms of composition known as “musical borrowings”, where composers use elements of existing songs—such as melodies, lyrics, or beats—to create new songs. Using Who Sampled data and Google Trends, we examine how the popularity of a borrowing song affects the original. Employing Regression Discontinuity Design (RDD) for short-term effects and Granger Causality for long-term impacts, we find evidence of causal popularity boosts in some cases. *Borrowee* songs can revive interest in older tracks, underscoring economic dynamics that may support fairer compensation in the music industry.

1. INTRODUCTION

Digitization has drastically transformed the production, distribution, and consumption of music. Advances in data storage, transmission, and processing technologies have brought about significant changes in the music industry, making digital music predominant. This shift has led to new forms of composition, known as “musical borrowings”, where composers and producers draw on existing musical influences to create new works, thereby expanding the boundaries of musical genres.

An example of musical borrowing is Beyoncé’s 2003 hit “Crazy in Love”, one of her most iconic songs. This work calls this song a *borrowee*. The *borrowee*’s introduction, marked by triumphant brass, drew significant attention. The brass section was sampled from The Chi-Lites’ 1970s song “Are You My Woman”, here called a *borrowed* song. Created in a pre-digital era, the *borrowed*’s influence on “Crazy in Love” is evident, although its precise contribution to the song’s success remains difficult to quantify.

We explore three primary forms of musical borrowing: sampling, remixes, and covers. Sampling, a technique popularized in the late 1970s, involves using fragments of pre-existing songs. Widely used in hip-hop, electronic, and experimental music, sampling allows for creative expression through the manipulation and combination of different

tracks. On the other hand, remixes involve transforming an existing song by altering its structure, rhythm, or effects, often revitalizing it for new audiences. Finally, covers involve re-recordings that usually offer a fresh interpretation and sometimes even surpass the original in popularity.

Musical borrowings raise significant legal issues, particularly regarding the copyright laws that govern these practices – debates around fair use often center on whether the borrowing negatively impacts the market for the original work. A study titled “Sampling Increases Music Sales: An Empirical Copyright Study” highlights that, in the U.S., the Supreme Court considers the market impact on the original song as a key factor in determining fair use [1]. The research suggests that digital sampling can boost sales of the original songs, indicating that musical borrowings may, in some cases, qualify as fair use. Understanding the consequences of using samples in original songs can contribute to fairer legal discussions and more informed decisions by musicians and managers.

Our research aligns with the above discussion, but it also presents counter-arguments, as we now detail. Specifically, we focus on how the popularity of a *borrowee* song (cover, sample, or remix) influences the popularity of the *borrowed* song. Data from WhoSampled^{1 2}, a crowd-sourced website cataloging samples, remixes, and covers, and Google Trends³, which measures relative search interest, provide the basis for this analysis.

To effectively measure an effect, we employ two techniques. First, we explore Regression Discontinuity Design [2] (RDD) to measure the immediate causal impact of the *borrowee*’s release. Next, we make use of Granger Causality [3] as a means to understand the lasting impact of a *borrower* after release. As an example of a *borrowee* that impacts a *borrowed* song, we show in Figure 1 the song “Somebody” by Natalie La Rose (released in late 2014) that sampled “Shots” from LMFAO. Here, we can see that “Somebody” renews interest in the LMFAO song.

In general, it is important to note that establishing causality is challenging. Among the approximately 884 instances of musical borrowings analyzed, we identified 182 cases (20%) showing any causal evidence. From these, our results show that only a fraction of the *borrowees* (45% or 82 out of 182) have an immediate causal impact via our RDD study. Out of these, 64% (or 117 out of 182) of these



¹ <https://www.whosampled.com>.

² Our dataset is available at: github.com/uai-ufmg/whosampleddata

³ <https://trends.google.com>




Figure 1: Search Interest Over Time for “Shots” (LMFAO) and “Somebody” (Natalie La Rose). The release month is the shaded region.

cases show a lasting effect via Granger Causality. It should be emphasized that this analysis focuses on web search interest, rather than sales metrics. Nonetheless, even with a relatively limited dataset, we observed some influence of *borrowees* on the originals. Thus, we hope that our research sheds light on the actual implications of sampling.

In the following sections, we begin with a review of related work (Section 2), followed by a description of our datasets and methodology (Section 3) and our results (Section 4). Finally, we conclude with a discussion of our findings and their broader implications (Section 5).

2. RELATED WORK

The term “musical borrowings” defined by Burkholder [4], refers to the techniques composers use to create music based on pre-existing melodies or structures. This practice has been studied not only in music but also in legal contexts, especially regarding copyright. Scholars have highlighted discrepancies between copyright law and the realities of musical composition. Arewa [5] critiques copyright law’s focus on “romanticized authorship”. This assumes originality and autonomy, disregarding musical borrowing, particularly in genres like hip-hop, where sampling has raised legal concerns.

Musical borrowing also intersects with social inequalities. Hesmondhalgh [6] explores how white musicians have historically appropriated works by Black artists, reinforcing cultural disparities. This issue extends to non-hegemonic markets, such as Latin America and Africa, where musical borrowing can have broader implications.

Analyzing musical influence presents challenges due to its subjective nature. Shalit et al. [7] tackled this by applying Dynamic Topic Models to a dataset of over 24,000 songs from 1922–2010, revealing that musical influence and innovation are not monotonically correlated. The study identified critical periods of high innovation in the early 1970s and mid-1990s.

Turning our attention to authors exploring a dataset similar to ours, Bryan et al. [8] also explore WhoSampled data to analyze to characterize networks. The study examined centrality in a network of 42,447 samples, applying complex network techniques and power-law distributions

to measure the influence of songs and artists. Katz Centrality was introduced as a new method to quantify musical influence, though the author’s study did not directly address how sampling affects popularity.

Ortega [9] analyzed cover versions using a network of over 106,000 artists and 855,000 covers from Second-Hand Songs, focusing on the impact of covers on popular music history. Key findings showed genre and language as primary factors in collaboration and influence.

3. METHODOLOGY AND DATA COLLECTION

Our work focuses on analyzing the popularity of songs involved in musical borrowings over time to measure the impact of the release of a new version on the original version. To do so, we used a dataset obtained from Who Sampled containing about 700,000 songs, including samples, covers, and remixes. This dataset was crawled in 2019, and we asked WhoSampled for permission to crawl those data. The website catalogs musical borrowings for more than 975,000 songs from 300,000 artists, contributed by a community of about 28,000 members and verified by moderators or staff⁴. In this sense, our original dataset comprises most of these songs.

The obtained data is structured as a directed graph, where nodes represent songs and edges represent musical borrowings, each labeled to indicate the type of musical borrowing. Since Who Sampled does not provide popularity data or external identifiers, we supplemented the data with Wikidata⁵ and Google Trends information.

Initially, we resorted to Wikidata to gather unique identifiers for each song. The song “Clube da Esquina vol. 2” by Milton Nascimento has a Wikidata identifier of Q20053386. More importantly, on Wikidata, this song also has its Freebase⁶ identifier listed: /m/0zjw3z... This Freebase identifier was later used for Google Trends queries. Before describing how we collected time series, we discuss how we gathered the Freebase id of our songs.

To find song identifiers, we used Wikidata with three search queries: one with the song title, one with the artist name, and one combining both. For each, we retrieved the top ten entities, removed duplicates, and processed entities in batches of 50 to extract relevant properties.

Several steps were taken to gather the correct IDs, as now detailed. Initially, it is essential to consider that some entities returned may not strictly represent a musical work during the textual search using the song and artist name. To address this issue, we filtered entities using the “**instance of (P31)**” property in Wikidata, selecting those classified as “**musical work (Q2188189)**” or subclasses like “**song (Q7366)**” or “**composition (Q204370)**”. This approach allows us to filter out entities not directly related to musical works.

From the original 699,123 songs, this filtering reduced the list to 111,238 entities, of which 66,373 (59.6%) lacked

⁴ <https://www.whosampled.com/about>

⁵ [wikidata.org](https://www.wikidata.org)

⁶ A discontinued ontology that is still used internally by Google Trends – [https://en.wikipedia.org/wiki/Freebase_\(database\)](https://en.wikipedia.org/wiki/Freebase_(database))

knowledge base identifiers. Focusing on entries with Freebase IDs, we finalized a set of 44,857 songs.

To ensure collected entities matched songs in our dataset, we refined them based on the normalized string similarity between the song and artist names, calculated separately with the `diffliib` library⁷ in Python. We retained only entries with a string distance above 0.55 for both names, reducing our dataset from 579,322 to 25,830 songs, representing 4,477 musical borrowings.

With the Freebase identifiers, we accessed Google Trends using each identifier as a URL parameter⁸. Our queries were configured to retrieve the global search interest over time for each song, starting from the inception of Google Trends in 2004. We successfully retrieved data for 4,360 songs, each represented as a monthly time series indicating relative popularity, normalized to peak at 100.

Estimating the causal effects of temporal data is challenging [10]. Initially, we can point out that several events may impact popularity (e.g., a song may be sampled twice). Furthermore, other seasonal effects (e.g., a good year for an artist) may be in place. To reduce this impact, all of our time series were filtered to have 24 points only, one per month. Twelve of these occurred before the release and Twelve after. Thus, we shall look at the impact on monthly interest. Moreover, we observed that some time series contained zero values in all time intervals, indicating the absence of search interest data for these musical entities throughout the analyzed period. After this last filter, we were left with 884 borrowings.

After this pre-processing step, we calculated the impact of the launch on the original in the short and medium term using RDD and Granger Causality.

4. EFFECT OF MUSICAL BORROWINGS ON INTEREST IN A SONG

To assess the impact of samples and re-recordings on audience interest in original music, we explored the hypothesis that these elements introduce discontinuity in nature. Considering that the original music already existed before and so did audience interest, we compared interest in the original music before and after the re-recording. This allows us to estimate the impact of the borrowing on audience interest in the original music.

4.1 Short-term impact

We used the Regression Discontinuity Design (RDD) method to measure the impact of a sample on the original music. RDD is a statistical method that allows us to estimate the causal effect of a treatment on an outcome variable when the probability of treatment changes discontinuously around a cutoff point. In this case, the treatment is the release date of the *borrowee* song. The outcome variable is the search interest in the original song, the *borrowed*. The cutoff point separates the observations into

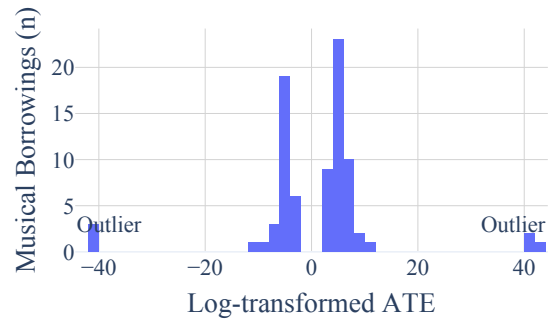


Figure 2: Distribution of the log-transformed average treatment effect (ATE) for the 82 statistically significant musical borrowings.

control and treatment groups. In detail, our RDD was as implemented as follows:

$$googletrends(t) \sim 1 + t + \mathbb{I}_{t>0} + t \cdot \mathbb{I}_{t>0} \quad (1)$$

In this context, t represents our time variable, shifted so that the release time of the *borrowee* occurs at time zero. This transformation is commonly used in RDD as it simplifies the interpretation of model parameters. The term $\mathbb{I}_{t>0}$ is an indicator function that equals one when $t > 0$ (after the release) and zero otherwise. Therefore, at $t = 0$, this term’s weight reflects the difference in intercepts, known as the Average Treatment Effect (ATE), which serves as a causal estimate of the release’s impact. We express this measurement in relative terms, where 100% indicates a doubling of the intercept.

Most importantly, we analyzed the ATE estimates only for regressions that achieved statistical significance at a level below 0.05. Our final dataset included 884 musical borrowings, of which 82 (approximately 9.3%) yielded statistically significant p -values. Figure 2 illustrates the ATE distribution for these 82 significant borrowings, which we will discuss next.

We initially point out borrowings with very high or very low ATEs from the figure. In some cases, this occurred because the original song had no prior search activity. However, after the sample’s release, the original song began to receive increased search interest. Figure 3a shows an example of this phenomenon. The song with the highest ATE in our set is “Something’s Got a Hold on Me” by Etta James, released in 1962, sampled in 2011 by American rapper Flo-Rida in his song “Good Feeling”. We manually inspected these extreme events; all were outliers like this one and we disregarded them.

We thus decided to focus on the non outlier cases, shown in the center of Figure 2. We begin with an interesting case of a negative ATE. Thus, Figure 3b presents the RDD analysis of Lana Del Rey’s “Summertime Sadness,” sampled in “Body Electric,” released five months later on the EP *Paraíso*. The ATE of -152% indicates a negative impact on interest in the original song following the sample release. We believe that the EP’s marketing and the similarity between the two tracks may have influenced the original song’s performance during that period.

⁷ <https://docs.python.org/3/library/diffliib.html>

⁸ https://trends.google.com/trends/explore?q=m/0zjw3z_

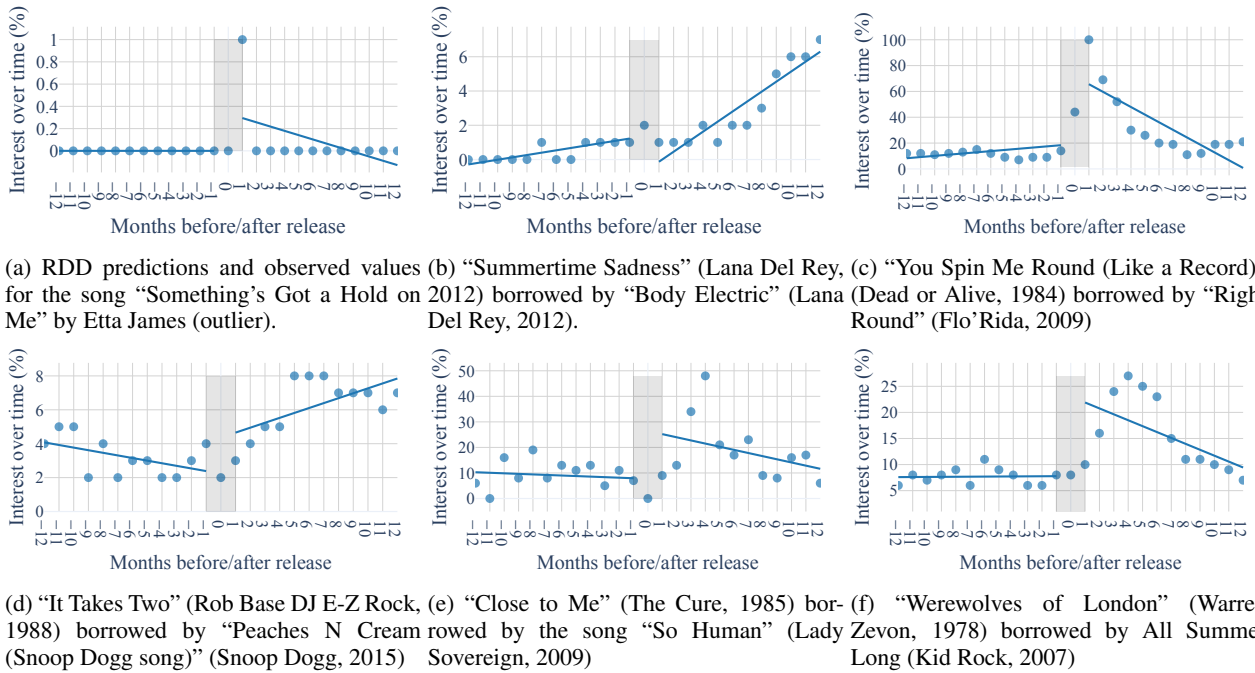


Figure 3: ATE Examples

Figures 3c, 3d, 3e and 3f all show examples of positive ATE. It is quite interesting that several of these examples show samples reviving interest in songs released in the late ’70s to late ’80s. This shows evidence that, in some cases, a *borrowee* may revive interest even in older songs. Nevertheless, we still need to check if this revival is long-lasting. Figures 3c and 3f show a sharp decay a few months after the *borrowee*’s release. This is why we complement this analysis with our Granger Causality analysis next.

4.2 Long-term impact

Granger causality is a statistical technique that investigates whether the past of one time series contains valuable information to predict the future of another. Thus, it considers the lags between observations.

We use the *statsmodels* [11] library to perform the Granger causality test, which offers robust statistical tools for time series analysis. The test involves formulating autoregressive (AR) models and comparing them to see if including lagged terms from one time series improves the prediction of another time series.

When applying the Granger test, it is necessary to specify the *max_lag* parameter, determining the maximum number of lags considered. We set the value of *max_lag* to 10, assuming that the causal relationship between the original song and the musical borrowing can be captured within 10 months (from one month to almost a year of maximum lag). The test is performed by combining two time series into a complete data frame and running the statistical procedure. The result of the Granger test is a *p*-value associated with the F-test of the sum of squares of the residuals (SSR F-test) for each lag up to the specified *max_lag*.

A low *p*-value (usually less than 0.05) indicates that the musical borrowing time series contains significant predic-

tive information about the original song time series. When this occurs, we have evidence of a Granger Causal relationship for the time series pair.

After the analysis, we observe that 64% (117) of the musical borrowings have a *p*-value below 0.05, establishing that, in most cases, the time series of the new version (cover, sample, or remix) contains significant predictive information about the time series of the original song. This suggests that *borrowees* generally influence *borrowers* over more extended periods of time.

5. CONCLUSIONS

This paper estimated causal evidence that musical borrowings impact the popularity of *borrowed* songs. Our motivation was delving deeper into the argument that borrowing (samples, remixes, and covers) increases the popularity of *borrowed* songs.

Our findings indicate some causal relationships between musical borrowings and the popularity of *borrowed* songs, though effects vary in strength. When present, a *borrowee* song can revitalize interest in older tracks, as shown by RDD analysis. Notably, sustained impacts assessed via Granger Causality suggest borrowings can have a lasting influence on original song popularity.

Our work has performed extensive care in collecting an accurate dataset that links several databases (Freebase, MusicBrainz, and Google Trends). Nevertheless, this does not come with some drawbacks. Out of approximately 700,000 songs analyzed, only 884 instances of musical borrowing were identified. Future work could refine data.

Finally, our results shed light on the complex economics of borrowings, with implications for fairer compensation practices [12].

6. REFERENCES

- [1] M. Schuster, D. Mitchell, and K. Brown, "Sampling increases music sales: An empirical copyright study," *American Business Law Journal*, vol. 56, no. 1, pp. 177–229, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ablj.12137>
- [2] J. Hahn, P. Todd, and W. Van der Klaauw, "Identification and estimation of treatment effects with a regression-discontinuity design," *Econometrica*, vol. 69, no. 1, pp. 201–209, 2001.
- [3] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969. [Online]. Available: <http://www.jstor.org/stable/1912791>
- [4] J. P. Burkholder, "The uses of existing music: Musical borrowing as a field," *Notes*, vol. 50, no. 3, pp. 851–870, 1994.
- [5] O. B. Arewa, "From jc bach to hip hop: Musical borrowing, copyright and cultural context," *NCL Rev.*, vol. 84, p. 547, 2005.
- [6] D. Hesmondhalgh, "Digital sampling and cultural inequality," *Social & legal studies*, vol. 15, no. 1, pp. 53–75, 2006.
- [7] U. Shalit, D. Weinshall, and G. Chechik, "Modeling musical influence with topic models," in *Proceedings of the 30th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, S. Dasgupta and D. McAllester, Eds., vol. 28, no. 2. Atlanta, Georgia, USA: PMLR, 17–19 Jun 2013, pp. 244–252. [Online]. Available: <https://proceedings.mlr.press/v28/shalit13.html>
- [8] N. J. Bryan and G. Wang, "Musical influence network analysis and rank of sample-based music." in *ISMIR*, 2011, pp. 329–334.
- [9] J. L. Ortega, "Cover versions as an impact indicator in popular music: A quantitative network analysis," *Plos one*, vol. 16, no. 4, p. e0250212, 2021.
- [10] C. Hausman and D. S. Rapson, "Regression discontinuity in time: Considerations for empirical applications," *Annual Review of Resource Economics*, vol. 10, no. 1, pp. 533–552, 2018.
- [11] S. Seabold and J. Perktold, "statsmodels: Econometric and statistical modeling with python," in *9th Python in Science Conference*, 2010.
- [12] S. Claflin, "How to get away with copyright infringement: Music sampling as fair use," *BUJ Sci. & Tech. L.*, vol. 26, p. 159, 2020.